

Body of Knowledge

	Knowledge Area	# Knowledge Units	CS Core Hours	KA Core Hours
AI	Artificial Intelligence	12	12	18
AL	Algorithmic Foundations	5	32	32
AR	Architecture and Organization	11	9	16
DM	Data Management	13	10	26
FPL	Foundations of Programming Languages	22	21	19
GIT	Graphics and Interactive Techniques	12	4	70
HCI	Human-Computer Interaction	6	8	16
MSF	Mathematical and Statistical Foundations	5	55	145
NC	Networking and Communication	8	7	24
OS	Operating Systems	14	8	14
PDC	Parallel and Distributed Computing	5	9	26
SDF	Software Development Fundamentals	5	43	
SE	Software Engineering	9	6	21
SEC	Security	7	6	35
SEP	Society, Ethics, and the Profession	11	18	14
SF	Systems Fundamentals	9	18	8
SPD	Specialized Platform Development	8	4	
	Total	162	270	N/A

Artificial Intelligence (AI)

Preamble

Artificial intelligence (AI) studies problems that are difficult or impractical to solve with traditional algorithmic approaches. These problems are often reminiscent of those considered to require human intelligence, and the resulting AI solution strategies typically generalize over classes of problems. AI techniques are now pervasive in computing, supporting everyday applications such as email, social media, photography, financial markets, and intelligent virtual assistants (e.g., Siri, Alexa). These techniques are also used in the design and analysis of autonomous agents that perceive their environment and interact rationally with it, such as self-driving vehicles and other robots.

Traditionally, AI has included a mix of symbolic and subsymbolic approaches. The solutions it provides rely on a broad set of general and specialized knowledge representation schemes, problem solving mechanisms, and optimization techniques. These approaches deal with perception (e.g., speech recognition, natural language understanding, computer vision), problem solving (e.g., search, planning, optimization), generation (e.g., narrative, conversation, images, models, recommendations), acting (e.g., robotics, task-automation, control), and the architectures needed to support them (e.g., single agents, multi-agent systems). Machine learning may be used within each of these aspects, and can even be employed end-to-end across all of them. The study of Artificial Intelligence prepares students to determine when an AI approach is appropriate for a given problem, identify appropriate representations and reasoning mechanisms, implement them, and evaluate them with respect to both performance and their broader societal impact.

Over the past decade, the term “artificial intelligence” has become commonplace within businesses, news articles, and everyday conversation, driven largely by a series of high-impact machine learning applications. These advances were made possible by the widespread availability of large datasets, increased computational power, and algorithmic improvements. In particular, there has been a shift from engineered representations to representations learned automatically through optimization over large datasets. The resulting advances have put such terms as “neural networks” and “deep learning” into everyday vernacular. Businesses now advertise AI-based solutions as value-additions to their services, so that “artificial intelligence” is now both a technical term and a marketing buzzword. Other disciplines, such as biology, art, architecture, and finance, increasingly use AI techniques to solve problems within their disciplines.

For the first time in our history, the broader population has access to sophisticated AI-driven tools, including tools to generate essays or poems from a prompt, artwork from a description, and fake photographs or videos depicting real people. AI technology is now in widespread use in stock trading, curating our news and social media feeds, automated evaluation of job applicants, detection of medical conditions, and influencing prison sentencing through recidivism prediction. Consequently, AI technology can have significant societal impacts and ethical considerations that must be understood and considered when developing and applying it.

Changes since CS2013

To reflect this recent growth and societal impact, the knowledge area has been revised from CS2013 in the following ways:

- The name has changed from “Intelligent Systems” to “Artificial Intelligence,” to reflect the most common terminology used for these topics within the field and its more widespread use outside the field.
- An increased emphasis on neural networks and representation learning reflects the recent advances in the field. Given its key role throughout AI, search is still emphasized but there is a slight reduction on symbolic methods in favor of understanding subsymbolic methods and learned representations. It is important, however, to retain knowledge-based and symbolic approaches within the AI curriculum because these methods offer unique capabilities, are used in practice, ensure a broad education, and because more recent neurosymbolic approaches integrate both learned and symbolic representations.
- There is an increased emphasis on practical applications of AI, including a variety of areas (e.g., medicine, sustainability, social media, etc.). This includes explicit discussion of tools that employ deep generative models (e.g., ChatGPT, DALL-E, Midjourney) and are now in widespread use, covering how they work at a high level, their uses, and their shortcomings/pitfalls.
- The curriculum reflects the importance of understanding and assessing the broader societal impacts and implications of AI methods and applications, including issues in AI ethics, fairness, trust, and explainability.
- The AI knowledge area includes connections to data science through 1) cross-connections with the Data Management and other knowledge areas and 2) a sample Data Science model course.
- There are explicit goals to develop basic AI literacy and critical thinking in every computer science student, given the breadth of interconnections between AI and other knowledge areas in practice.

Consider recent AI advances when using this curriculum

The field of AI is undergoing rapid development and increasingly widespread applications. Since the first draft of this document, several new techniques (e.g., generative networks, large language models) have become widely used and so were added to the CS or KA Cores. This document is as current as we can make it in 2023. However, we expect such rapid changes to continue in the subfield of AI during the expected life of this document. Consequently, it is imperative that faculty teaching AI understand current advances and consider whether these advances should be taught in order to keep the curriculum current.

Core Hours

Knowledge Units	CS Core	KA Core
Fundamental Issues	2	1
Search	2 + 3 (AL)	6
Fundamental Knowledge Representation and Reasoning	1 + 1 (MSF)	2
Machine Learning	4	6
Applications and Societal Impact	3	3
Probabilistic Representation and Reasoning		
Planning		
Logical Representation and Reasoning		
Agents and Cognitive Systems		
Natural Language Processing		
Robotics		
Perception and Computer Vision		
Total	12 (16)	18

Note: The CS Core includes 3 hours that are shared with and counted under [Algorithm Foundations](#) (Uninformed search) and 1 hour that is shared with and counted under [Mathematical Foundations](#) (Probability). The [AI](#) KA contributes 12 hours in total toward the complete CS Core; 16 hours if you include those hours counted under other KAs.

Knowledge Units

AI-Introduction: Fundamental Issues

CS Core:

1. Overview of AI problems, Examples of successful recent AI applications
2. Definitions of agents with examples (e.g., reactive, deliberative)
3. What is intelligent behavior?
 - a. The Turing test and its flaws
 - b. Multimodal input and output
 - c. Simulation of intelligent behavior
 - d. Rational versus non-rational reasoning

4. Problem characteristics
 - a. Fully versus partially observable
 - b. Single versus multi-agent
 - c. Deterministic versus stochastic
 - d. Static versus dynamic
 - e. Discrete versus continuous
5. Nature of agents
 - a. Autonomous, semi-autonomous, mixed-initiative autonomy
 - b. Reflexive, goal-based, and utility-based
 - c. Decision making under uncertainty and with incomplete information
 - d. The importance of perception and environmental interactions
 - e. Learning-based agents
 - f. Embodied agents
 - i. sensors, dynamics, effectors
6. Overview of AI Applications, growth, and impact (economic, societal, ethics)

KA Core:

7. Practice identifying problem characteristics in example environments
8. Additional depth on nature of agents with examples
9. Additional depth on AI Applications, growth, and Impact (economic, societal, ethics, security)

Non-core:

10. Philosophical issues
11. History of AI

Illustrative Learning Outcomes:

1. Describe the Turing test and the “Chinese Room” thought experiment.
2. Differentiate between optimal reasoning/behavior and human-like reasoning/behavior.
3. Differentiate the terms: AI, machine learning, and deep learning.
4. Enumerate the characteristics of a specific problem.

AI-Search: Search

CS Core:

1. State space representation of a problem
 - a. Specifying states, goals, and operators
 - b. Factoring states into representations (hypothesis spaces)
 - c. Problem solving by graph search
 - i. e.g., Graphs as a space, and tree traversals as exploration of that space
 - ii. Dynamic construction of the graph (you’re not given it upfront)
2. Uninformed graph search for problem solving (See also: [AL-Foundational](#))
 - a. Breadth-first search
 - b. Depth-first search
 - i. With iterative deepening

- c. Uniform cost search
- 3. Heuristic graph search for problem solving (See also: [AL-Strategies](#))
 - a. Heuristic construction and admissibility
 - b. Hill-climbing
 - c. Local minima and the search landscape
 - i. Local vs global solutions
 - d. Greedy best-first search
 - e. A* search
- 4. Space and time complexities of graph search algorithms

KA Core:

- 5. Bidirectional search
- 6. Beam search
- 7. Two-player adversarial games
 - a. Minimax search
 - b. Alpha-beta pruning
 - i. Ply cutoff
- 8. Implementation of A* search
- 9. Constraint satisfaction

Non-core:

- 10. Understanding the search space
 - a. Constructing search trees
 - b. Dynamic search spaces
 - c. Combinatorial explosion of search space
 - d. Search space topology (ridges, saddle points, local minima, etc.)
- 11. Local search
- 12. Tabu search
- 13. Variations on A* (IDA*, SMA*, RBFS)
- 14. Two-player adversarial games
 - a. The horizon effect
 - b. Opening playbooks / endgame solutions
 - c. What it means to “solve” a game (e.g., checkers)
- 15. Implementation of minimax search, beam search
- 16. Expectimax search (MDP-solving) and chance nodes
- 17. Stochastic search
 - a. Simulated annealing
 - b. Genetic algorithms
 - c. Monte-Carlo tree search

Illustrative Learning Outcomes:

- 1. Design the state space representation for a puzzle (e.g., N-queens or 3-jug problem)
- 2. Select and implement an appropriate uninformed search algorithm for a problem (e.g., tic-tac-toe), and characterize its time and space complexities.

3. Select and implement an appropriate informed search algorithm for a problem after designing a helpful heuristic function (e.g., a robot navigating a 2D gridworld).
4. Evaluate whether a heuristic for a given problem is admissible/can guarantee an optimal solution.
5. Apply minimax search in a two-player adversarial game (e.g., connect four), using heuristic evaluation at a particular depth to compute the scores to back up. [KA Core]
6. Design and implement a genetic algorithm solution to a problem.
7. Design and implement a simulated annealing schedule to avoid local minima in a problem.
8. Design and implement A*/beam search to solve a problem, and compare it against other search algorithms in terms of the solution cost, number of nodes expanded, etc.
9. Apply minimax search with alpha-beta pruning to prune search space in a two-player adversarial game (e.g., connect four).
10. Compare and contrast genetic algorithms with classic search techniques, explaining when it is most appropriate to use a genetic algorithm to learn a model versus other forms of optimization (e.g., gradient descent).
11. Compare and contrast various heuristic searches vis-a-vis applicability to a given problem.
12. Model a logic or Sudoku puzzle as a constraint satisfaction problem, solve it with backtrack search, and determine how much arc consistency can reduce the search space.

AI-KRR: Fundamental Knowledge Representation and Reasoning

CS Core:

1. Types of representations
 - a. Symbolic, logical
 - i. Creating a representation from a natural language problem statement
 - b. Learned subsymbolic representations
 - c. Graphical models (e.g., naive Bayes, Bayesian network)
2. Review of probabilistic reasoning, Bayes theorem (See also: [MSF-Probability](#))
3. Bayesian reasoning
 - a. Bayesian inference

KA Core:

4. Random variables and probability distributions
 - a. Axioms of probability
 - b. Probabilistic inference
 - c. Bayes' Rule (derivation)
 - d. Bayesian inference (more complex examples)
5. Independence
6. Conditional Independence
7. Markov chains and Markov models
8. Utility and decision making

Illustrative Learning Outcomes:

1. Given a natural language problem statement, encode it as a symbolic or logical representation.
2. Explain how we can make decisions under uncertainty, using concepts such as Bayes theorem and utility.

3. Compute a probabilistic inference in a real-world problem using Bayes' theorem to determine the probability of a hypothesis given evidence.
4. Apply Bayes' rule to determine the probability of a hypothesis given evidence.
5. Compute the probability of outcomes and test whether outcomes are independent.

AI-ML: Machine Learning

CS Core:

1. Definition and examples of a broad variety of machine learning tasks
 - a. Supervised learning
 - i. Classification
 - ii. Regression
 - b. Reinforcement learning
 - c. Unsupervised learning
 - i. Clustering
2. Fundamental ideas:
 - a. No free lunch theorem: no one learner can solve all problems; representational design decisions have consequences
 - b. Sources of error and undecidability in machine learning
3. A simple statistical-based supervised learning such as linear regression or decision trees
 - a. Focus on how they work without going into mathematical or optimization details; enough to understand and use existing implementations correctly
4. The overfitting problem / controlling solution complexity (regularization, pruning – intuition only)
 - a. The bias (underfitting) - variance (overfitting) tradeoff
5. Working with Data
 - a. Data preprocessing
 - i. Importance and pitfalls of preprocessing choices
 - b. Handling missing values (imputing, flag-as-missing)
 - i. Implications of imputing vs flag-as-missing
 - c. Encoding categorical variables, encoding real-valued data
 - d. Normalization/standardization
 - e. Emphasis on real data, not textbook examples
6. Representations
 - a. Hypothesis spaces and complexity
 - b. Simple basis feature expansion, such as squaring univariate features
 - c. Learned feature representations
7. Machine learning evaluation
 - a. Separation of train, validation, and test sets
 - b. Performance metrics for classifiers
 - c. Estimation of test performance on held-out data
 - d. Tuning the parameters of a machine learning model with a validation set
 - e. Importance of understanding what your model is actually doing, where its pitfalls/shortcomings are, and the implications of its decisions
8. Basic neural networks

- a. Fundamentals of understanding how neural networks work and their training process, without details of the calculations
- b. Basic introduction to generative neural networks (large language models, etc.)
- 9. Ethics for Machine Learning (See also: [SEP-Context](#))
 - a. Focus on real data, real scenarios, and case studies
 - b. Dataset/algorithmic/evaluation bias and unintended consequences

KA Core:

- 10. Formulation of simple machine learning as an optimization problem, such as least squares linear regression or logistic regression
 - a. Objective function
 - b. Gradient descent
 - c. Regularization to avoid overfitting (mathematical formulation)
- 11. Ensembles of models
 - a. Simple weighted majority combination
- 12. Deep learning
 - a. Deep feed-forward networks (intuition only, no mathematics)
 - b. Convolutional neural networks (intuition only, no mathematics)
 - c. Visualization of learned feature representations from deep nets
 - d. Other architectures (generative NN, recurrent NN, transformers, etc.)
- 13. Performance evaluation
 - a. Other metrics for classification (e.g., error, precision, recall)
 - b. Performance metrics for regressors
 - c. Confusion matrix
 - d. Cross-validation
 - i. Parameter tuning (grid/random search, via cross-validation)
- 14. Overview of reinforcement learning methods
- 15. Two or more applications of machine learning algorithms
 - a. E.g., medicine and health, economics, vision, natural language, robotics, game play
- 16. Ethics for Machine Learning
 - a. Continued focus on real data, real scenarios, and case studies (See also: [SEP-Context](#))
 - b. Privacy (See also: [SEP-Privacy](#))
 - c. Fairness (See also: [SEP-Privacy](#))
 - d. Intellectual property
 - e. Explainability

Non-core:

- 17. General statistical-based learning, parameter estimation (maximum likelihood)
- 18. Supervised learning
 - a. Decision trees
 - b. Nearest-neighbor classification and regression
 - c. Learning simple neural networks / multi-layer perceptrons
 - d. Linear regression
 - e. Logistic regression

- f. Support vector machines (SVMs) and kernels
- g. Gaussian Processes
- 19. Overfitting
 - a. The curse of dimensionality
 - b. Regularization (mathematical computations, L_2 and L_1 regularization)
- 20. Experimental design
 - a. Data preparation (e.g., standardization, representation, one-hot encoding)
 - b. Hypothesis space
 - c. Biases (e.g., algorithmic, search)
 - d. Partitioning data: stratification, training set, validation set, test set
 - e. Parameter tuning (grid/random search, via cross-validation)
 - f. Performance evaluation
 - i. Cross-validation
 - ii. Metric: error, precision, recall, confusion matrix
 - iii. Receiver operating characteristic (ROC) curve and area under ROC curve
- 21. Bayesian learning (Cross-Reference AI/Reasoning Under Uncertainty)
 - a. Naive Bayes and its relationship to linear models
 - b. Bayesian networks
 - c. Prior/posterior
 - d. Generative models
- 22. Deep learning
 - a. Deep feed-forward networks
 - b. Neural tangent kernel and understanding neural network training
 - c. Convolutional neural networks
 - d. Autoencoders
 - e. Recurrent networks
 - f. Representations and knowledge transfer
 - g. Adversarial training and generative adversarial networks
 - h. Attention mechanisms
- 23. Representations
 - a. Manually crafted representations
 - b. Basis expansion
 - c. Learned representations (e.g., deep neural networks)
- 24. Unsupervised learning and clustering
 - a. K-means
 - b. Gaussian mixture models
 - c. Expectation maximization (EM)
 - d. Self-organizing maps
- 25. Graph analysis (e.g., PageRank)
- 26. Semi-supervised learning
- 27. Graphical models (See also: [AI-Probability](#))
- 28. Ensembles
 - a. Weighted majority
 - b. Boosting/bagging

- c. Random forest
 - d. Gated ensemble
29. Learning theory
 - a. General overview of learning theory / why learning works
 - b. VC dimension
 - c. Generalization bounds
 30. Reinforcement learning
 - a. Exploration vs exploitation trade-off
 - b. Markov decision processes
 - c. Value and policy iteration
 - d. Policy gradient methods
 - e. Deep reinforcement learning
 - f. Learning from demonstration and inverse RL
 31. Explainable / interpretable machine learning
 - a. Understanding feature importance (e.g., LIME, Shapley values)
 - b. Interpretable models and representations
 32. Recommender systems
 33. Hardware for machine learning
 - a. GPUs / TPUs
 34. Application of machine learning algorithms to:
 - a. Medicine and health
 - b. Economics
 - c. Education
 - d. Vision
 - e. Natural language
 - f. Robotics
 - g. Game play
 - h. Data mining (Cross-reference DM/Data Analytics)
 35. Ethics for Machine Learning
 - a. Continued focus on real data, real scenarios, and case studies (See also: [SEP-Context](#))
 - b. In depth exploration of dataset/algorithmic/evaluation bias, data privacy, and fairness (See also: [SEP-Privacy](#), [SEP-Context](#))
 - c. Trust / explainability

Illustrative Learning Outcomes:

1. Describe the differences among the three main styles of learning (supervised, reinforcement, and unsupervised) and determine which is appropriate to a particular problem domain.
2. Differentiate the terms of AI, machine learning, and deep learning.
3. Frame an application as a classification problem, including the available input features and output to be predicted (e.g., identifying alphabetic characters from pixel grid input).
4. Apply two or more simple statistical learning algorithms to a classification task and measure the classifiers' accuracy.
5. Identify overfitting in the context of a problem and learning curves and describe solutions to overfitting.

6. Explain how machine learning works as an optimization/search process.
7. Implement a statistical learning algorithm and the corresponding optimization process to train the classifier and obtain a prediction on new data.
8. Describe the neural network training process and resulting learned representations
9. Explain proper ML evaluation procedures, including the differences between training and testing performance, and what can go wrong with the evaluation process leading to inaccurate reporting of ML performance.
10. Compare two machine learning algorithms on a dataset, implementing the data preprocessing and evaluation methodology (e.g., metrics and handling of train/test splits) from scratch.
11. Visualize the training progress of a neural network through learning curves in a well-established toolkit (e.g., TensorBoard) and visualize the learned features of the network.
12. Compare and contrast several learning techniques (e.g., decision trees, logistic regression, naive Bayes, neural networks, and belief networks), providing examples of when each strategy is superior.
13. Evaluate the performance of a simple learning system on a real-world dataset.
14. Characterize the state of the art in learning theory, including its achievements and shortcomings.
15. Explain the problem of overfitting, along with techniques for detecting and managing the problem.
16. Explain the triple tradeoff among the size of a hypothesis space, the size of the training set, and performance accuracy.
17. Given a real-world application of machine learning, describe ethical issues regarding the choices of data, preprocessing steps, algorithm selection, and visualization/presentation of results.

AI-SEP: Applications and Societal Impact

Note: There is substantial benefit to studying applications and ethics/fairness/trust/explainability in a curriculum alongside the methods and theory that they apply to, rather than covering ethics in a separate, dedicated class session. Whenever possible, study of these topics should be integrated alongside other modules, such as exploring how decision trees could be applied to a specific problem in environmental sustainability such as land use allocation, then assessing the social/environmental/ethical implications of doing so.

CS Core:

1. At least one application of AI to a specific problem and field, such as medicine, health, sustainability, social media, economics, education, robotics, etc. (choose at least one for the CS Core).
 - a. Formulating and evaluating a specific application as an AI problem
 - i. How to deal with underspecified or ill-posed problems
 - b. Data availability/scarcity and cleanliness
 - i. Basic data cleaning and preprocessing
 - ii. Data set bias
 - c. Algorithmic bias
 - d. Evaluation bias
 - e. Assessment of societal implications of the application
2. Deployed deep generative models

- a. High-level overview of deep image generative models (e.g. as of 2023, DALL-E, Midjourney, Stable Diffusion, etc.), their uses, and their shortcomings/pitfalls.
- b. High-level overview of large language models (e.g. as of 2023, ChatGPT, Bard, etc.), their uses, and their shortcomings/pitfalls.
- 3. Overview of societal impact of AI
 - a. Ethics (See also: [SEP-Context](#))
 - b. Fairness (See also: [SEP-Privacy](#), [SEP-DEIA](#))
 - c. Trust / explainability (See also: [SEP-Context](#))
 - d. Privacy and usage of training data (See also: [SEP-Privacy](#))
 - e. Human autonomy and oversight/regulations/legal requirements (See also: [SEP-Context](#))
 - f. Sustainability (See also: [SEP-Sustainability](#))

KA Core:

- 4. One or more additional applications of AI to a broad set of problems and diverse fields, such as medicine, health, sustainability, social media, economics, education, robotics, etc. (choose a different area from that chosen for the CS Core).
 - a. Formulating and evaluating a specific application as an AI problem
 - i. How to deal with underspecified or ill-posed problems
 - b. Data availability/scarcity and cleanliness
 - i. Basic data cleaning and preprocessing
 - ii. Data set bias
 - c. Algorithmic bias
 - d. Evaluation bias
 - e. Assessment of societal implications of the application
- 5. Additional depth on deployed deep generative models
 - a. Introduction to how deep image generative models (e.g. as of 2023, DALL-E, Midjourney, Stable Diffusion, etc.) work, including discussion of attention
 - b. Introduction to how large language models (e.g. as of 2023, ChatGPT, Bard, etc.) work, including discussion of attention
 - c. Idea of foundational models, how to use them, and the benefits / issues with training them from big data
- 6. Analysis and discussion of the societal impact of AI
 - a. Ethics (See also: [SEP-Context](#))
 - b. Fairness (See also: [SEP-Privacy](#), [SEP-DEIA](#))
 - c. Trust / explainability (See also: [SEP-Context](#))
 - d. Privacy and usage of training data (See also: [SEP-Privacy](#))
 - e. Human autonomy and oversight/regulations/legal requirements (See also: [SEP-Context](#))
 - f. Sustainability (See also: [SEP-Sustainability](#))

Illustrative Learning Outcomes:

- 1. Given a real-world application domain and problem, formulate an AI solution to it, identifying proper data/input, preprocessing, representations, AI techniques, and evaluation metrics/methodology.
- 2. Analyze the societal impact of one or more specific real-world AI applications, identifying issues regarding ethics, fairness, bias, trust, and explainability.

3. Describe some of the failure modes of current deep generative models for language or images, and how this could affect their use in an application.

AI-LRR: Logical Representation and Reasoning

Non-core:

1. Review of propositional and predicate logic (See also: [MSF-Discrete](#))
2. Resolution and theorem proving (propositional logic only)
 - a. Forward chaining, backward chaining
3. Knowledge representation issues
 - a. Description logics
 - b. Ontology engineering
4. Semantic web
5. Non-monotonic reasoning (e.g., non-classical logics, default reasoning)
6. Argumentation
7. Reasoning about action and change (e.g., situation and event calculus)
8. Temporal and spatial reasoning
9. Logic programming
 - a. Prolog, Answer Set Programming
10. Rule-based Expert Systems
11. Semantic networks
12. Model-based and Case-based reasoning

Illustrative Learning Outcomes:

1. Translate a natural language (e.g., English) sentence into a predicate logic statement.
2. Convert a logic statement into clausal form.
3. Apply resolution to a set of logic statements to answer a query.
4. Compare and contrast the most common models used for structured knowledge representation, highlighting their strengths and weaknesses.
5. Identify the components of non-monotonic reasoning and its usefulness as a representational mechanism for belief systems.
6. Compare and contrast the basic techniques for representing uncertainty.
7. Compare and contrast the basic techniques for qualitative representation.
8. Apply situation and event calculus to problems of action and change.
9. Explain the distinction between temporal and spatial reasoning, and how they interrelate.
10. Explain the difference between rule-based, case-based and model-based reasoning techniques.
11. Define the concept of a planning system and how it differs from classical search techniques.
12. Describe the differences between planning as search, operator-based planning, and propositional planning, providing examples of domains where each is most applicable.
13. Explain the distinction between monotonic and non-monotonic inference.

AI-Probability: Probabilistic Representation and Reasoning

Non-core:

1. Conditional Independence review
2. Knowledge representations
 - a. Bayesian Networks
 - i. Exact inference and its complexity
 - ii. Markov blankets and d-separation
 - iii. Randomized sampling (Monte Carlo) methods (e.g. Gibbs sampling)
 - b. Markov Networks
 - c. Relational probability models
 - d. Hidden Markov Models
3. Decision Theory
 - a. Preferences and utility functions
 - b. Maximizing expected utility
 - c. Game theory

Illustrative Learning Outcomes:

1. Compute the probability of a hypothesis given the evidence in a Bayesian network.
2. Explain how conditional independence assertions allow for greater efficiency of probabilistic systems.
3. Identify examples of knowledge representations for reasoning under uncertainty.
4. State the complexity of exact inference. Identify methods for approximate inference.
5. Design and implement at least one knowledge representation for reasoning under uncertainty.
6. Describe the complexities of temporal probabilistic reasoning.
7. Design and implement an HMM as one example of a temporal probabilistic system.
8. Describe the relationship between preferences and utility functions.
9. Explain how utility functions and probabilistic reasoning can be combined to make rational decisions.

AI-Planning: Planning

Non-core:

1. Review of propositional and first-order logic
2. Planning operators and state representations
3. Total order planning
4. Partial-order planning
5. Plan graphs and GraphPlan
6. Hierarchical planning
7. Planning languages and representations
 - a. PDDL
8. Multi-agent planning
9. MDP-based planning
10. Interconnecting planning, execution, and dynamic replanning
 - a. Conditional planning
 - b. Continuous planning
 - c. Probabilistic planning

Illustrative Learning Outcomes:

1. Construct the state representation, goal, and operators for a given planning problem.
2. Encode a planning problem in PDDL and use a planner to solve it.
3. Given a set of operators, initial state, and goal state, draw the partial-order planning graph and include ordering constraints to resolve all conflicts
4. Construct the complete planning graph for GraphPlan to solve a given problem.

AI-Agents: Agents and Cognitive Systems**Non-core:**

1. Agent architectures (e.g., reactive, layered, cognitive)
2. Agent theory (including mathematical formalisms)
3. Rationality, Game Theory
 - a. Decision-theoretic agents
 - b. Markov decision processes (MDP)
 - c. Bandit algorithms
4. Software agents, personal assistants, and information access
 - a. Collaborative agents
 - b. Information-gathering agents
 - c. Believable agents (synthetic characters, modeling emotions in agents)
5. Learning agents
6. Cognitive systems
 - a. Cognitive architectures (e.g., ACT-R, SOAR, ICARUS, FORR)
 - b. Capabilities (perception, decision making, prediction, knowledge maintenance, etc.)
 - c. Knowledge representation, organization, utilization, acquisition, and refinement
 - d. Applications and evaluation of cognitive systems
7. Multi-agent systems
 - a. Collaborating agents
 - b. Agent teams
 - c. Competitive agents (e.g., auctions, voting)
 - d. Swarm systems and biologically inspired models
 - e. Multi-agent learning
8. Human-agent interaction (See also: [HCI-User](#), [HCI-Accessibility](#))
 - a. Communication methodologies (verbal and non-verbal)
 - b. Practical issues
 - c. Applications
 - i. Trading agents, supply chain management
 - ii. Ethical issues of AI interactions with humans
 - iii. Regulation and legal requirements of AI systems for interacting with humans

Illustrative Learning Outcomes:

1. Characterize and contrast the standard agent architectures.
2. Describe the applications of agent theory to domains such as software agents, personal assistants, and believable agents, and discuss associated ethical implications.
3. Describe the primary paradigms used by learning agents.

4. Demonstrate using appropriate examples how multi-agent systems support agent interaction.
5. Construct an intelligent agent using a well-established cognitive architecture (ACT-R, SOAR) for solving a specific problem.

AI-NLP: Natural Language Processing

Non-core:

1. Deterministic and stochastic grammars
2. Parsing algorithms
 - a. CFGs and chart parsers (e.g. CYK)
 - b. Probabilistic CFGs and weighted CYK
3. Representing meaning / Semantics
 - a. Logic-based knowledge representations
 - b. Semantic roles
 - c. Temporal representations
 - d. Beliefs, desires, and intentions
4. Corpus-based methods
5. N-grams and HMMs
6. Smoothing and backoff
7. Examples of use: POS tagging and morphology
8. Information retrieval (See also: [DM-Unstructured](#))
 - a. Vector space model
 - i. TF & IDF
 - b. Precision and recall
9. Information extraction
10. Language translation
11. Text classification, categorization
 - a. Bag of words model
12. Deep learning for NLP (See also: [AI-ML](#))
 - a. RNNs
 - b. Transformers
 - c. Multi-modal embeddings (e.g., images + text)
 - d. Generative language models

Illustrative Learning Outcomes:

1. Define and contrast deterministic and stochastic grammars, providing examples to show the adequacy of each.
2. Simulate, apply, or implement classic and stochastic algorithms for parsing natural language.
3. Identify the challenges of representing meaning.
4. List the advantages of using standard corpora. Identify examples of current corpora for a variety of NLP tasks.
5. Identify techniques for information retrieval, language translation, and text classification.
6. Implement a TF/IDF transform, use it to extract features from a corpus, and train an off-the-shelf machine learning algorithm using those features to do text classification.

AI-Robotics: Robotics

(See also: [SPD-Robot](#))

Non-core:

1. Overview: problems and progress
 - a. State-of-the-art robot systems, including their sensors and an overview of their sensor processing
 - b. Robot control architectures, e.g., deliberative vs reactive control and Braitenberg vehicles
 - c. World modeling and world models
 - d. Inherent uncertainty in sensing and in control
2. Sensors and effectors
 - a. Sensors: LIDAR, sonar, vision, depth, stereoscopic, event cameras, microphones, haptics, etc.
 - b. Effectors: wheels, arms, grippers, etc.
3. Coordinate frames, translation, and rotation (2D and 3D)
4. Configuration space and environmental maps
5. Interpreting uncertain sensor data
6. Localization and mapping
7. Navigation and control
8. Forward and inverse kinematics
9. Motion path planning and trajectory optimization
10. Manipulation and grasping
11. Joint control and dynamics
12. Vision-based control
13. Multiple-robot coordination and collaboration
14. Human-robot interaction (See also: [HCI-User](#), [HCI-Accessibility](#))
 - a. Shared workspaces
 - b. Human-robot teaming and physical HRI
 - c. Social assistive robots
 - d. Motion/task/goal prediction
 - e. Collaboration and communication (explicit vs implicit, verbal or symbolic vs non-verbal or visual)
 - f. Trust
15. Applications and Societal, Economic, and Ethical Issues
 - a. Societal, economic, right-to-work implications
 - b. Ethical and privacy implications of robotic applications
 - c. Liability in autonomous robotics
 - d. Autonomous weapons and ethics
 - e. Human oversight and control

Illustrative Learning Outcomes:

(Note: Due to the expense of robot hardware, all of these could be done in simulation or with low-cost educational robotic platforms.)

1. List capabilities and limitations of today's state-of-the-art robot systems, including their sensors and the crucial sensor processing that informs those systems.
2. Integrate sensors, actuators, and software into a robot designed to undertake some task.

3. Program a robot to accomplish simple tasks using deliberative, reactive, and/or hybrid control architectures.
4. Implement fundamental motion planning algorithms within a robot configuration space.
5. Characterize the uncertainties associated with common robot sensors and actuators; articulate strategies for mitigating these uncertainties.
6. List the differences among robots' representations of their external environment, including their strengths and shortcomings.
7. Compare and contrast at least three strategies for robot navigation within known and/or unknown environments, including their strengths and shortcomings.
8. Describe at least one approach for coordinating the actions and sensing of several robots to accomplish a single task.
9. Compare and contrast a multi-robot coordination and a human-robot collaboration approach, and attribute their differences to differences between the problem settings.
10. Analyze the societal, economic, and ethical issues of a real-world robotics application

AI-Vision: Perception and Computer Vision

Non-core:

1. Computer vision
 - a. Image acquisition, representation, processing and properties
 - b. Shape representation, object recognition, and segmentation
 - c. Motion analysis
 - d. Generative models
2. Audio and speech recognition
3. Touch and proprioception
4. Other modalities (e.g., olfaction)
5. Modularity in recognition
6. Approaches to pattern recognition. (See also: [AI-ML](#))
 - a. Classification algorithms and measures of classification quality
 - b. Statistical techniques
 - c. Deep learning techniques

Illustrative Learning Outcomes:

1. Summarize the importance of image and object recognition in AI and indicate several significant applications of this technology.
2. List at least three image-segmentation approaches, such as thresholding, edge-based and region-based algorithms, along with their defining characteristics, strengths, and weaknesses.
3. Implement 2d object recognition based on contour- and/or region-based shape representations.
4. Distinguish the goals of sound-recognition, speech-recognition, and speaker-recognition and identify how the raw audio signal will be handled differently in each of these cases.
5. Provide at least two examples of a transformation of a data source from one sensory domain to another, e.g., tactile data interpreted as single-band 2d images.
6. Implement a feature-extraction algorithm on real data, e.g., an edge or corner detector for images or vectors of Fourier coefficients describing a short slice of audio signal.

7. Implement an algorithm combining features into higher-level percepts, e.g., a contour or polygon from visual primitives or phoneme hypotheses from an audio signal.
8. Implement a classification algorithm that segments input percepts into output categories and quantitatively evaluates the resulting classification.
9. Evaluate the performance of the underlying feature-extraction, relative to at least one alternative possible approach (whether implemented or not) in its contribution to the classification task (8), above.
10. Describe at least three classification approaches, their prerequisites for applicability, their strengths, and their shortcomings.
11. Implement and evaluate a deep learning solution to problems in computer vision, such as object or scene recognition.

Professional Dispositions

- **Meticulousness:** Attention must be paid to details when implementing AI and machine learning algorithms, requiring students to be meticulous to detail.
- **Persistence:** AI techniques often operate in partially observable environments and optimization processes may have cascading errors from multiple iterations. Getting AI techniques to work predictably takes trial and error, and repeated effort. These call for persistence on the part of the student.
- **Inventive:** Applications of AI involve creative problem formulation and application of AI techniques, while balancing application requirements and societal and ethical issues.
- **Responsible:** Applications of AI can have significant impacts on society, affecting both individuals and large populations. This calls for students to understand the implications of work in AI to society, and to make responsible choices for when and how to apply AI techniques.

Mathematics Requirements

Required:

- Algebra
- Precalculus
- Discrete Math: (See also: [MSF-Discrete](#))
 - sets, relations, functions, graphs
 - predicate and first-order logic, logic-based proofs
- Linear Algebra: (See also: [MSF-Linear](#))
 - Matrix operations, matrix algebra
 - Basis sets
- Probability and Statistics: (See also: [MSF-Statistics](#))
 - Basic probability theory, conditional probability, independence
 - Bayes theorem and applications of Bayes theorem
 - Expected value, basic descriptive statistics, distributions
 - Basic summary statistics and significance testing

- All should be applied to real decision making examples with real data, not “textbook” examples

Desirable:

- Calculus-based probability and statistics
- Calculus: single-variable and partial derivatives
- Other topics in probability and statistics
 - Hypothesis testing, data resampling, experimental design techniques
- Optimization
- Linear algebra (all other topics)

Course Packaging Suggestions

Artificial Intelligence to include the following:

- [AI-Introduction](#) (4 hours)
- [AI-Search](#) (9 hours)
- [AI-KRR](#) (4 hours)
- [AI-ML](#) (12 hours)
- [AI-Probability](#) (5 hours)
- [AI-SEP](#) (4 hours - should be integrated throughout the course)

Prerequisites:

- [SDF-Fundamentals](#)
- [SDF-Data-Structures](#)
- [SDF-Algorithms](#)
- [MSF-Discrete](#)
- [MSF-Probability](#)

Course objective: A student who completes this course should understand the basic areas of AI and be able to understand, develop, and apply techniques in each. They should be able to solve problems using search techniques, basic Bayesian reasoning, and simple machine learning methods. They should understand the various applications of AI and associated ethical and societal implications.

Machine Learning to include the following:

- [AI-ML](#) (32 hours)
- [AI-KRR](#) (4 hours)
- [AI-NLP](#) (4 hours - selected topics, e.g., TF-IDF, bag of words, and text classification)
- [AI-SEP](#) (4 hours - should be integrated throughout the course)

Prerequisites:

- [SDF-Fundamentals](#)
- [SDF-Data-Structures](#)
- [SDF-Algorithms](#)
- [MSF-Discrete](#)
- [MSF-Probability](#)
- [MSF-Statistics](#)
- [MSF-Linear](#) (optional)

Course objective: A student who completes this course should be able to understand, develop, and apply mechanisms for supervised, unsupervised, and reinforcement learning. They should be able to select the proper machine learning algorithm for a problem, preprocess the data appropriately, apply proper evaluation techniques, and explain how to interpret the resulting models, including the model's shortcomings. They should be able to identify and compensate for biased data sets and other sources of error, and be able to explain ethical and societal implications of their application of machine learning to practical problems.

Robotics to include the following:

- [AI-Robotics](#) (25 hours)
- [SPD-Robot](#) (4 hours - focusing on hardware, constraints/considerations, and software architectures; some other topics in SPD/Robot Platforms overlap with AI/Robotics)
- [AI-Search](#) (4 hours - selected topics well-integrated with robotics, e.g., A* and path search)
- [AI-ML](#) (6 hours - selected topics well-integrated with robotics, e.g., neural networks for object recognition)
- [AI-SEP](#) (3 hours - should be integrated throughout the course; robotics is already a huge application, so this really should focus on societal impact and specific robotic applications)

Prerequisites:

- [SDF-Fundamentals](#)
- [SDF-Data-Structures](#)
- [SDF-Algorithms](#)
- [MSF-Linear](#)

Course objective: A student who completes this course should be able to understand and use robotic techniques to perceive the world using sensors, localize the robot based on features and a map, and plan paths and navigate in the world in simple robot applications. They should understand and be able to apply simple computer vision, motion planning, and forward and inverse kinematics techniques.

Introduction to Data Science to include the following:

- [GIT-Visualization](#) (6 hours): types of visualization, libraries, foundations
- [GIT-SEP](#) (2 hours): ethically responsible visualization
- [DM-Core](#) (2 hours): Parallel and distributed processing (MapReduce, cloud frameworks, etc.)
- [DM-Modeling](#) (2 hours): Graph representations, entity resolution
- [DM-Querying](#) (4 hours): SQL, query formation
- [DM-NoSQL](#) (2 hours): Graph DBs, data lakes, data consistency
- [DM-Security](#) (2 hours): privacy, personally identifying information and its protection
- [DM-Analytics](#) (1 hour): exploratory data techniques, data science lifecycle
- [DM-SEP](#) (2 hours): Data provenance
- [AI-ML](#) (15 hours): Data preprocessing, missing data imputation, supervised/semi-supervised/unsupervised learning, text analysis, graph analysis and PageRank, experimental methodology, evaluation, and ethics
- [AI-SEP](#) (3 hours): Applications specific to data science, interspersed throughout the course
- [MSF-Statistics](#) (3 hours): Statistical analysis, hypothesis testing, experimental design

Prerequisites:

- [SDF-Fundamentals](#)

Course objective: A student who completes this course should be able to formulate questions as data analysis problems, understand and use statistical techniques to achieve that analysis from real data, apply visualization techniques to convey the results, and analyze the ethical and societal implications of data science applications. Students should also be able to understand and effectively use data management techniques for preprocessing, storage, security, and retrieval of data in current systems.

Committee

Chair: Eric Eaton, University of Pennsylvania, Philadelphia, PA, USA

Members:

- Zachary Dodds, Harvey Mudd College, Claremont, CA, USA
- Susan L. Epstein, Hunter College and The Graduate Center of The City University of New York, New York, NY, USA
- Laura Hiatt, US Naval Research Laboratory, Washington, DC, USA
- Amruth N. Kumar, Ramapo College of New Jersey, Mahwah, NJ, USA
- Peter Norvig, Google, Mountain View, CA, USA
- Meinolf Sellmann, GE Research, Niskayuna, NY, USA
- Reid Simmons, Carnegie Mellon University, Pittsburgh, PA, USA

Contributors:

- Nate Derbinsky, Northeastern University, Boston, MA, USA
- Eugene Freuder, Insight Centre for Data Analytics, University College Cork, Cork, Ireland
- Ashok Goel, Georgia Institute of Technology, Atlanta, GA, USA
- Claudia Schulz, Thomson Reuters, Zurich, Switzerland

Algorithmic Foundations (AL)

Preamble

Algorithms and data structures are fundamental to computer science, since every theoretical computation and applied program consists of algorithms that operate on data elements possessing some underlying structure. Selecting appropriate computational solutions to real-world problems benefits from understanding the theoretical and practical capabilities and limitations of available algorithms and paradigms, including their impact on the environment and society. Moreover, this understanding provides insight into the intrinsic nature of computation, computational problems, and computational problem-solving as well as possible solution techniques independent of programming language, programming paradigm, computer hardware, or other implementation aspects.

This knowledge area focuses on the nature of computation including the concepts and skills required to design and analyze algorithms for solving real-world computational problems. It complements the implementation of algorithms and data structures found in the Software Development Foundations ([SDF](#)) knowledge area. As algorithms and data structures are essential in all advanced areas of computer science, this area provides the algorithmic foundations that every computer science graduate is expected to know. Exposure to the breadth of these foundational AL topics is designed to provide students with the basis for studying these topics in more depth, additional computation and algorithm topics, and for learning advanced algorithms across a variety of CS knowledge areas and CS-X disciplines.

Changes since CS2013

This area has been renamed from Algorithms and Complexity to better reflect its foundational scope since topics in this area focus on the practical and theoretical foundations of algorithms, complexity and computability. These topics also provide the foundational prerequisites for advanced study in computer science. Additionally, topics focused on complexity and computability have been cleanly separated into respective knowledge units. To reinforce the important impact of computation on society, a Society, Ethics, and the Profession ([AL-SEP](#)) knowledge unit has been added with the expectation that SEP implications be addressed in some manner during every lecture hour of focus in this AL knowledge area.

The increase of four CS Core hours acknowledges the importance of this foundational area in the CS curriculum and returns it to the 2001 level (less than one course). Despite this increase, there is a significant overlap in hours with the Software Development Fundamentals ([SDF](#)) and Mathematical Foundations ([MSF](#)) areas. There is also a complementary nature of the units in this area since, for example, while linear search of an array covers topics in [AL-Foundational](#), it can be used to simultaneously explain [AL-Complexity](#) $O(n)$ and [AL-Strategies](#) Brute-Force topics. The KA topics and hours primarily reflect topics studied in a stand-alone computational theory course and the availability of additional hours when such a course is included in the curriculum.

Core Hours

Knowledge Unit	CS Core	KA Core
Foundational Data Structures and Algorithms	11	6
Algorithmic Strategies	6	
Complexity Analysis	6	3
Computational Models and Formal Languages	9	23
Society, Ethics, and the Profession	Included in SEP hours	
Total	32	32

The 11 CS Core hours in AL-Foundational are in addition to the 9 hours found in [SDF](#) and 3 in [MSF](#).

Knowledge Units

AL-Foundational: Foundational Data Structures and Algorithms

CS Core: (See also: [SDF-Data-Structures](#), [SDF-Algorithms](#))

1. Abstract Data Type (ADT) and operations on an ADT (See also: [FPL-Types](#))
 - a. Dictionary operations (insert, delete, find)
2. Arrays
 - a. Numeric vs non-numeric, character strings
 - b. Single (vector) vs multidimensional (matrix)
3. Records/Structs/Tuples and Objects (See also: [FPL-OOP](#))
4. Linked lists (for historical reasons)
 - a. Single vs Double and Linear vs Circular
5. Stacks
6. Queues and dequeues
 - a. Heap-based priority queue
7. Hash tables / maps
 - a. Collision resolution and complexity (e.g., probing, chaining, rehash)
8. Graphs (e.g., [un]directed, [a]cyclic, [un]connected, and [un]weighted)
(See also: [MSF-Discrete](#))
 - a. Graph representation: adjacency list vs matrix
9. Trees (See also: [MSF-Discrete](#))
 - a. Binary, n-ary, and search trees
 - b. Balanced (e.g., AVL, Red-Black, Heap)
10. Sets (See also: [MSF-Discrete](#))

11. Search algorithms
 - a. $O(n)$ complexity (e.g., linear/sequential array/list search)
 - b. $O(\log_2 n)$ complexity (e.g., binary search)
 - c. $O(\log_b n)$ complexity (e.g. uninformed depth/breadth-first tree search)
12. Sorting algorithms (e.g., stable, unstable)
 - a. $O(n^2)$ complexity (e.g., insertion, selection),
 - b. $O(n \log n)$ complexity (e.g., quicksort, merge, timsort)
13. Graph algorithms
 - a. Shortest path (e.g., Dijkstra's, Floyd's)
 - b. Minimal spanning tree (e.g., Prim's, Kruskal's)

KA Core:

14. Sorting algorithms
 - a. $O(n \log n)$ complexity heapsort
 - b. Pseudo $O(n)$ complexity (e.g., bucket, counting, radix)
15. Graph algorithms
 - a. Transitive closure (e.g., Warshall's)
 - b. Topological sort
16. Matching
 - a. Efficient string matching (e.g., Boyer-Moore, Knuth-Morris-Pratt)
 - b. Longest common subsequence matching
 - c. Regular expression matching

Non-core:

17. Cryptography algorithms (e.g., SHA-256) (See also: [SEC-Crypto](#))
18. Parallel algorithms (See also: [PDC-Algorithms](#), [FPL-Parallel](#))
19. Consensus algorithms (e.g., Blockchain) (See also: [SEC-Crypto](#))
 - a. Proof of work vs proof of stake (See also: [SEP-Sustainability](#))
20. Quantum computing algorithms (See also: [AL-Models](#), [AR-Quantum](#))
 - a. Oracle-based (e.g. Deutsch-Jozsa, Bernstein-Vazirani, Simon)
 - b. Superpolynomial speed-up via QFT (e.g., Shor's)
 - c. Polynomial speed-up via amplitude amplification (e.g., Grover's)
21. Fast-Fourier Transform (FFT) algorithm
22. Differential evolution algorithm

Illustrative Learning Outcomes:

CS Core:

1. For each ADT/Data-Structure in this unit:
 - a. Explain its definition, properties, representation(s), and associated ADT operations,
 - b. Explain step-by-step how the ADT operations associated with the data structure transform it.
2. For each algorithm in this unit:
 - a. Explain step-by-step how the algorithm operates.
3. For each algorithmic approach (e.g., sorting) in this unit:
 - a. Apply a prototypical example of the approach (e.g., merge sort)

4. Given requirements for a problem, develop multiple solutions using various data structures and algorithms. Subsequently, evaluate the suitability, strengths, and weaknesses selecting an approach that best satisfies the requirements.
5. Explain how collision avoidance and collision resolution is handled in hash tables.
6. Explain factors beyond computational efficiency that influence the choice of algorithms, such as: programming time, maintainability, and the use of application-specific patterns in the input data.
7. Explain the heap property and the use of heaps as an implementation of a priority queue.

KA Core:

8. For each of the algorithms and algorithmic approaches in the KA Core topics:
 - a. Explain a prototypical example of the algorithm,
 - b. Explain step-by-step how the algorithm operates.

Non-core:

9. An appreciation of quantum computation and its application to certain problems.

AL-Strategies: Algorithmic Strategies

CS Core:

1. Paradigms
 - a. Brute-Force (e.g., linear search, selection sort, traveling salesperson, knapsack)
 - b. Decrease-and-Conquer
 - i. By a Constant (e.g., insertion sort, topological sort),
 - ii. By a Constant Factor (e.g., binary search),
 - iii. By a Variable Size (e.g., Euclid's)
 - c. Divide-and-Conquer (e.g., binary search, quicksort, mergesort, Strassen's)
 - d. Greedy (e.g., Dijkstra's, Kruskal's, Knapsack)
 - e. Transform-and-Conquer
 - i. Instance simplification (e.g. find duplicates via list presort)
 - ii. Representation change (e.g., heapsort)
 - iii. Problem reduction (e.g., least-common-multiple, linear programming)
 - iv. Dynamic programming (e.g., Floyd's, Warshall, Bellman-Ford)
 - f. Space vs time tradeoffs (e.g., hashing)
2. Handling exponential growth (e.g., heuristic A*, branch-and-bound, backtracking)
3. Iteration vs recursion (e.g., factorial, tree search)

KA Core:

4. Paradigms
 - a. Approximation algorithms
 - b. Iterative improvement (e.g., Ford-Fulkerson, simplex)
 - c. Randomized/Stochastic algorithms (e.g., max-cut, balls and bins)

Non-core:

5. Quantum computing

Illustrative Learning Outcomes:

CS Core:

1. For each of the paradigms in this unit
 - a. Explain its definitional characteristics,
 - b. Explain an example that demonstrates the paradigm including how this example satisfies the paradigm's characteristics
2. For each of the algorithms in the [AL-Foundational](#) unit:
 - a. Explain the paradigm used by the algorithm and how it exemplifies this paradigm
3. Given an algorithm, explain the paradigm used by the algorithm and how it exemplifies this paradigm
4. Give a real-world problem, evaluate appropriate algorithmic paradigms and algorithms from these paradigms that address the problem including evaluating the tradeoffs among the paradigms and algorithms selected.
5. Give examples of iterative and recursive algorithms that solve the same problem, explain the benefits and disadvantages of each approach.
6. Evaluate whether a greedy approach leads to an optimal solution.
7. Explain various approaches for addressing computational problems whose algorithmic solutions are exponential.

AL-Complexity: Complexity

CS Core:

1. Complexity Analysis Framework
 - a. Best, average, and worst case performance of an algorithm
 - b. Empirical and relative (Order of Growth) measurements
 - c. Input size and primitive operations
 - d. Time and space efficiency
2. Asymptotic complexity analysis (average and worst case bounds)
 - a. Big-O, Big-Omega, and Big-Theta formal notations
 - b. Foundational Complexity Classes and Representative Examples/Problems
 - i. $O(1)$ *Constant* (e.g., array access)
 - ii. $O(\log_2 n)$ *Logarithmic* (e.g., binary search)
 - iii. $O(n)$ *Linear*(e.g., linear search)
 - iv. $O(n \log_2 n)$ *Log Linear* (e.g., mergesort)
 - v. $O(n^2)$ *Quadratic* (e.g., selection sort)
 - vi. $O(n^c)$ *Polynomial* (e.g., $O(n^3)$ Gaussian elimination)
 - vii. $O(2^n)$ *Exponential* (e.g., Knapsack, Satisfiability (SAT), Traveling Sales-Person (TSP), all subsets)
 - viii. $O(n!)$ *Factorial* (e.g., Hamiltonian circuit, all permutations)
3. Empirical measurements of performance
4. Tractability and intractability
 - a. P, NP and NP-Complete Complexity Classes
 - b. NP-Complete Problems (e.g., SAT, Knapsack, TSP)

- c. Reductions
- 5. Time and space trade-offs in algorithms.

KA Core:

- 6. Little-o, Little-Omega, and Little Theta notations
- 7. Formal recursive analysis
- 8. Amortized analysis
- 9. Turing Machine-based models of complexity
 - a. Time complexity
 - i. P, NP, NP-C, and EXP classes
 - ii. Cook-Levin theorem
 - b. Space Complexity
 - i. NSpace and PSpace
 - ii. Savitch's theorem

Illustrative Learning Outcomes:

CS Core:

- 1. Prepare a presentation that explains to first year students the basic concepts of algorithmic complexity including: best, average, and worst case algorithm behavior, Big- O, Omega, and Theta notations, complexity classes, time and space tradeoffs, empirical measurement, and impact on practical problems.
- 2. Using examples, explain each of the foundational complexity classes in this unit.
- 3. For each foundational complexity class in this unit:
 - a. Explain an algorithm that demonstrates the associated runtime complexity.
- 4. For each algorithm in the [AL-Foundational](#) unit:
 - a. Explain its runtime complexity class and why it belongs to this class.
- 5. Informally evaluate the foundational complexity class of simple algorithms.
- 6. Given a problem to program for which there may be several algorithmic approaches, evaluate them and determine which are feasible, and select one that is optimal in implementation and run-time behavior.
- 7. Develop empirical studies to determine and validate hypotheses about the runtime complexity of various algorithms by running algorithms on input of various sizes and comparing actual performance to the theoretical analysis.
- 8. Explain examples that illustrate time-space trade-offs of algorithms.
- 9. Explain how tree balance affects the efficiency of binary search tree operations.
- 10. Explain to a non-technical audience the significance of tractable versus intractable algorithms using an intuitive explanation of Big-O complexity.
- 11. Explain the significance of NP-Completeness.
- 12. Explain how NP-Hard is a lower bound and NP is an upper bound for NP-Completeness.
- 13. Explain examples of NP-complete problems.

KA Core:

- 14. Use recurrence relations to evaluate the time complexity of recursively defined algorithms.
- 15. Apply elementary recurrence relations using some form of the Master Theorem.

16. Apply Big-O notation to give upper case bounds on time/space complexity of algorithms.
17. Explain the Cook-Levin Theorem and the NP-Completeness of SAT.
18. Explain the classes P and NP.
19. Prove that a problem is NP-Complete by reducing a classic known NP-C problem to it (e.g., 3SAT and Clique)
20. Explain the P-space class and its relation to the EXP class.

AL-Models: Computational Models and Formal Languages

CS Core:

1. Formal automata
 - a. Finite State
 - b. Pushdown
 - c. Linear Bounded
 - d. Turing Machine
2. Formal languages, grammars and Chomsky Hierarchy
(See also: [FPL-Translation](#), [FPL-Syntax](#))
 - a. Regular (Type-3)
 - i. Regular Expressions
 - b. Context-Free (Type-2)
 - c. Context-Sensitive (Type-1)
 - d. Recursively Enumerable (Type-0)
3. Relations among formal automata, languages, and grammars
4. Decidability, (un)computability, and halting
5. The Church-Turing thesis
6. Algorithmic correctness
 - a. Invariants (e.g., in: iteration, recursion, tree search)

KA Core:

7. Deterministic and nondeterministic automata
8. Pumping Lemma proofs
 - a. Proof of Finite State/Regular-Language limitation
 - b. Pushdown Automata/Context-Free-Language limitation
9. Decidability
 - a. Arithmetization and diagonalization
10. Reducibility and reductions
11. Time complexity based on Turing Machine
12. Space complexity (e.g., PSPACE, Savitch's Theorem)
13. Equivalent models of algorithmic computation
 - a. Turing Machines and Variations (e.g., multi-tape, non-deterministic)
 - b. Lambda Calculus (See also: [FPL-Functional](#))
 - c. Mu-Recursive Functions

Non-core:

14. Quantum computation (See also: [AR-Quantum](#))
 - a. Postulates of quantum mechanics
 - i. State space
 - ii. State evolution
 - iii. State composition
 - iv. State measurement
 - b. Column vector representations of qubits
 - c. Matrix representations of quantum operations
 - d. Simple quantum gates (e.g., XNOT, CNOT)

Illustrative Learning Outcomes:

CS Core:

1. For each formal automata in this unit:
 - a. Explain its definition comparing its characteristics with this unit's other automata,
 - b. Using an example, explain step-by-step how the automata operates on input including whether it accepts the associated input,
 - c. Explain an example of inputs that can and cannot be accepted by the automata.
2. Given a problem, develop an appropriate automaton that addresses the problem.
3. Develop a regular expression for a given regular language expressed in natural language.
4. Explain the difference between regular expressions (Type-3 acceptors) and the regular expressions (Type-2 acceptors) used in programming languages.
5. For each formal model in this unit
 - a. Explain its definition comparing its characteristics with the others in this unit,
 - b. Explain example inputs that are and cannot be accepted by the language/grammar.
6. Explain a universal Turing Machine and its operation.
7. Present to an audience of co-workers and managers the impossibility of providing them a program that checks all other programs, including some seemingly simple ones, for infinite loops including an explanation of the Halting problem, why it has no algorithmic solution, and its significance for real-world algorithmic computation.
8. Explain examples of classic uncomputable problems.
9. Explain the Church-Turing Thesis and its significance for algorithmic computation.
10. Explain how (loop) invariants can be used to prove the correctness of an algorithm.

Illustrative Learning Outcomes:

KA Core:

11. For each formal automata in this unit
 - a. Explain (compare/contrast) its deterministic and nondeterministic capabilities.
12. Apply pumping lemmas, or alternative means, to prove the limitations of Finite State and Pushdown automata.
13. Apply arithmetization and diagonalization to prove the Halting Problem for Turing Machines is Undecidability.
14. Given a known undecidable language, apply a mapping reduction or computational history to prove that another language is undecidable.

15. Convert among equivalently powerful notations for a language, including among DFAs, NFAs, and regular expressions, and between PDAs and CFGs
16. Explain Rice's theorem and its significance.
17. Explain an example proof of a problem that is uncomputable by reducing a classic known uncomputable problem to it.
18. Explain the Primitive and General Recursive functions (zero, successor, selection, primitive recursion, composition, and Mu), their significance, and Turing Machine implementations.
19. Explain how computation is performed in Lambda Calculus (e.g., Alpha conversion and Beta reduction)

Non-core:

20. For a quantum system give examples that explain the following postulates:
 - a. State Space: system state represented as a unit vector in Hilbert space,
 - b. State Evolution: the use of unitary operators to evolve system state,
 - c. State Composition: the use of tensor product to compose systems states,
 - d. State Measurement: the probabilistic output of measuring a system state.
21. Explain the operation of a quantum XNOT or CNOT gate on a quantum bit represented as a matrix and column vector respectively

AL-SEP: Society, Ethics, and the Profession

CS Core: (See also: [SEP-Context](#), [SEP-Sustainability](#))

1. Social, ethical, and secure algorithms
2. Algorithmic fairness
3. Anonymity (e.g., Differential Privacy)
4. Accountability/Transparency
5. Responsible algorithms
6. Economic and other impacts of inefficient algorithms
7. Sustainability

KA Core:

8. Context aware computing

Illustrative Learning Outcomes:

CS Core:

1. Develop algorithmic solutions to real-world societal problems, such as routing an ambulance to a hospital.
2. Explain the impact that an algorithm may have on the environment and society when used to solve a real-world problem taking into account its sustainability and that it can affect different societal groups in different ways.
3. Prepare a presentation that justifies the selection of appropriate data structures and/or algorithms to solve a given real-world problem.

4. Explain an example that articulates how differential privacy protects knowledge of an individual's data.
5. Explain the environmental impacts of design choices that relate to algorithm design.
6. Explain the tradeoffs involved in proof-of-work and proof-of-stake algorithms.

Professional Dispositions

- **Meticulous:** As an algorithm is a formal solution to a computational problem, attention to detail is important when developing and combining algorithms.
- **Persistent:** As developing algorithmic solutions to computational problems can be challenging, computer scientists must be resolute in pursuing such solutions
- **Inventive:** As computer scientists develop algorithmic solutions to real-world problems, they must be inventive in developing solutions to these problems.


Mathematics Requirements

Required:

- [MSF-Discrete](#)

Course Packaging Suggestions

As depicted in the following figure, the committee envisions two common approaches for addressing foundational AL topics in CS courses. Both approaches included required introductory Programming (CS1) and Data Structures (CS2) courses. In a three-course approach, all CS Core topics are covered with additional unused hours to cover other topics. Alternatively, in the four-course approach, the AL-Model knowledge unit CS and KA Core topics are addressed in a Computational Theory focused course, which leaves room to address additional KA topics in the third Algorithms course. Both approaches assume Big-O analysis is introduced in the Data Structures (CS2) course and that graphs are taught in the third Algorithms course. The committee recognizes that there are many different approaches for packaging AL topics into courses including, for example, introducing graphs in CS2 Data Structures, Backtracking in an AI course, and AL-Model topics in a theory course that also addresses, for instance, FPL topics. The given example is simply one way to cover the entire AL CS Core in three introductory courses with additional lecture hours to spare.



Courses Common to Three and Four Course Exemplars

Programming 1 (CS1)

- [AL-Foundational](#) (2 hours)
 - Arrays and Strings
 - Search Algorithms (e.g., $O(n)$ Linear Search)
- [AL-SEP](#) (In SEP hours)

Note: the following AL topics are demonstrated in CS1, but not explicitly taught as such:

- [AL-Strategies](#) (less than hour)
 - Brute Force (e.g., linear search)
 - Iteration (e.g., linear search)
- [AL-Complexity](#) (less than 1 hour)
 - Foundational Complexity Classes
 - $O(1)$ *Constant* and $O(n)$ *Linear* runtime complexities

Course objectives: Students should be able to explain, evaluate, and apply arrays in a variety of problem-solving contexts including using linear search for elements in an array. They should also be able to begin to explain the impact algorithmic design and use has on society.

Data Structures (CS2)

- [AL-Foundational](#) (12 hours)
 - Abstract Data Types and Operations (ADTs)
 - Binary Search
 - Multi-dimensional Arrays
 - Linked Lists
 - Hash Tables/Maps including conflict resolution strategies
 - Records/Structs/Tuples and Objects
 - Sets
 - Stacks, Queues, and Deques
 - Trees: Binary, Ordered, Breadth- and Depth-first search
 - Search Algorithms (e.g., $O(n^2)$ Selection Sort, $O(\log_2 n)$ binary search)
 - Sorting Algorithms (e.g., $O(n \log n)$ Mergesort, $O(\log_b n)$ tree search)
- [AL-Strategies](#) (3 hours)
 - Brute Force (e.g., selection sort)
 - Decrease-and-Conquer (e.g., depth/breadth tree search)
 - Divide-and-Conquer (e.g., mergesort, quicksort)
 - Iteration vs Recursion (e.g., factorial, tree search)
 - Space vs Time tradeoff (e.g., hashing)
- [AL-Complexity](#) (3 hours)
 - Complexity Analysis Framework
 - Foundational Complexity Classes
 - $O(\log_2 n)$ *Logarithmic*, $O(n \log_2 n)$ *Log Linear*, and $O(n^2)$ *Quadratic*
 - Time and Space Trade-Offs in Algorithms
- [AL-SEP](#) (In SEP hours)

Course objectives: Students should be able to explain, evaluate, and apply the specified data structures and algorithms in a variety of problem-solving contexts. Additionally, they should be able to demonstrate the use of different data structures, algorithms, and algorithmic strategies (paradigms) to solve the same problem. Also, they will continue to enhance and refine their understanding of the impact that algorithmic design and use has on society.

Three Course Exemplar Approach

Algorithms-C

- [AL-Foundational](#) (3 hours)
 - Graphs including Graph Algorithms
- [AL-Complexity](#) (3 hours)
 - Asymptotic Complexity Analysis
 - Foundational Complexity Classes
 - $O(2^n)$ Exponential and $O(n!)$ Factorial
 - Empirical Measurements of Performance
 - Tractability and Intractability
- [AL-Strategies](#) (3 hours)
 - Brute Force (e.g., traveling salesperson, knapsack)
 - Decrease-and-Conquer (e.g., topological sort)
 - Divide-and-Conquer (e.g., Strassen's)
 - Greedy (e.g., Dijkstra's, Kruskal's)
 - Transform-and-Conquer/Reduction (e.g., heapsort, trees (2-3, AVL, Red-Black))
 - Dynamic Programming (e.g. Warshall's, Floyd's, Bellman-Ford)
 - Handling Exponential Growth (e.g., heuristic A^* , branch-and-bound, backtracking)
- [AL-Models](#) (9 hours)
 - All CS Core topics
- [AL-SEP](#) (In SEP hours)

Course objectives: Students should be able to explain, evaluate, and apply the specified data structures and algorithms in a variety of problem-solving contexts. Additionally, they should be able to formally explain complexity analysis and the importance of tractability including approaches for handling intractable problems. Finally, they should also be able to summarize formal models of computation, grammars, and languages including the definition of a computer as a Turing Machine and the undecidability of the Halting problem..

Four Course Exemplar Approach

Algorithms-C (third course)

- [AL-Foundational](#) (3 hours)
 - Graphs including Graph Algorithms
 - Sorting Algorithms
 - $O(n \log n)$ heapsort
 - Pseudo $O(n)$ complexity (e.g., bucket, counting, radix)
 - Graph Algorithms

- Transitive closure (e.g., Warshall's)
 - Topological sort
 - Matching
 - Efficient String Matching (e.g., Boyer-Moore, Knuth-Morris-Pratt)
 - Longest common subsequence matching
 - Regular expression matching
- [AL-Complexity](#) (3 hours)
 - Asymptotic Complexity Analysis
 - Foundational Complexity Classes
 - $O(2^n)$ Exponential and $O(n!)$ Factorial
 - Empirical Measurements of Performance
 - Tractability and Intractability
- [AL-Strategies](#) (3 hours)
 - Brute Force (e.g., traveling salesperson, knapsack)
 - Decrease-and-Conquer (e.g., topological sort)
 - Divide-and-Conquer (e.g., Strassen's algorithm)
 - Greedy (e.g., Dijkstra's, Kruskal's)
 - Transform-and-Conquer/Reduction (e.g., heapsort, trees (2-3, AVL, Red-Black))
 - Dynamic Programming (e.g. Warshall's, Floyd's, Bellman-Ford)
 - Handling Exponential Growth (e.g., heuristic A*, branch-and-bound, backtracking)

Course objectives: Students should be able to explain, evaluate, and apply the specified data structures and algorithms in a variety of problem-solving contexts. Additionally, they should be able to formally explain complexity analysis and the importance of tractability including approaches for handling intractable problems.

Computation Theory (fourth course)

- [AL-Complexity](#) (3 hours)
 - Turing Machine-based models of complexity (P, NP, and NP-C classes)
 - Space complexity (NSpace, PSpace Savitch' Theorem)
- [AL-Models](#) (29 hours)
 - All CS and KA Core topics
- [AL-SEP](#) (In SEP hours)

Course objectives: Students should be able to explain, evaluate, and apply models of computation, grammars, and languages. Additionally, they should be able to explain formal proofs that demonstrate the capability and limitations of various automata. Students should be able to relate the complexity of Random Access Models of Computation to Turing Machine models. Finally, students should be able to summarize decidability and reduction proofs.

Committee

Chair: Richard Blumenthal, Regis University, Denver, CO, USA

Members:

- Cathy Bareiss, Bethel University, Mishawaka, MN, USA
- Tom Blanchet, Sci Tec., Boulder, CO, USA
- Doug Lea, State University of New York at Oswego, Oswego, NY, USA
- Sara Miner More, John Hopkins University, Baltimore, MD, USA
- Mia Minnes, University of California San Diego, San Diego, CA, USA
- Atri Rudra, University at Buffalo, Buffalo, NY, USA
- Christian Servin, El Paso Community College, El Paso, TX, USA

Architecture and Organization (AR)

Preamble

Computing professionals spend considerable time writing efficient code to solve a particular problem in an application domain. As the shift from sequential to parallel processing occurs, a deeper understanding of the underlying computer architectures is necessary. Architecture can no longer be viewed as a black box where principles from one architecture can be applied to another. Instead, programmers should look inside the black box and use particular components to enhance system performance and energy efficiency.

The Architecture and Organization (AR) knowledge area aims to develop a deeper understanding of the hardware environments upon which almost all computing is based, and the relevant interfaces provided to higher software layers. The target hardware comprises low-end embedded system processors up to high-end enterprise multiprocessors.

The topics in this knowledge area will benefit students by enabling them to appreciate the fundamental architectural principles of modern computer systems, including the challenge of harnessing parallelism to sustain performance and energy improvements into the future. This KA will help computer science students depart from the black box approach and become more aware of the underlying computer system and the efficiencies specific architectures can achieve.

Changes since CS2013

Changes and additions are summarized as follows:

- Topics have been revised, particularly AR/Memory Hierarchy and AR/Performance and Energy Efficiency. This update brings recent advances in memory caching and energy consumption.
- The newly created AR/Heterogeneous Architectures covers emerging topics in Computer Architecture: Processing In-Memory (PIM) and domain-specific architectures (e.g., neural network processors).
- The new AR/Quantum Architectures offers a "toolbox" covering introductory topics in quantum computing.
- Knowledge units have been merged to better deal with overlaps:
 - AR/Multiprocessing and Alternative Architectures were merged into newly created AR/Heterogeneous Architectures.
- The new AR/Secure Processor Architectures covers hardware support for multi-stack security applications.

Core Hours

Knowledge Unit	CS Core		KA Core
----------------	---------	--	---------

Digital Logic and Digital Systems			2 + 1 (SF)
Machine-Level Data Representation	1		
Assembly Level Machine Organization	1		1 + 1 (PDC)
Memory Hierarchy	4+2 (OS)		
Interfacing and Communication	1		
Functional Organization			2
Performance and Energy Efficiency			3
Heterogeneous Architectures			2
Secure Processor Architectures			2
Quantum Architectures			2
Sustainability Issues	Included in SEP hours		
Total	9		16

The hours shared with OS include overlapping topics and are counted here.

Knowledge Units

AR-Logic: Digital Logic and Digital Systems

KA Core:

1. Combinational vs sequential logic/field programmable gate arrays (FPGAs) (See also: [SF-Overview](#), [SF-Foundations](#), [SPD-Embedded](#))
 - a. Fundamental combinational
 - b. Sequential logic building block
2. Computer-aided design tools that process hardware and architectural representations
3. High-level synthesis
 - a. Register transfer notation
 - b. Hardware description language (e.g., Verilog/VHDL/Chisel)
4. System-on-chip (SoC) design flow
5. Physical constraints
 - a. Gate delays
 - b. Fan-in and fan-out
 - c. Energy/power
 - d. Speed of light

Illustrative Learning Outcomes:

KA Core:

1. Discuss the progression of computer technology components from vacuum tubes to VLSI, from mainframe computer architectures to the organization of warehouse-scale computers.
2. Describe parallelism and data dependencies between and within components in a modern heterogeneous computer architecture.
3. Explain how the relationship between parallelism and power consumption.
4. Construct the design of basic building blocks for a computer: arithmetic-logic unit (gate-level), registers (gate-level), central processing unit (register transfer-level), and memory (register transfer-level).
5. Evaluate simple building blocks (e.g., arithmetic-logic unit, registers, movement between registers) of a simple computer design.
6. Analyze the timing behavior of a pipelined processor, identifying data dependency issues.

AR-Representation: Machine-Level Data Representation

CS Core:

1. Overview and history of computer architecture (See also: [SPD-Game](#))
2. Bits, bytes, and words
3. Unsigned, signed and two's complement representations
4. Numeric data representation and number bases
 - a. Fixed-point
 - b. Floating-point
5. Representation of non-numeric data
6. Representation of records, arrays and UTF data types (See also: [AL-Foundational](#))

Illustrative Learning Outcomes:

CS Core:

1. Discuss why everything is data in computers, including instructions.
2. Explain how fixed-length number representations can affect accuracy and precision.
3. Describe how negative integers are stored in sign-magnitude and two's-complement representations.
4. Discuss how different formats can represent numerical data.
5. Explain the bit-level representation of non-numeric data, such as characters, strings, records, and arrays.
6. Translate numerical data from one format to another.
7. Describe how a single adder (without overflow detection) can handle both signed (two's complement) and unsigned (binary) input without "knowing" which format a given input is using.

AR-Assembly: Assembly Level Machine Organization

CS Core:

1. von Neumann machine architecture
2. Control unit: instruction fetch, decode, and execution (See also: [OS-Principles](#))
3. Introduction to SIMD vs MIMD and the Flynn taxonomy (See also: [PDC-Programs](#), [OS-Scheduling](#), [OS-Process](#))
4. Shared memory multiprocessors/multicore organization (See also: [PDC-Programs](#), [OS-Scheduling](#))

KA Core:

5. Instruction set architecture (ISA) (e.g., x86, ARM and RISC-V)
 - a. Fixed vs variable-width instruction sets
 - b. Instruction formats
 - c. Data manipulation, control, I/O
 - d. Addressing modes
 - e. Machine language programming
 - f. Assembly language programming
6. Subroutine call and return mechanisms (See also: [FPL-Translation](#), [OS-Principles](#))
7. I/O and interrupts (See also: [OS-Principles](#))
8. Heap, static, stack, and code segments (See also: [FPL-Translation](#), [OS-Process](#))

Illustrative Learning Outcomes:

CS Core:

1. Discuss how the classical von Neumann functional units are implemented in embedded systems, particularly on-chip and off-chip memory.
2. Describe how instructions are executed in a classical von Neumann machine, with extensions for threads, multiprocessor synchronization, and SIMD execution.
3. Assess an example diagram with instruction-level parallelism and hazards to describe how they are managed in typical processor pipelines.

KA Core:

4. Discuss how instructions are represented at the machine level and in the context of a symbolic assembler.
5. Map an example of high-level language patterns into assembly/machine language notations.
6. Contrast different instruction formats considering aspects such as addresses per instruction and variable-length vs fixed-length formats.
7. Analyze a subroutine diagram to comment on how subroutine calls are handled at the assembly level.
8. Describe basic concepts of interrupts and I/O operations.
9. Write a simple assembly language program for string/array processing and manipulation.

AR-Memory: Memory Hierarchy

CS Core:

1. Memory hierarchy: the importance of temporal and spatial locality (See also: [SF-Performance](#), [OS-Memory](#))
2. Main memory organization and operations (See also: [OS-Memory](#))

3. Persistent memory (e.g., SSD, standard disks)
4. Latency, cycle time, bandwidth and interleaving (See also: [SF-Performance](#))
5. Cache memories (See also: [SF-Performance](#))
 - a. Address mapping
 - b. Block size
 - c. Replacement and store policy
 - d. Prefetching
6. Multiprocessor cache coherence (See also: [OS-Scheduling](#))
7. Virtual memory (hardware support) (See also: [OS-Memory](#))
8. Fault handling and reliability (See also: [SF-Reliability](#))
9. Reliability (See also: [SF-Reliability](#), [OS-Faults](#))
 - a. Error coding
 - b. Data compression
 - c. Data integrity

KA Core:

10. Processing In-Memory (PIM)

Illustrative Learning Outcomes:

CS Core:

1. Using a memory system diagram, identify the main types of memory technology (e.g., SRAM, DRAM) and their relative cost and performance.
2. Measure the effect of memory latency on running time.
3. Enumerate the functions of a system with virtual memory management.
4. Compute average memory access time under various cache and memory configurations and mixes of instruction and data references.

AR-IO: Interfacing and Communication

CS Core:

1. I/O fundamentals (See also: [OS-Devices](#), [PDC-Communication](#))
 - a. Handshaking and buffering
 - b. Programmed I/O
 - c. Interrupt-driven I/O ([OS-Principles](#))
2. Interrupt structures: vectored and prioritized, interrupt acknowledgment (See also: [OS-Principles](#))
3. I/O devices (e.g., mouse, keyboard, display, camera, sensors, actuators) (See also: [GIT-Fundamentals](#), [GIT-Interaction](#), [OS-Advanced-Files](#), [PDC-Programs](#))
4. External storage, physical organization, and drives
5. Buses fundamentals (See also: [OS-Devices](#))
 - a. Bus protocols
 - b. Arbitration
 - c. Direct-memory access (DMA)

Illustrative Learning Outcomes:

CS Core:

1. Analyze an interrupt control diagram to comment on how interrupts are used to implement I/O control and data transfers.
2. Enumerate various types of buses in a computer system.
3. List the advantages of magnetic disks and contrast them with those of solid-state disks.

AR-Organization: Functional Organization

KA Core:

1. Implementation of simple datapaths, including instruction pipelining, hazard detection, and resolution (e.g., stalls, forwarding)
2. Control unit
 - a. Hardwired implementation
 - b. Microprogrammed realization
3. Instruction pipelining (See also: [SF-Overview](#))
4. Introduction to instruction-level parallelism (ILP) (See also: [PDC-Programs](#))

Illustrative Learning Outcomes:

KA Core:

1. Compare alternative implementation of datapaths in modern computer architectures.
2. Produce a set of control signals for adding two integers using hardwired and microprogrammed implementations.
3. Discuss instruction-level parallelism using pipelining and significant hazards that may occur.
4. Design a complete processor, including datapath and control.
5. Compute the average cycles per instruction for a given processor and memory system implementation.

AR-Performance-Energy: Performance and Energy Efficiency

KA Core:

1. Performance-energy evaluation (introduction): performance, power consumption, memory, and communication costs (See also: [SF-Evaluation](#), [OS-Scheduling](#), [SPD-Game](#))
2. Branch prediction, speculative execution, out-of-order execution, Tomasulo's algorithm
3. Enhancements for vector processors and GPUs (See also: [SPD-Game](#))
4. Hardware support for multithreading (See also: [OS-Concurrency](#), [OS-Scheduling](#), [PDC-Programs](#))
 - a. Race conditions
 - b. Lock implementations
 - c. Point-to-point synchronization
 - d. Barrier implementation
5. Scalability
6. Alternative architectures including VLIW/EPIC, accelerators, and other special purpose processors
7. Dynamic voltage and frequency scaling (DVFS)
8. Dark Silicon

Illustrative Learning Outcomes:

KA Core:

1. Discuss performance and energy efficiency evaluation metrics.
2. Analyze a speculative execution diagram and write about the decisions that can be made.
3. Create a GPU performance-watt benchmarking diagram.
4. Write a multithreaded program that adds (in parallel) elements of two integer vectors.
5. Recommend a set of design choices for alternative computer architectures.
6. Enumerate key concepts associated with dynamic voltage and frequency scaling.
7. Measure energy savings improvement for an 8-bit integer quantization compared to a 32-bit quantization.

AR-Heterogeneity: Heterogeneous Architectures

KA Core:

1. SIMD and MIMD architectures (e.g., General-Purpose GPUs, TPUs, and NPU) (See also: [PDC-Programs](#), [SPD-Embedded](#), [GIT-Shading](#), [SPD-Game](#))
2. Heterogeneous memory systems (See also: [OS-Process](#), [PDC-Communication](#))
 - a. Shared memory versus distributed memory
 - b. Volatile vs non-volatile memory
 - c. Coherence protocols
3. Domain-Specific Architectures (DSAs) (See also: [HCI-Accountability](#), [GIT-Shading](#))
 - a. Machine Learning Accelerator
 - b. In-networking computing (See also: NC-Applications)
 - c. Embedded systems for emerging applications
 - d. Neuromorphic computing
 - e. Edge computing devices
4. Packaging and integration solutions such as 3DIC and Chiplets
5. Machine learning in architecture design
 - a. AI algorithms for workload analysis
 - b. Optimization of architecture configurations for performance and power efficiency

Illustrative Learning Outcomes:

KA Core

1. Analyze a system diagram with alternative parallel architectures, e.g., SIMD and MIMD, and identify the key differences.
2. Discuss what memory-management issues are found in multiprocessors that are not present in uniprocessors and how these issues might be resolved.
3. Indicate the differences between memory backplane, processor memory interconnect, and remote memory via networks, their implications for access latency, and their impact on program performance.
4. Discuss how you would determine when to use a domain-specific accelerator instead of a general-purpose CPU.

5. Enumerate key differences in architectural design principles between a vector and scalar-based processing unit.
6. List the advantages and disadvantages of PIM architectures.

AR-Security: Secure Processor Architectures

KA core:

1. Principles of Secure Hardware
 - a. Security Risk Analysis, Asset Protection, and Threat Model
 - b. Cryptographic Acceleration with Hardware (See also: [SEC-Crypto](#))
 - c. Support for virtualization (e.g., OS isolation)
2. Roots of trust in hardware, Physically Unclonable Functions (PUF)
3. Hardware Random Number Generators
4. Memory protection extensions
 - a. Runtime pointer bounds checking (e.g., buffer overflow)
 - b. Protection at the microarchitectural level
 - c. Protection at the ISA level
5. Trusted Execution Environment (TEE)
 - a. Trusted Computer Base Protections
 - b. Protecting virtual machines
 - c. Protecting containers
 - d. Trusted software modules (Enclaves)
6. Homomorphic encryption for privacy-preserving data processing

Illustrative Learning Outcomes

KA Core:

1. Discuss principles of secure hardware, exploring a framework for risk analysis and asset protection.
2. Summarize how Physically Unclonable Functions (PUF) can be a unique device identifier in security applications.
3. Distinguish a random number generator with dedicated hardware support from generators without hardware dedicated to generating entropy.
4. List the advantages and disadvantages of memory protection at the ISA level.
5. Describe key design issues of a trusted execution environment (TEE) to support virtual machines.

AR-Quantum: Quantum Architectures

KA Core:

1. Principles (See also: AL-Models: 8)
 - a. The wave-particle duality principle
 - b. The uncertainty principle in the double-slit experiment
 - c. What is a Qubit? Superposition, interference and measurement. Photons as qubits.
 - d. Systems of two qubits. Entanglement. Bell states. The No-Signaling theorem.
2. Axioms of QM: superposition principle, measurement axiom, unitary evolution

3. Single qubit gates for the circuit model of quantum computation: X, Z, H.
4. Two qubit gates and tensor products. Working with matrices.
5. The No-Cloning Theorem. The Quantum Teleportation protocol.
6. Algorithms (See also: [AL-Foundational](#))
 - a. Simple quantum algorithms: Bernstein-Vazirani, Simon's algorithm.
 - b. Implementing Deutsch-Josza with Mach-Zehnder Interferometers.
 - c. Quantum factoring (Shor's Algorithm)
 - d. Quantum search (Grover's Algorithm)
7. Implementation aspects (See also: [SPD-Interactive](#))
 - a. The physical implementation of qubits
 - b. Classical control of a Quantum Processing Unit (QPU)
 - c. Error mitigation and control. NISQ and beyond.
 - d. Measurement approaches
8. Emerging Applications
 - a. Post-quantum encryption
 - b. The Quantum Internet
 - c. Adiabatic quantum computation (AQC) and quantum annealing

Illustrative Learning Outcomes:

KA Core:

1. Discuss how a quantum object produced as a particle propagates like a wave and is detected as a particle with a probability distribution corresponding to the wave.
2. Discuss the quantum-level nature that is inherently probabilistic.
3. Express your view on entanglement that can be used to create non-classical correlations, but there is no way to use quantum entanglement to send messages faster than the speed of light.
4. Describe quantum parallelism and the role of constructive vs destructive interference in quantum algorithms given the probabilistic nature of measurement(s).
5. Analyze a code snippet providing the role of quantum Fourier transform (QFT) in Shor's algorithm.
6. Write a program to implement Shor's algorithm in a simulator, highlighting the classical components and aspects of Shor's algorithm.
7. Enumerate the specifics of each qubit modality (e.g., trapped ion, superconducting, silicon spin, photonic, quantum dot, neutral atom, topological, color center, electron-on-helium, etc.).
8. Contrast AQC with the gate model of quantum computation and the problems each is better suited to solve.

AR-SEP: Sustainability Issues

Non-core:

1. Environmental impacts of implementation decisions
 - a. Sustainability goals, resource consumption, and economic viability.
 - b. Carbon footprint, hardware electronic waste
 - c. The energy footprint of data centers at various workloads (e.g., AI model training and use)
 - d. Guidelines for sustainable design standards

Illustrative Learning Outcomes:

Non-core:

1. Assess the environmental impacts of a given project's deployment (e.g., the energy consumption of CPUs and GPUs, contribution to e-waste, and effect of hardware virtualization in data centers).

Professional Dispositions

- **Self-directed:** Students should increasingly become self-motivated to acquire complementary knowledge.
- **Proactive:** Students should exercise control and anticipate issues related to the underlying computer system.

Mathematics Requirements

- [MSF-Discrete](#), [MSF-Linear](#), [MSF-Statistics](#), [MSF-Calculus](#), [MSF-Probability](#)

Course Packaging Suggestions

Computer Architecture - Introductory Course to include the following:

- [SEP-History](#) (2 hours)
- [AR-Representation](#) (2 hours)
- [AR-Assembly](#) (2 hours)
- [AR-Memory](#) (10 hours)
- [OS-Memory](#) (10 hours)
- [AR-IO](#) (4 hours)
- [AR-Heterogeneity](#) (5 hours)
- [PDC-Programs](#) (4 hours)
- [SEP-Ethical-Analysis](#) (3 hours)

Course objectives: Students should understand the fundamentals of modern computer architectures, including the challenges associated with memory caches, memory management and pipelining.

Prerequisites:

- [MSF-Discrete](#)

Computer Architecture - Advanced Topics Course to include the following:

- [AR-Logic](#) (4 hours)
- [AR-Representation](#) (2 hours)
- [AR-Assembly](#) (2 hours)
- [AR-Memory](#) (10 hours)
- [AR-IO](#) (2 hours)
- [SF-Performance](#) (4 hours)
- [AR-Heterogeneity](#) (4 hours)
- [AR-Performance-Energy](#) (5 hours)

- [AR-Security](#) (4 hours)
- [AR-Quantum](#) (4 hours)

Course objectives: Students should understand how computer architectures evolved into today's heterogeneous systems and to what extent choices made in the past can influence the design of future high-performance computing systems.

Prerequisites:

- [MSF-Discrete](#)

Systems Course to include the following:

- [SEP-History](#) (2 hours)
- [SF-Design](#) (2 hours)
- [SF-Reliability](#) (2 hours)
- [OS-Purpose](#) (2 hours)
- [AR-Representation](#) (2 hours)
- [AR-Assembly](#) (2 hours)
- [AR-Memory](#) (8 hours)
- [AR-IO](#) (2 hours)
- [PDC-Algorithms](#) (4 hours)
- [AR-Heterogeneity](#) (4 hours)
- [AR-Performance-Energy](#) (5 hours)
- [NC-Applications](#) (5 hours)

Course objectives: Students should understand the advanced architectural aspects of modern computer systems, including heterogeneous architectures and the required hardware and software interfaces to improve the performance and energy footprint of applications.

Prerequisites:

- [MSF-Discrete](#), [MSF-Statistics](#)

Committee

Chair: Marcelo Pias, Federal University of Rio Grande (FURG), Rio Grande-RS, Brazil

Members:

- Brett A. Becker, University College Dublin, Dublin, Ireland
- Mohamed Zahran, New York University, New York, NY, USA
- Monica D. Anderson, University of Alabama, Tuscaloosa, AL, USA
- Qiao Xiang, Xiamen University, Xiamen, China
- Adrian German, Indiana University, Bloomington, IN, USA

Data Management (DM)

Preamble

Since the mid-1970s, the study of Data Management (DM) has meant an almost exclusive study of relational database systems. Depending on institutional context, students have studied, in varying proportions:

- Data modeling and database design: for example, E-R Data model, relational model, normalization theory
- Query construction: e.g., relational algebra, SQL
- Query processing: e.g., indices (B+tree, hash), algorithms (e.g., external sorting, select, project, join), query optimization (transformations, index selection)
- DBMS internals: e.g., concurrency/locking, transaction management, buffer management

Today's graduates are expected to possess DBMS user (as opposed to implementor) skills. These primarily include data modeling and query construction; ability to take an unorganized collection of data, organize it using a DBMS, and access/update the collection via queries.

Additionally, students need to study:

- The role data plays in an organization. This includes:
 - o The Data Life Cycle: Creation-Processing-Review/Reporting-Retention/Retrieval-Destruction.
 - o The social/legal aspects of data collection: e.g., scale, data privacy, database privacy (compliance) by design, de-identification, ownership, reliability, database security, and intended and unintended applications.
- Emerging and advanced technologies that are augmenting/replacing traditional relational systems, particularly those used to support (big) data analytics, including NoSQL (e.g., JSON, XML, key-value store databases), cloud databases, MapReduce, and dataframes.
- We recognize the existing and emerging roles for those involved with data management, which include:
 - Product feature engineers: those who use both SQL and NoSQL operational databases.
 - Analytical engineers/data engineers: those who write analytical SQL, Python, and Scala code to build data assets for business groups.
 - Business analysts: those who build/manage data most frequently with Excel spreadsheets.
 - Data infrastructure engineers: those who implement a data management system in a variety of data applications (e.g., OLTP).
 - "Everyone:" those who produce or consume data need to understand the associated social, ethical, and professional issues.

One role that transcends all the above categories is that of data custodian. Previously, data was seen as a resource to be managed (Information Systems Management) just like other enterprise resources. Today, data is seen in a larger context. Data about customers can now be seen as belonging to (or in some national contexts, as owned by) those customers. There is now an accepted understanding that the safe and ethical storage, and use, of institutional data is part of being a responsible data custodian.

Furthermore, we acknowledge the tension between a curricular focus on professional preparation versus the study of a knowledge area as a scientific endeavor. This is particularly true with Data Management. For example, proving (or at least knowing) the completeness of Armstrong's Axioms is fundamental in functional dependency theory. However, most computer science graduates will never utilize this concept during their professional careers. The same can be said for many other topics in the Data Management canon. Conversely, if our graduates can only normalize data into Boyce-Codd normal form (using an automated tool) and write SQL queries, without understanding the role that indices play in efficient query execution, we have done them and society a disservice.

To this end, the number of CS Core hours is relatively small relative to the KA Core hours. This approach is designed to allow institutions with differing contexts to customize their curricula appropriately. An institution that focuses on OLTP implementation, for example, would prioritize efficient storage and data access, while an institution that focuses on product features would prioritize programmatic access to extant databases.

However an institution manages this tension we wish to give voice to one of the ironies of computer science curricula. Students typically spend much of their educational life reading (and writing) data from a file or interactively, while outside of the academy the predominant data comes from databases accessed programmatically. Perhaps in the not-too-distant future students will learn programmatic database access early on and then continue this practice as they progress through their curriculum.

Finally, we understand that while the Data Management KA may be orthogonal to the SEC (Security) and SEP (Society, Ethics, and the Profession) KAs, it is also ground zero for these (and other) knowledge areas. When designing persistent data stores, the question of what should be stored must be examined from both legal and ethical perspectives. Are there privacy concerns? And just as importantly, how well protected is the data?

Changes since CS2013

- Rename the knowledge area from Information Management to Data Management. This renaming does not represent any kind of philosophical shift. It is simply an effort to avoid confusion with the similar definitions used in Information Systems and Information Technology curricula.
- Inclusion of NoSQL approaches and MapReduce as CS Core topics.
- Increased attention to SEP and SEC topics in both the CS Core and KA Core areas.

Core Hours

Knowledge Unit	CS Core Hours	KA Core Hours
The Role of Data	2	
Core Database Systems Concepts	2	1

Data Modeling	2	3
Relational Databases	1	3
Query Construction	2	4
Query Processing		4
DBMS Internals		4
NoSQL Systems		2
Data Security & Privacy	1	2
Data Analytics		3
Distributed Databases/Cloud Computing		
Semi-structured and Unstructured Databases		
Society, Ethics, and the Profession	Included in SEP hours	
Total	10	26

The CS Core hours in Data Security & Privacy are shared with SEC include overlapping topics and are counted here.

Knowledge Units

DM-Data: The Role of Data and the Data Life Cycle

CS Core:

1. The Data Life Cycle: Creation-Processing-Review/Reporting-Retention/Retrieval-Destruction (See also: [SEP-Context](#), [SEP-Ethical-Analysis](#), [SEP-Professional-Ethics](#), [SEP-Privacy](#), [SEP-Security](#), [SEC-Foundations](#))

Illustrative Learning Outcomes:

CS Core:

1. Identify the five stages of the Data Life Cycle

DM-Core: Core Database System Concepts

CS Core:

1. Purpose and advantages of database systems
2. Components of database systems
3. Design of core DBMS functions (e.g., query mechanisms, transaction management, buffer management, access methods)
4. Database architecture, data independence, and data abstraction
5. Transaction management
6. Normalization
7. Approaches for managing large volumes of data (e.g., NoSQL database systems, use of MapReduce) (See also: [PDC-Algorithms](#))
8. How to support CRUD-only applications
9. Distributed databases/cloud-based systems
10. Structured, semi-structured, and unstructured data
11. Use of a declarative query language

KA Core:

12. Systems supporting structured and/or stream content

Illustrative Learning Outcomes:

CS Core:

1. Identify at least four advantages that using a database system provides.
2. Enumerate the components of a (relational) database system.
3. Follow a query as it is processed by the components of a (relational) database system.
4. Defend the value of data independence.
5. Compose a simple select-project-join query in SQL.
6. Enumerate the four properties of a correct transaction manager.
7. Describe the advantages for eliminating duplicate repeated data.
8. Outline how MapReduce uses parallelism to process data efficiently.
9. Evaluate the differences between structured and semi/unstructured databases.

DM-Modeling: Data Modeling

CS Core:

1. Data modeling (See also: [SE-Requirements](#))
2. Relational data model (See also: [MSF-Discrete](#))

KA Core:

3. Conceptual models (e.g., entity-relationship, UML diagrams)
4. Semi-structured data models (expressed using DTD, XML, or JSON Schema, for example)

Non-core:

5. Spreadsheet models
6. Object-oriented models (See also: [FPL-OOP](#))
 - a. GraphQL
7. New features in SQL

8. Specialized Data Modeling topics
 - a. Time series data (aggregation, and join)
 - b. Graph data (link traversal)
 - c. Techniques for avoiding inefficient raw data access (e.g., “avg daily price”): materialized views and special data structures (e.g., Hyperloglog, bitmap)
 - d. Geo-Spatial data (e.g., GIS databases) (See also: [SPD-Interactive](#))

Illustrative Learning Outcomes:

CS Core:

1. Describe the components of the relational data model
2. Model 1:1, 1:n, and n:m relationships using the relational data model

KA Core:

3. Describe the components of the E-R (or some other non-relational) data model.
4. Model a given environment using a conceptual data model.
5. Model a given environment using the document-based or key-value store-based data model.

DM-Relational: Relational Databases

CS Core:

1. Entity and referential integrity
 - a. Candidate key, superkeys
2. Relational database design

KA Core:

3. Mapping conceptual schema to a relational schema
4. Physical database design: file and storage structures (See also: [OS-Files](#))
5. Introduction to Functional dependency theory
6. Normalization Theory
 - a. Decomposition of a schema; lossless-join and dependency-preservation properties of a decomposition
 - b. Normal forms (BCNF)
 - c. Denormalization (for efficiency)

Non-core:

7. Functional dependency theory
 - a. Closure of a set of attributes
 - b. Canonical Cover
8. Normalization theory
 - a. Multi-valued dependency (4NF)
 - b. Join dependency (PJNF, 5NF)
 - c. Representation theory

Illustrative Learning Outcomes:

CS Core:

1. Describe the defining characteristics behind the relational data model.
2. Comment on the difference between a foreign key and a superkey.
3. Enumerate the different types of integrity constraints.

KA Core:

4. Compose a relational schema from a conceptual schema which contains 1:1, 1:n, and n:m relationships.
5. Map appropriate file structure to relations and indices.
6. Describe how functional dependency theory generalizes the notion of key.
7. Defend a given decomposition as lossless and or dependency preserving.
8. Detect which normal form a given decomposition yields.
9. Comment on reasons for denormalizing a relation.

DM-Querying: Query Construction**CS Core:**

1. SQL Query Formation
 - a. Interactive SQL execution
 - b. Programmatic execution of an SQL query

KA Core:

2. Relational Algebra
3. SQL
 - a. Data definition including integrity and other constraints specification
 - b. Update sublanguage

Non-core:

4. Relational Calculus
5. QBE and 4th-generation environments
6. Different ways to invoke non-procedural queries in conventional languages
7. Introduction to other major query languages (e.g., XPATH, SPARQL)
8. Stored procedures

Illustrative Learning Outcomes:**CS Core:**

1. Compose SQL queries that incorporate select, project, join, union, intersection, set difference, and set division.
2. Determine when a nested SQL query is correlated or not.
3. Iterate over data retrieved programmatically from a database via an SQL query.

KA Core:

4. Define, in SQL, a relation schema, including all integrity constraints and delete/update triggers.

5. Compose an SQL query to update a tuple in a relation.

DM-Processing: Query Processing

KA Core:

1. Page structures
2. Index structures
 - a. B+ trees (See also: [AL-Foundational](#))
 - b. Hash indices: static and dynamic (See also: [AL-Foundational](#), [SEC-Foundations](#))
 - c. Index creation in SQL
3. File structures (See also: [OS-Files](#))
 - a. Heap files
 - b. Hash files
4. Algorithms for query operators
 - a. External Sorting (See also: [AL-Foundational](#))
 - b. Selection
 - c. Projection; with and without duplicate elimination
 - d. Natural Joins: Nested loop, Sort-merge, Hash join
 - e. Analysis of algorithm efficiency (See also: [AL-Complexity](#))
5. Query transformations
6. Query optimization
 - a. Access paths
 - b. Query plan construction
 - c. Selectivity estimation
 - d. Index-only plans
7. Parallel Query Processing (e.g., parallel scan, parallel join, parallel aggregation) (See also: [PDC-Algorithms](#))
8. Database tuning/performance
 - a. Index selection
 - b. Impact of indices on query performance (See also: [SF-Performance](#), [SEP-Sustainability](#))
 - c. Denormalization

Illustrative Learning Outcomes:

KA Core:

1. Describe the purpose and organization of both B+ tree and hash index structures.
2. Compose an SQL command to create an index (any kind).
3. Specify the steps for the various query operator algorithms: external sorting, projection with duplicate elimination, sort-merge join, hash-join, block nested-loop join.
4. Derive the run-time (in I/O requests) for each of the above algorithms
5. Transform a query in relational algebra to its equivalent appropriate for a left-deep, pipelined execution.
6. Compute selectivity estimates for a given selection and/or join operation.
7. Describe how to modify an index structure to facilitate an index-only operation for a given relation.

8. For a given scenario decide on which indices to support for the efficient execution of a set of queries.
9. Describe how DBMSs leverage parallelism to speed up query processing by dividing the work across multiple processors or nodes.

DM-Internals: DBMS Internals

KA Core:

1. DB Buffer Management (See also: [OS-Memory](#), [SF-Resource](#))
2. Transaction Management (See also: [PDC-Coordination](#))
 - a. Isolation Levels
 - b. ACID
 - c. Serializability
 - d. Distributed Transactions
3. Concurrency Control: (See also: [OS-Concurrency](#))
 - a. 2-Phase Locking
 - b. Deadlocks handling strategies
 - c. Quorum-based consistency models
4. Recovery Manager
 - a. Relation with Buffer Manager

Non-core:

5. Concurrency Control:
 - a. Optimistic concurrency control
 - b. Timestamp concurrency control
6. Recovery Manager
 - a. Write-Ahead logging
 - b. ARIES recovery system (Analysis, REDO, UNDO)

Illustrative Learning Outcomes:

KA Core:

1. Describe how a DBMS manages its Buffer Pool
2. Describe the four properties for a correct transaction manager
3. Outline the principle of serializability

DM-NoSQL: NoSQL Systems

KA Core:

1. Why NoSQL? (e.g. Impedance mismatch between Application [CRUD] and RDBMS)
2. Key-Value and Document data model

Non-core:

3. Storage systems (e.g., Key-Value systems, Data Lakes)

4. Distribution Models (Sharding and Replication) (See also: [PDC-Communication](#))
5. Graph Databases
6. Consistency Models (Update and Read, Quorum consistency, CAP theorem) (See also: [PDC-Communication](#))
7. Processing model (e.g., Map-Reduce, multi-stage map-reduce, incremental map-reduce) (See also: [PDC-Communication](#))
8. Case Studies: Cloud storage system (e.g., S3); Graph databases; “When not to use NoSQL” (See also: [SPD-Web](#))

Illustrative Learning Outcomes:

KA Core:

1. Develop a use case for the use of NoSQL over RDBMS.
2. Describe the defining characteristics behind Key-Value and Document-based data models.

DM-Security: Data Security and Privacy

CS Core:

1. Differences between data security and data privacy (See also: [SEC-Foundations](#))
2. Protecting data and database systems from attacks, including injection attacks such as SQL injection (See also: [SEC-Foundations](#))
3. Personally identifying information (PII) and its protection (See also: [SEC-Foundations](#), [SEP-Security](#), [SEP-Privacy](#))
4. Ethical considerations in ensuring the security and privacy of data (See also: [SEC-SEP](#), [SEP-Ethical-Analysis](#), [SEP-Security](#), [SEP-Privacy](#))

KA Core:DM

5. Need for, and different approaches to securing data at rest, in transit, and during processing (See also: [SEC-Foundations](#), [SEC-Crypto](#))
6. Database auditing and its role in digital forensics (See also: [SEC-Forensics](#))
7. Data inferencing and preventing attacks (See also: [SEC-Crypto](#))
8. Laws and regulations governing data security and data privacy (See also: [SEP-Security](#), [SEP-Privacy](#), [SEC-Foundations](#), [SEC-Governance](#))

Non-core:

9. Typical risk factors and prevention measures for ensuring data integrity (See also: [SEC-Governance](#))
10. Ransomware and prevention of data loss and destruction (See also: [SEC-Coding](#), [SEC-Forensics](#))

Illustrative Learning Outcomes:

CS Core:

1. Describe the differences in the goals for data security and data privacy.
2. Identify and mitigate risks associated with different approaches to protecting data.

3. Describe legal and ethical considerations of end-to-end data security and privacy.

KA Core:

4. Develop a database auditing system given risk considerations.
5. Apply several data exploration approaches to understanding unfamiliar datasets.

DM-Analytics: Data Analytics

KA Core:

1. Exploratory data techniques (motivation, representation, descriptive statistics, visualizations)
2. Data science lifecycle: business understanding, data understanding, data preparation, modeling, evaluation, deployment, and user acceptance (See also: [AI-ML](#))
3. Data mining and machine learning algorithms: e.g., classification, clustering, association, regression (See also: [AI-ML](#))
4. Data acquisition and governance (See also: [SEC-Governance](#))
5. Data security and privacy considerations (See also: [SEP-Security](#), [SEP-Privacy](#), [SEC-Foundations](#))
6. Data fairness and bias (See also: [SEP-Security](#), [AI-SEP](#))
7. Data visualization techniques and their use in data analytics (See also: [GIT-Visualization](#))
8. Entity Resolution

Illustrative Learning Outcomes:

KA Core:

1. Describe several data exploration approaches, including visualization, to understanding unfamiliar datasets.
2. Apply several data exploration approaches to understanding unfamiliar datasets.
3. Describe basic machine learning/data mining algorithms and when they are appropriate for use.
4. Apply several machine learning/data mining algorithms.
5. Describe legal and ethical considerations in acquiring, using, and modifying datasets.
6. Describe issues of fairness and bias in data collection and usage

DM-Distributed: Distributed Databases/Cloud Computing

Non-core:

1. Distributed DBMS (See also [PDC-Communications](#))
 - a. Distributed data storage
 - b. Distributed query processing
 - c. Distributed transaction model
 - d. Homogeneous and heterogeneous solutions
 - e. Client-server distributed databases (See also: [NC-Fundamentals](#))
2. Parallel DBMS (See also: [PDC-Algorithms](#))
 - a. Parallel DBMS architectures: shared memory, shared disk, shared nothing;
 - b. Speedup and scale-up, e.g., use of the MapReduce processing model (See also: [PDC-Programs](#), [SF-Foundations](#))

- c. Data replication and weak consistency models (See also: [PDC-Coordination](#))

DM-Unstructured: Semi-structured and Unstructured Databases

Non-core:

1. Vectorized unstructured data (text, video, audio, etc.) and vector storage
 - a. TF-IDF Vectorizer with ngram
 - b. Word2Vec
 - c. Array database or array data type handling
2. Semi-structured databases (e.g., JSON)
 - a. Storage
 - i. Encoding and compression of nested data types
 - b. Indexing
 - i. Btree, skip index, Bloom filter
 - ii. Inverted index and bitmap compression
 - iii. Space filling curve indexing for semi-structured geo-data
 - c. Query processing for OLTP and OLAP use cases
 - i. Insert, Select, update/delete trade offs
 - ii. Case studies on Postgres/JSON, MongoDB and Snowflake/JSON

DM-SEP: Society, Ethics, and the Profession

CS Core:

1. Issues related to scale (See also: [SEP-Economies](#))
2. Data privacy overall (See also: [SEP-Privacy](#), [SEP-Ethical-Analysis](#))
 - a. Privacy compliance by design (See also: [SEP-Privacy](#))
3. Data anonymity (See also: [SEP-Privacy](#))
4. Data ownership/custodianship (See also: [SEP-Professional-Ethics](#))
5. Intended and unintended applications of stored data (See also: [SEP-Professional-Ethics](#), [SEC-Foundations](#))

KA Core:

6. Reliability of data (See also: [SEP-Security](#))
7. Provenance, data lineage, and metadata management (See also: [SEP-Professional-Ethics](#))
8. Data security (See also: [DM-Security](#), [SEP-Security](#))

Illustrative Learning Outcomes:

CS Core:

1. Enumerate three social and three legal issues related to large data collections
2. Describe the value of data privacy
3. Identify the competing stakeholders with respect to data ownership
4. Enumerate three negative unintended consequences from a given (well known) data-centric application (e.g., Facebook, LastPass, Ashley Madison)

KA Core:

5. Describe the meaning of data provenance and lineage
6. Identify how a database might contribute to data security as well as how it may introduce insecurities.

Professional Dispositions

- **Meticulous:** Those who either access or store data collections must be meticulous in fulfilling data ownership responsibilities.
- **Responsible:** In conjunction with the professional management of (personal) data, it is equally important that data be managed responsibly. Protection from unauthorized access as well as prevention of irresponsible, though legal, use of data is paramount. Furthermore, data custodians need to protect data not only from outside attack, but from crashes and other foreseeable dangers.
- **Collaborative:** Data managers and data users must behave in a collaborative fashion to ensure that the correct data is accessed, and is used only in an appropriate manner.
- **Responsive:** The data that gets stored and is accessed is always in response to an institutional need/request.

Mathematics Requirements

Required:

- Discrete Mathematics
 - Set theory (union, intersection, difference, cross-product) (See also: [MSF-Discrete](#))

Desired:

- Probability and Statistics for those studying [DM-Analytics](#). (See also: [MSF-Probability](#), [MSF-Statistics](#))

Desirable Data Structures:

- Hash functions and tables (See also: [AL-Foundational](#))
- Balanced (binary) trees (e.g., AVL, 2-3-4, Red-Black) (See also: [AL-Foundational](#))
- B and B+-trees

Course Packaging Suggestions

For those implementing a single course on Database Systems, there are a variety of options. As described in [27], there are four primary perspectives from which to approach databases:

- Database design/modeling
- Database use
- Database administration
- Database development, which includes implementation algorithms

Course design proceeds by focusing on topics from each perspective in varying degrees according to one's institutional context. For example, in [27], one of the courses described can be characterized as design/modeling (20%), use (20%), development/internals (30%), and administration/tuning/advanced topics (30%). The topics might include:

- [DM-SEP](#) (3 hours)
- [DM-Data](#) (1 hour)
- [DM-Core](#) (3 hours)
- [DM-Modeling](#) (5 hours)
- [DM-Relational](#) (4 hours)
- [DM-Querying](#) (6 hours)
- [DM-Processing](#) (5 hours)
- [DM-Internals](#) (5 hours)
- [DM-NoSQL](#) (4 hours)
- [DM-Security](#) (3 hours)
- [DM-Distributed](#) (2 hours)

Possibly, the more interesting question is how to cover the CS Core concepts in the absence of a dedicated database course. Perhaps the key to accomplishing this is to *normalize* database access. Starting with the introductory course students could be accessing a database versus file I/O or interactive data entry, to acquire the data needed for introductory-level programming. As students progress through their curriculum, additional CS Core topics can be introduced. For example, introductory students would be given the code to access the database along with the SQL query. By the intermediate level, they could be writing their own queries. Finally, in a Software Engineering or capstone course, they are practicing database design. One advantage of this *databases across the curriculum* approach is that allows for the inclusion of database-related SEP topics to also be spread across the curriculum.

In a similar vein one might have a whole course on the Role of Data from either a Security ([SEC](#)) perspective, or an Ethics ([SEP](#)) perspective.

Committee

Chair: Mikey Goldweber, Denison University, Granville, OH, USA

Members:

- Sherif Aly, The American University in Cairo, Cairo, Egypt
- Sara More, Johns Hopkins University, Baltimore, MD, USA
- Mohamed Mokbel, University of Minnesota, Minneapolis, MN, USA
- Rajendra K. Raj, Rochester Institute of Technology, Rochester, NY, USA
- Avi Silberschatz, Yale University, New Haven, CT, USA
- Min Wei, Microsoft, Seattle, WA, USA
- Qiao Xiang, Xiamen University, Xiamen, China

Foundations of Programming Languages (FPL)

Preamble

The foundations of programming languages are rooted in discrete mathematics, logic and formal languages, and provide a basis for the understanding of complex modern programming languages. Although programming languages vary according to the language paradigm and the problem domain and evolve in response to both societal needs and technological advancement, they share an underlying abstract model of computation and program development. This remains true even as processor hardware and their interface with programming tools become increasingly intertwined and progressively more complex. An understanding of the common abstractions and programming paradigms enables faster learning of programming languages.

The Foundations of Programming Languages knowledge area is concerned with articulating the underlying concepts and principles of programming languages, the formal specification of a programming language and the behavior of a program, explaining how programming languages are implemented, comparing the strengths and weaknesses of various programming paradigms, and describing how programming languages interface with entities such as operating systems and hardware. The concepts covered here are applicable to several languages and an understanding of these principles assists a learner to move readily from one language to another, as well as select a programming paradigm and language that best suits the problem at hand.

Programming languages are the medium through which programmers precisely describe concepts, formulate algorithms, and reason about solutions. Over the course of a career, a computer scientist will learn and work with many different languages, separately or together. Software developers must understand different programming models, programming features and constructs, and underlying to make informed design choices among languages that support multiple complementary approaches. It would be useful to know how programming language features are defined, composed, and implemented to improve execution efficiency and long-term maintenance of developed software. Also useful is a basic knowledge of language translation, program analysis, run-time behavior, memory management and interplay of concurrent processes communicating with each other through message-passing, shared memory, and synchronization. Finally, some developers and researchers will need to design new languages, an exercise which requires greater familiarity with basic principles.

Changes since CS2013

Changes since 2013 include a change in name of the KA from Programming Languages to Foundations of Programming Languages to reflect the fact that the KA is about the fundamentals underpinning programming languages, and related concepts, not about any specific programming language. Changes also include a redistribution of content formerly identified as core Tier-1 and core Tier-2 within the Programming Language Knowledge Area (KA). In CS2013, graduates were expected to complete all Tier-1 topics and 80% of Tier-2 topics, for a total of 24 required hours. These are now designated CS Core topics. The remaining Tier-2 topics are now designated KA Core topics. The change in core topics (Tier-1 plus 80% of Tier-2 hours) from 2013 reflects the change in importance or

relevance of topics over the past decade. The inclusion of new topics was driven by their current prominence in the programming language landscape, or the anticipated impact of emerging areas on the profession in general. Specifically, the changes are:

- Object-Oriented Programming -4 CS Core hours
- Functional Programming -2 CS Core hours
- Event-Driven and Reactive Programming +1 CS Core hour
- Parallel and Distributed Computing +3 CS Core hours
- Type Systems -1 CS Core hour
- Program Representation -1 CS Core hour

In addition, some knowledge units from CS2013 were renamed to reflect their content more accurately:

- Static Analysis was renamed Program Analysis and Analyzers
- Concurrency and Parallelism was renamed Parallel and Distributed Computing
- Program Representation was renamed Program Abstraction and Representation
- Runtime Systems was renamed Runtime Behavior and Systems
- Basic Type Systems and Type Systems were merged into a single topic and named Type Systems

Six new knowledge units have been added to reflect their continuing and growing importance as we look toward the 2030s:

- Shell Scripting +2 CS Core hours
- Systems Execution and Memory Model +3 CS Core hours
- Formal Development Methodologies
- Design Principles of Programming Languages
- Fundamentals of Programming Languages and Society, Ethics, and the Profession

Notes:

- Several topics within this knowledge area either build on or overlap with content covered in other knowledge areas such as the Software Development Fundamentals Knowledge Area in a curriculum's introductory courses. Curricula will differ on which topics are integrated in this fashion and which are delayed until later courses on software development and programming languages.
- Different programming paradigms correspond to different problem domains. Most languages have evolved to integrate more than one programming paradigm such as imperative with object-oriented, functional programming with object-oriented, logic programming with object-oriented, and event and reactive modeling with object-oriented programming. Hence, the emphasis is not on just one programming paradigm but on a balance of all major programming paradigms.
- While the number of CS Core and KA Core hours is identified for each major programming paradigm (object-oriented, functional, logic), the distribution of hours across the paradigms may differ depending on the curriculum and programming languages students have been exposed to leading up to coverage of this knowledge area. This document assumes that students have exposure to an object-oriented programming language leading into this knowledge area.
- Imperative programming is not listed as a separate paradigm to be examined. Instead it is treated as a subset of the object-oriented paradigm.

- With multicore computing, cloud computing, and computer networking becoming commonly available in the market, it has become critical to understand the integration of “distribution, concurrency, parallelism” along with other programming paradigms as a core area. This paradigm is integrated with almost all other major programming paradigms.
- With ubiquitous computing and real-time temporal computing applications increasing in daily human life within domains such as health, transportation, smart homes, it has become important to cover the software development aspects of event-driven and reactive programming as well as parallel and distributed computing. Some of the topics covered will require and overlap with concepts in knowledge areas such as Architecture and Organization, Operating Systems, and Systems Fundamentals.
- Some topics from the Parallel and Distributed Computing knowledge unit are likely to be integrated within the curriculum with topics from the Parallel and Distributed Programming kKnowledge area.
- There is an increasing interest in formal methods to prove program correctness and other properties. To support this, additional coverage of topics related to formal methods is included, but all these topics are identified as non-core.
- When introducing these topics, it is also important that an instructor provides context for this material including why we have an interest in programming languages and what they do for us in terms of providing a human readable version of instructions for a computer to execute.

Core Hours

Knowledge Units	CS Core	KA Core
Object-Oriented Programming	4 + 1 (SDF)	1
Functional Programming	4	3
Logic Programming		2 + 1 (MSF)
Shell Scripting	2	
Event-Driven and Reactive Programming	2	2
Parallel and Distributed Computing	2 + 1 (PDC)	2
Aspect-Oriented Programming		
Type Systems	3	3
Systems Execution and Memory Model	2 + 1 (AR and OS)	
Language Translation and Execution	2	3
Program Abstraction and Representation		3
Syntax Analysis		
Compiler Semantic Analysis		
Program Analysis and Analyzers		
Code Generation		
Runtime Behavior and Systems		
Advanced Programming Constructs		
Language Pragmatics		
Formal Semantics		

Formal Development Methodologies		
Design Principles of Programming Languages		
Society, Ethics, and the Profession	Included in SEP hours	
Total	21	19

Compared to CS2013 which had a total of 24 CS Core hours (Tier-1 hours plus 80% of Tier-2 hours), and 4 KA Core hours (20% of Tier-2 hours), the current recommendation has a total of 24 CS core hours (of which 3 are shared with other knowledge areas, and counted toward the CS Core hour total in that knowledge area) and 20 KA core hours (of which 1 is shared with another knowledge area and counted there).

Knowledge Units

FPL-OOP: Object-Oriented Programming

CS Core:

1. Imperative programming as a subset of object-oriented programming.
2. Object-oriented design:
 - a. Decomposition into objects carrying state and having behavior.
 - b. Class-hierarchy design for modeling.
3. Definition of classes: fields, methods, and constructors. (See also: [SDF-Fundamentals](#))
4. Subclasses, inheritance (including multiple inheritance), and method overriding.
5. Dynamic dispatch: definition of method-call.
6. Exception handling. (See also: [SDF-Fundamentals](#), [PDC-Coordination](#), [SE-Construction](#))
7. Object-oriented idioms for encapsulation:
 - a. Privacy, data hiding, and visibility of class members.
 - b. Interfaces revealing only method signatures.
 - c. Abstract base classes, traits and mixins.
8. Dynamic vs static properties.
9. Composition vs inheritance.
10. Subtyping:
 - a. Subtype polymorphism; implicit upcasts in typed languages.
 - b. Notion of behavioral replacement: subtypes acting like supertypes.
 - c. Relationship between subtyping and inheritance.

KA Core:

11. Collection classes, iterators, and other common library components.
12. Metaprogramming and reflection.

Illustrative Learning Outcomes:

CS Core:

1. Enumerate the differences between imperative and object-oriented programming paradigms.
2. Compose a class through design, implementation, and testing to meet behavioral requirements.

3. Build a simple class hierarchy utilizing subclassing that allows code to be reused for distinct subclasses.
4. Predict and validate control flow in a program using dynamic dispatch.
5. Compare and contrast how computational solutions to a problem differ in procedural, functional, and object-oriented approaches.
6. Compare and contrast mechanisms to define and protect data elements within procedural, functional, and object-oriented approaches.
7. Compare and contrast the benefits and costs/impact of using inheritance (subclasses) and composition (in particular how to base composition on higher order functions).
8. Explain the relationship between object-oriented inheritance (code-sharing and overriding) and subtyping (the idea of a subtype being usable in a context that expects the supertype).
9. Use object-oriented encapsulation mechanisms such as interfaces and private members.
10. Define and use iterators and other operations on aggregates, including operations that take functions as arguments, in multiple programming languages, selecting the most natural idioms for each language. (See also: [FPL-Functional](#))

KA Core:

11. Use collection classes and iterators effectively to solve a problem.

FPL-Functional: Functional Programming

CS Core:

1. Lambda expressions and evaluation: (See also: [AL-Models](#), [FPL-Formalism](#))
 - a. Variable binding and scope rules. (See also: [SDF-Fundamentals](#))
 - b. Parameter-passing. (See also: [SDF-Fundamentals](#))
 - c. Nested lambda expressions and reduction order.
2. Effect-free programming:
 - a. Function calls have no side effects, facilitating compositional reasoning.
 - b. Immutable variables and data copying vs reduction.
 - c. Use of recursion vs loops vs pipelining (map/reduce).
3. Processing structured data (e.g., trees) via functions with cases for each data variant:
 - a. Functions defined over compound data in terms of functions applied to the constituent pieces.
 - b. Persistent data structures.
4. Using higher-order functions (taking, returning, and storing functions).

KA Core:

5. Metaprogramming and reflection.
6. Function closures (functions using variables in the enclosing lexical environment):
 - a. Basic meaning and definition - creating closures at run-time by capturing the environment.
 - b. Canonical idioms: call-backs, arguments to iterators, reusable code via function arguments.
 - c. Using a closure to encapsulate data in its environment.
 - d. Lazy versus eager evaluation.

Non-core:

7. Graph reduction machine and call-by-need.

8. Implementing lazy evaluation.
9. Integration with logic programming paradigm using concepts such as equational logic, narrowing, residuation and semantic unification. (See also: [FPL-Logic](#))
10. Integration with other programming paradigms such as imperative and object-oriented.

Illustrative learning outcomes:

CS Core:

1. Develop basic algorithms that avoid assigning to mutable state or considering reference equality.
2. Develop useful functions that take and return other functions.
3. Compare and contrast how computational solutions to a problem differ in procedural, functional, and object-oriented approaches.
4. Compare and contrast mechanisms to define and protect data elements within procedural, functional, and object-oriented approaches.

KA Core:

5. Explain a simple example of lambda expression being implemented using a virtual machine, such as a SECD machine, showing storage and reclaim of the environment.
6. Correctly interpret variables and lexical scope in a program using function closures.
7. Use functional encapsulation mechanisms such as closures and modular interfaces.
8. Compare and contrast stateful vs stateless execution.
9. Define and use iterators and other operations on aggregates, including operations that take functions as arguments, in multiple programming languages, selecting the most natural idioms for each language. (See also: [FPL-OOP](#))

Non-core:

10. Illustrate graph reduction using a λ -expression using a shared subexpression.
11. Illustrate the execution of a simple nested λ -expression using an abstract machine, such as an ABC machine.
12. Illustrate narrowing, residuation and semantic unification using simple illustrative examples.
13. Illustrate the concurrency constructs using simple programming examples of known concepts such as a buffer being read and written concurrently or sequentially. (See also: [FPL-OOP](#))

FPL-Logic: Logic Programming

KA Core:

1. Universal vs existential quantifiers. (See also: [AI-LRR](#), [MSF-Discrete](#))
2. First order predicate logic vs higher order logic. (See also: [AI-LRR](#), [MSF-Discrete](#))
3. Expressing complex relations using logical connectives and simpler relations.
4. Definitions of Horn clause, facts, goals and subgoals.
5. Unification and unification algorithm; unification vs assertion vs expression evaluation.
6. Mixing relations with functions. (See also: [MSF-Discrete](#))
7. Cuts, backtracking, and non-determinism.
8. Closed-world vs open-world assumptions.

Non-core:

9. Memory overhead of variable copying in handling iterative programs.
10. Programming constructs to store partial computation and pruning search trees.
11. Mixing functional programming and logic programming using concepts such as equational logic, narrowing, residuation and semantic unification. (See also: [FPL-Functional](#))
12. Higher-order, constraint, and inductive logic programming. (See also: [AI-LRR](#))
13. Integration with other programming paradigms such as object-oriented programming.
14. Advance programming constructs such as difference-lists, creating user defined data structures, set of, etc.

Illustrative learning outcomes:

KA Core:

1. Use a logic language to implement a conventional algorithm.
2. Use a logic language to implement an algorithm employing implicit search using clauses, relations, and cuts.
3. Use a simple illustrative example to show correspondence between First Order Predicate Logic (FOPL) and logic programs using Horn clauses.
4. Use examples to illustrate the unification algorithm and its role of parameter-passing in query reduction.
5. Use simple logic programs interleaving relations, functions, and recursive programming such as factorial and Fibonacci numbers and simple complex relationships between entities, and illustrate execution and parameter-passing using unification and backtracking.

Non-core:

6. Illustrate computation of simple programs such as Fibonacci and show overhead of recomputation, and then show how to improve execution overhead.

FPL-Scripting: Shell Scripting

CS Core:

1. Error/exception handling.
2. Piping. (See also: [AR-Organization](#), [SF-Overview](#), [OS-Process](#))
3. System commands: (See also: [SF-Overview](#))
 - a. Interface with operating systems. (See also: [SF-Overview](#), [OS-Principles](#))
4. Environment variables. (See also: [SF-Overview](#))
5. File abstraction and operators. (See also: [SDF-Fundamentals](#), [OS-Files](#), [SF-Resource](#))
6. Data structures, such as arrays and lists. (See also: [AL-Foundational](#), [SDF-Fundamentals](#), [SDF-Data-Structures](#))
7. Regular expressions. (See also: [AL-Models](#))
8. Programs and processes. (See also: [OS-Process](#))
9. Workflow.

Illustrative learning outcomes:

CS Core:

1. Create and execute automated scripts to manage various system tasks.
2. Solve various text processing problems through scripting.

FPL-Event-Driven: Event-Driven and Reactive Programming

CS Core:

1. Procedural programming vs reactive programming: advantages of reactive programming in capturing events.
2. Components of reactive programming: event-source, event signals, listeners and dispatchers, event objects, adapters, event-handlers. (See also: [GIT-Interaction](#), [SPD-Web](#), [SPD-Mobile](#), [SPD-Robot](#), [SPD-Embedded](#), [SPD-Game](#), [SPD-Interactive](#))
3. Stateless and state-transition models of event-based programming.
4. Canonical uses such as GUIs, mobile devices, robots, servers. (See also: [GIT-Interaction](#), [GIT-Image](#), [SPD-Web](#), [SPD-Mobile](#), [SPD-Robot](#), [SPD-Embedded](#), [SPD-Game](#), [SPD-Interactive](#))

KA Core:

5. Using a reactive framework:
 - a. Defining event handlers/listeners.
 - b. Parameterization of event senders and event arguments.
 - c. Externally generated events and program-generated events.
6. Separation of model, view, and controller.
7. Event-driven and reactive programs as state-transition systems.

Illustrative learning outcomes:

CS Core:

1. Implement event handlers for use in reactive systems, such as GUIs.
2. Examine why an event-driven programming style is natural in domains where programs react to external events.

KA Core:

3. Define and use a reactive framework.
4. Describe an interactive system in terms of a model, a view, and a controller.

FPL-Parallel: Parallel and Distributed Computing

CS Core:

1. Safety and liveness: (See also: [PDC-Evaluation](#))
 - a. Race conditions. (See also: [OS-Concurrency](#))
 - b. Dependencies/preconditions.
 - c. Fault models. (See also: [OS-Faults](#))
 - d. Termination. (See also: [PDC-Coordination](#))
2. Programming models. (See also: [PDC-Programs](#))

One or more of the following:

 - a. Actor models.
 - b. Procedural and reactive models.
 - c. Synchronous/asynchronous programming models.
 - d. Data parallelism.
3. Properties (See also: [PDC-Programs](#), [PDC-Coordination](#))

- a. Order-based properties:
 - i. Commutativity.
 - ii. Independence.
- b. Consistency-based properties:
 - i. Atomicity.
 - ii. Consensus.
- 4. Execution control: (See also: [PDC-Coordination](#), [SF-Foundations](#))
 - a. Async await.
 - b. Promises.
 - c. Threads.
- 5. Communication and coordination: (See also: [OS-Process](#), [PDC-Communication](#), [PDC-Coordination](#))
 - a. Mutexes.
 - b. Message-passing.
 - c. Shared memory.
 - d. Cobegin-coend.
 - e. Monitors.
 - f. Channels.
 - g. Threads.
 - h. Guards.

KA Core:

- 6. Futures.
- 7. Language support for data parallelism such as forall, loop unrolling, map/reduce.
- 8. Effect of memory-consistency models on language semantics and correct code generation.
- 9. Representational State Transfer Application Programming Interfaces (REST APIs).
- 10. Technologies and approaches: cloud computing, high performance computing, quantum computing, ubiquitous computing
- 11. Overheads of message-passing
- 12. Granularity of program for efficient exploitation of concurrency.
- 13. Concurrency and other programming paradigms (e.g., functional).

Illustrative learning outcomes:

CS Core:

- 1. Explain why programming languages do not guarantee sequential consistency in the presence of data races and what programmers must do as a result.
- 2. Implement correct concurrent programs using multiple programming models, such as shared memory, actors, futures, synchronization constructs, and data-parallelism primitives.
- 3. Use a message-passing model to analyze a communication protocol.
- 4. Use synchronization constructions such as monitor/synchronized methods in a simple program.
- 5. Modeling data dependency using simple programming constructs involving variables, read and write.
- 6. Modeling control dependency using simple constructs such as selection and iteration.

KA Core:

7. Explain how REST API's integrate applications and automate processes.
8. Explain benefits, constraints and challenges related to distributed and parallel computing.

FPL-Aspect: Aspect-Oriented Programming**Non-core:**

1. Aspects.
2. Join points.
3. Advice:
 - a. Before.
 - b. After (as finally, returning or throwing).
 - c. Around.
4. Point cuts:
 - a. Designators.
5. Weaving - static and dynamic.
6. Alternatives including annotations and IDEs.

FPL-Types: Type Systems**CS Core:**

1. A type as a set of values together with a set of operations:
 - b. Primitive types (e.g., numbers, Booleans). (See also: [SDF-Fundamentals](#))
 - c. Compound types built from other types (e.g., records/structs, unions, arrays, lists, functions, references using set operations). (See also: [SDF-Data-Structures](#))
2. Association of types to variables, arguments, results, and fields.
3. Type safety as an aspect of program correctness. (See also: [FPL-Formalism](#))
4. Type safety and errors caused by using values inconsistently given their intended types.
5. Goals and limitations of static and dynamic typing:
 - a. Detecting and eliminating errors as early as possible.
6. Generic types (parametric polymorphism):
 - a. Definition and advantages of polymorphism: parametric, subtyping, overloading and coercion.
 - b. Comparison of monomorphic and polymorphic types.
 - c. Comparison with ad-hoc polymorphism (overloading) and subtype polymorphism.
 - d. Generic parameters and typing.
 - e. Use of generic libraries such as collections.
 - f. Comparison with ad hoc polymorphism (overloading) and subtype polymorphism.
 - g. Prescriptive vs descriptive polymorphism.
 - h. Implementation models of polymorphic types.
 - i. Subtyping.

KA Core:

7. Type equivalence: structural vs name equivalence.
8. Complementary benefits of static and dynamic typing:
 - a. Errors early vs errors late/avoided.

- b. Enforce invariants during code development and code maintenance vs postpone typing decisions while prototyping and conveniently allow flexible coding patterns such as heterogeneous collections.
- c. Typing rules:
 - i. Rules for function, product, and sum types.
- d. Avoid misuse of code vs allow more code reuse.
- e. Detect incomplete programs vs allow incomplete programs to run.
- f. Relationship to static analysis.
- g. Decidability.

Non-core:

- 9. Compositional type constructors, such as product types (for aggregates), sum types (for unions), function types, quantified types, and recursive types.
- 10. Type checking.
- 11. Subtyping: (See also: [FPL-OOP](#))
 - a. Subtype polymorphism; implicit upcasts in typed languages.
 - b. Notion of behavioral replacement: subtypes acting like supertypes.
 - c. Relationship between subtyping and inheritance.
- 12. Type safety as preservation plus progress.
- 13. Type inference.
- 14. Static overloading.
- 15. Propositions as types (implication as a function, conjunction as a product, disjunction as a sum). (See also: [FPL-Formalism](#))
- 16. Dependent types (universal quantification as dependent function, existential quantification as dependent product). (See also: [FPL-Formalism](#))

Illustrative learning outcomes:

CS Core:

- 1. Describe, for both a primitive and a compound type, the values that have that type.
- 2. Describe, for a language with a static type system, the operations that are forbidden statically, such as passing the wrong type of value to a function or method.
- 3. Describe examples of program errors detected by a type system.
- 4. Identify program properties, for multiple programming languages, that are checked statically and program properties that are checked dynamically.
- 5. Describe an example program that does not type-check in a particular language and yet would have no error if run.
- 6. Use types and type-error messages to write and debug programs.

KA Core:

- 7. Explain how typing rules define the set of operations that are legal for a type.
- 8. List the type rules governing the use of a particular compound type.
- 9. Explain why undecidability requires type systems to conservatively approximate program behavior.
- 10. Define and use program pieces (such as functions, classes, methods) that use generic types, including for collections.

11. Discuss the differences among generics, subtyping, and overloading.
12. Explain multiple benefits and limitations of static typing in writing, maintaining, and debugging software.

Non-core:

13. Define a type system precisely and compositionally.
14. For various foundational type constructors, identify the values they describe and the invariants they enforce.
15. Precisely describe the invariants preserved by a sound type system.
16. Prove type safety for a simple language in terms of preservation and progress theorems.
17. Implement a unification-based type-inference algorithm for a simple language.
18. Explain how static overloading and associated resolution algorithms influence the dynamic behavior of programs.

FPL-Systems: Systems Execution and Memory Model

CS Core:

1. Data structures for translation, execution, translation and code mobility such as stack, heap, aliasing (sharing using pointers), indexed sequence and string.
2. Direct, indirect, and indexed access to memory location.
3. Run-time representation of data abstractions such as variables, arrays, vectors, records, pointer-based data elements such as linked-lists and trees, and objects.
4. Abstract low-level machine with simple instruction, stack and heap to explain translation and execution.
5. Run-time layout of memory: activation record (with various pointers), static data, call-stack, heap: (See also: [AR-Memory](#), [OS-Memory](#))
 - a. Translating selection and iterative constructs to control-flow diagrams.
 - b. Translating control-flow diagrams to low level abstract code.
 - c. Implementing loops, recursion, and tail calls.
 - d. Translating function/procedure calls and return from calls, including different parameter-passing mechanisms using an abstract machine.
6. Memory management: (See also: [AR-Memory](#), [OS-Memory](#))
 - a. Low level allocation and accessing of high-level data structures such as basic data types, n-dimensional array, vector, record, and objects.
 - b. Return from procedure as automatic deallocation mechanism for local data elements in the stack.
 - c. Manual memory management: allocating, de-allocating, and reusing heap memory.
 - d. Automated memory management: garbage collection as an automated technique using the notion of reachability.
7. Green computing. (See also: [SEP-Sustainability](#))

Illustrative learning outcomes:

CS Core:

1. Explain how a core language construct, such as data abstractions and control abstractions, is executed.

2. Explain how programming language implementations typically organize memory into global data, text, heap, and stack sections and how features such as recursion and memory management map to this memory model.
3. Explain why memory leaks and dangling pointer problems occur, and what can be done by a programmer to avoid/fix them.

FPL-Translation: Language Translation and Execution

CS Core:

1. Execution models for JIT (Just-In-Time), compiler, interpreter.
2. Use of intermediate code, e.g., bytecode.
3. Limitations and benefits of JIT, compiler, and interpreter.
4. Cross compilers/transpilers.
5. BNF and extended BNF representation of context-free grammar.
6. Parse tree using a simple sentence such as arithmetic expression or if-then-else statement.
7. Execution as native code or within a virtual machine.
8. Language translation pipeline: syntax analysis, parsing, optional type-checking, translation/code generation and optimization, linking, loading, execution.

KA Core:

9. Run-time representation of core language constructs such as objects (method tables) and first-class functions (closures).
10. Secure compiler development. (See also: [SEC-Foundations](#), [SEC-Coding](#))

Illustrative learning outcomes:

CS Core:

1. Explain and understand the differences between compiled, JIT, and interpreted language implementations, including the benefits and limitations of each.
2. Differentiate syntax and parsing from semantics and evaluation.
3. Use BNF and extended BNF to specify the syntax of simple constructs such as if-then-else, type declaration and iterative constructs for known languages such as C++ or Python.
4. Illustrate the parse tree using a simple sentence/arithmetic expression.
5. Illustrate translation of syntax diagrams to BNF/extended BNF for simple constructs such as if-then-else, type declaration, iterative constructs, etc.
6. Illustrate ambiguity in parsing using nested if-then-else/arithmetic expression and show resolution using precedence order.

KA-Core:

7. Discuss the benefits and limitations of garbage collection, including the notion of reachability.

FPL-Abstraction: Program Abstraction and Representation

KA Core:

1. BNF and regular expressions.

2. Programs that take (other) programs as input such as interpreters, compilers, type-checkers, documentation generators.
3. Components of a language:
 - a. Definitions of alphabets, delimiters, sentences, syntax and semantics.
 - b. Syntax vs semantics.
4. Program as a set of non-ambiguous meaningful sentences.
5. Basic programming abstractions: constants, variables, declarations (including nested declarations), command, expression, assignment, selection, definite and indefinite iteration, iterators, function, procedure, modules, exception handling. (See also: [SDF-Fundamentals](#))
6. Mutable vs immutable variables: advantages and disadvantages of reusing existing memory location vs advantages of copying and keeping old values; storing partial computation vs recomputation.
7. Types of variables: static, local, nonlocal, global; need and issues with nonlocal and global variables.
8. Scope rules: static vs dynamic; visibility of variables; side-effects.
9. Side-effects induced by nonlocal variables, global variables and aliased variables.

Non-core:

10. L-values and R-values: mapping mutable variable-name to L-values; mapping immutable variable-names to R-values.
11. Environment vs store and their properties.
12. Data and control abstraction.
13. Mechanisms for information exchange between program units such as procedures, functions and modules: nonlocal variables, global variables, parameter-passing, import-export between modules.
14. Data structures to represent code for execution, translation, or transmission.
15. Low level instruction representation such as virtual machine instructions, assembly language, and binary representation. (See also: [AR-Representation](#), [AR-Assembly](#))
16. Lambda calculus, variable binding, and variable renaming. (See also: [AL-Models](#), [FPL-Formalism](#))
17. Types of semantics: operational, axiomatic, denotational, behavioral; define and use abstract syntax trees; contrast with concrete syntax.

Illustrative learning outcomes:

KA Core:

1. Illustrate the scope of variables and visibility using simple programs.
2. Illustrate different types of parameter-passing using simple pseudo programming language.
3. Explain side-effect using global and nonlocal variables and how to fix such programs.
4. Explain how programs that process other programs treat the other programs as their input data.
5. Describe a grammar and an abstract syntax tree for a small language.
6. Describe the benefits of having program representations other than strings of source code.
7. Implement a program to process some representation of code for some purpose, such as an interpreter, an expression optimizer, or a documentation generator.

FPL-Syntax: Syntax Analysis

Non-core:

1. Regular grammars vs context-free grammars. (See also: [AL-Models](#))
2. Scanning and parsing based on language specifications.
3. Lexical analysis using regular expressions.
4. Tokens and their use.
5. Parsing strategies including top-down (e.g., recursive descent, or LL) and bottom-up (e.g., LR or GLR) techniques.
 - a. Lookahead tables and their application to parsing.
6. Language theory:
 - a. Chomsky hierarchy. (See also: [AL-Models](#))
 - b. Left-most/right-most derivation and ambiguity.
 - c. Grammar transformation.
7. Parser error recovery mechanisms.
8. Generating scanners and parsers from declarative specifications.

Illustrative learning outcomes:

Non-core:

1. Use formal grammars to specify the syntax of languages.
2. Illustrate the role of lookahead tables in parsing.
3. Use declarative tools to generate parsers and scanners.
4. Recognize key issues in syntax definitions: ambiguity, associativity, precedence.

FPL-Semantics: Compiler Semantic Analysis

Non-core:

1. Abstract syntax trees; contrast with concrete syntax.
2. Defining, traversing and modifying high-level program representations.
3. Scope and binding resolution.
4. Static semantics:
 - a. Type checking.
 - b. Define before use.
 - c. Annotation and extended static checking frameworks.
5. L-values/R-values. (See also: [SDF-Fundamentals](#))
6. Call semantics.
7. Types of parameter-passing with simple illustrations and comparison: call by value, call by reference, call by value-result, call by name, call by need and their variations.
8. Declarative specifications such as attribute grammars and their applications in handling limited context-base grammar.

Illustrative learning outcomes:

Non-core:

1. Draw the abstract syntax tree for a small language.
2. Implement context-sensitive, source-level static analyses such as type-checkers or resolving identifiers to identify their binding occurrences.
3. Describe semantic analyses using an attribute grammar.

FPL-Analysis: Program Analysis and Analyzers

Non-core:

4. Relevant program representations, such as basic blocks, control-flow graphs, def-use chains, and static single assignment.
5. Undecidability and consequences for program analysis.
6. Flow-insensitive analysis, such as type-checking and scalable pointer and alias analysis.
7. Flow-sensitive analysis, such as forward and backward dataflow analyses.
8. Path-sensitive analysis, such as software model checking and software verification.
9. Tools and frameworks for implementing analyzers.
10. Role of static analysis in program optimization and data dependency analysis during exploitation of concurrency. (See also: [FPL-Code](#))
11. Role of program analysis in (partial) verification and bug-finding. (See also: [FPL-Code](#))
12. Parallelization:
 - a. Analysis for auto-parallelization.
 - b. Analysis for detecting concurrency bugs.

Illustrative learning outcomes:

Non-core:

1. Explain the difference between dataflow graph and control flow graph.
2. Explain why non-trivial sound program analyses must be approximate.
3. Argue why an analysis is correct (sound and terminating).
4. Explain why potential aliasing limits sound program analysis and how alias analysis can help.
5. Use the results of a program analysis for program optimization and/or partial program correctness.

FPL-Code: Code Generation

Non-core:

1. Instruction sets. (See also: [AR-Assembly](#))
2. Control flow.
3. Memory management. (See also: [AR-Memory](#), [OS-Memory](#))
4. Procedure calls and method dispatching.
5. Separate compilation; linking.
6. Instruction selection.
7. Instruction scheduling (e.g., pipelining).
8. Register allocation.
9. Code optimization as a form of program analysis. (See also: [FPL-Analysis](#))
10. Program generation through generative AI.

Illustrative learning outcomes:

Non-core:

1. Identify all essential steps for automatically converting source code into assembly or other low-level languages.
2. Explain the low-level code necessary for calling functions/methods in modern languages.
3. Discuss why separate compilation requires uniform calling conventions.

4. Discuss why separate compilation limits optimization because of unknown effects of calls.
5. Discuss opportunities for optimization introduced by naive translation and approaches for achieving optimization, such as instruction selection, instruction scheduling, register allocation, and peephole optimization.

FPL-Run-Time: Run-time Behavior and Systems

Non-core:

1. Process models using stacks and heaps to allocate and deallocate activation records and recovering environments using frame pointers and return addresses during a procedure call including parameter-passing examples.
2. Schematics of code lookup using hash tables for methods in implementations of object-oriented programs.
3. Data layout for objects and activation records.
4. Object allocation in heap.
5. Implementing virtual entities and virtual methods; virtual method tables and their application.
6. Run-time behavior of object-oriented programs.
7. Compare and contrast allocation of memory during information exchange using parameter-passing and non-local variables (using chain of static links).
8. Dynamic memory management approaches and techniques: malloc/free, garbage collection (mark-sweep, copying, reference counting), regions (also known as arenas or zones).
9. Just-in-time compilation and dynamic recompilation.
10. Interface to operating system (e.g., for program initialization).
11. Interoperability between programming languages including parameter-passing mechanisms and data representation: (See also: [AR-Representation](#))
 - a. Big endian, little endian.
 - b. Data layout of composite data types such as arrays.
12. Other common features of virtual machines, such as class loading, threads, and security checking.
13. Sandboxing.

Illustrative learning outcomes:

Non-core:

1. Discuss benefits and limitations of automatic memory management.
2. Explain the use of metadata in run-time representations of objects and activation records, such as class pointers, array lengths, return addresses, and frame pointers.
3. Compare and contrast static allocation vs stack-based allocation vs heap-based allocation of data elements.
4. Explain why some data elements cannot be automatically deallocated at the end of a procedure/method call (need for garbage collection).
5. Discuss advantages, disadvantages, and difficulties of just-in-time and dynamic recompilation.
6. Discuss use of sandboxing in mobile code.
7. Identify the services provided by modern language run-time systems.

FPL-Constructs: Advanced Programming Constructs

Non-core:

1. Encapsulation mechanisms.
2. Lazy evaluation and infinite streams.
3. Compare and contrast lazy evaluation vs eager evaluation.
4. Unification vs assertion vs expression evaluation.
5. Control abstractions: exception handling, continuations, monads.
6. Object-oriented abstractions: multiple inheritance, mixins, traits, multimethods.
7. Metaprogramming: macros, generative programming, model-based development.
8. String manipulation via pattern-matching (regular expressions).
9. Dynamic code evaluation ("eval").
10. Language support for checking assertions, invariants, and pre/post-conditions.
11. Domain specific languages, such as database languages, data science languages, embedded computing languages, synchronous languages, hardware interface languages.
12. Massive parallel high performance computing models and languages.

Illustrative learning outcomes:**Non-core:**

1. Use various advanced programming constructs and idioms correctly.
2. Discuss how various advanced programming constructs aim to improve program structure, software quality, and programmer productivity.
3. Discuss how various advanced programming constructs interact with the definition and implementation of other language features.

FPL-Pragmatics: Language Pragmatics**Non-core:**

1. Effect of technology needs and software requirements on programming language development and evolution.
2. Problems domains and programming paradigm.
3. Criteria for good programming language design:
 - a. Principles of language design such as orthogonality.
 - b. Defining control and iteration constructs.
 - c. Modularization of large software.
4. Evaluation order, precedence, and associativity.
5. Eager vs delayed evaluation.
6. Defining control and iteration constructs.
7. External calls and system libraries.

Illustrative learning outcomes:**Non-core:**

1. Discuss the role of concepts such as orthogonality and well-chosen defaults in language design.
2. Objectively evaluate and justify language-design decisions.
3. Implement an example program whose result can differ under different rules for evaluation order, precedence, or associativity.
4. Illustrate uses of delayed evaluation, such as user-defined control abstractions.

5. Discuss the need for allowing calls to external calls and system libraries and the consequences for language implementation.

FPL-Formalism: Formal Semantics

Non-core:

1. Syntax vs semantics.
2. Approaches to semantics: axiomatic, operational, denotational, type-based.
3. Axiomatic semantics of abstract constructs such as assignment, selection, iteration using pre-condition, post-conditions and loop invariant.
4. Operational semantics analysis of abstract constructs and sequence of such as assignment, expression evaluation, selection, iteration using environment and store:
 - a. Symbolic execution.
 - b. Constraint checkers.
5. Denotational semantics:
 - a. Lambda Calculus. (See also: [AL-Models](#), [FPL-Functional](#))
6. Proofs by induction over language semantics.
7. Formal definitions and proofs for type systems: (See also: [FPL-Types](#))
 - a. Propositions as types (implication as a function, conjunction as a product, disjunction as a sum).
 - b. Dependent types (universal quantification as dependent function, existential quantification as dependent product).
 - c. Parametricity.

Illustrative learning outcomes:

Non-core:

1. Construct formal semantics for a small language.
2. Write a lambda-calculus program and show its evaluation to a normal form.
3. Discuss the different approaches of operational, denotational, and axiomatic semantics.
4. Use induction to prove properties of all programs in a language.
5. Use induction to prove properties of all programs in a language that is well-typed according to a formally defined type system.
6. Use parametricity to establish the behavior of code given only its type.

FPL-Methodologies: Formal Development Methodologies

1. Formal specification languages and methodologies.
2. Theorem provers, proof assistants, and logics.
3. Constraint checkers. See also: [FPL-Formalism](#))
4. Dependent types (universal quantification as dependent function, existential quantification as dependent product). (See also: [FPL-Types](#), [FPL-Formalism](#))
5. Specification and proof discharge for fully verified software systems using pre/post conditions, refinement types, etc.
6. Formal modeling and manual refinement/implementation of software systems.
7. Use of symbolic testing and fuzzing in software development.
8. Model checking.

9. Understanding of situations where formal methods can be effectively applied and how to structure development to maximize their value.

Illustrative learning outcomes:

Non-core:

1. Use formal modeling techniques to develop and validate architectures.
2. Use proof assisted programming languages to develop fully specified and verified software artifacts.
3. Use verifier and specification support in programming languages to formally validate system properties.
4. Integrate symbolic validation tooling into a programming workflow.
5. Discuss when and how formal methods can be effectively used in the development process.

FPL-Design: Design Principles of Programming Languages

Non-core:

1. Language design principles:
 - a. Simplicity.
 - b. Security. (See also: [SEC-Coding](#))
 - c. Fast translation.
 - d. Efficient object code.
 - e. Orthogonality.
 - f. Readability.
 - g. Completeness.
 - h. Implementation strategies.
2. Designing a language to fit a specific domain or problem.
3. Interoperability between programming languages.
4. Language portability.
5. Formal description of a programming language.
6. Green computing principles. (See also: [SEP-Sustainability](#))

Illustrative Learning Outcomes:

Non-core:

1. Understand what constitutes good language design and apply that knowledge to evaluate a real programming language.

FPL-SEP: Society, Ethics, and the Profession

Non-core:

1. Impact of English-centric programming languages.
2. Enhancing accessibility and inclusivity for people with disabilities:
 - a. Supporting assistive technologies.
3. Human factors related to programming languages and usability:
 - a. Impact of syntax on accessibility.
 - b. Supporting cultural differences (e.g., currency, decimals, dates).
 - c. Neurodiversity.

4. Etymology of terms such as “class”, “master”, “slave” in programming languages.
5. Increasing accessibility by supporting multiple languages within applications (UTF).

Illustrative learning outcomes:

Non-core:

1. Consciously design programming languages to be inclusive and non-offensive.

Professional Dispositions

1. **Professional:** Students must demonstrate and apply the highest standards when using programming languages and formal methods to build safe systems that are fit for their purpose.
2. **Meticulous:** Attention to detail is essential when using programming languages and applying formal methods.
3. **Inventive:** Programming and approaches to formal proofs is inherently a creative process, students must demonstrate innovative approaches to problem solving. Students are accountable for their choices regarding the way a problem is solved.
4. **Proactive:** Programmers are responsible for anticipating all forms of user input and system behavior and to design solutions that address each one.
5. **Persistent:** Students must demonstrate perseverance since the correct approach is not always self-evident and a process of refinement may be necessary to reach the solution.

Mathematics Requirements

Required:

- Discrete Mathematics – Boolean algebra, proof techniques, digital logic, sets and set operations, mapping, functions and relations, states and invariants, graphs and relations, trees, counting, recurrence relations, finite state machine, regular grammar. (See also: [MSF-Discrete](#))
- Logic – propositional logic (negations, conjunctions, disjunctions, conditionals, biconditionals), first-order logic, logical reasoning (induction, deduction, abduction). (See also: [MSF-Discrete](#))
- Mathematics – Matrices, probability, statistics. (See also: [MSF-Probability](#), [MSF-Statistics](#))

Course Packaging Suggestions

Two example courses are presented illustrating how the content may be covered. The first is an introductory course which covers the CS Core and KA Core content. This course focuses on the different programming paradigms and ensures familiarity with each to a level sufficient to be able to decide which paradigm is appropriate in each circumstance.

The second course is an advanced course focused on the implementation of a programming language, the formal description of a programming language and a formal description of the behavior of a program.

While these two courses have been the predominant way to cover this knowledge area over the past decade, it is by no means the only way that this content can be covered. Institutions can, for example,

choose to cover only the CS Core content (24 hours) as part of one or spread over multiple courses (e.g. Software Engineering). Natural combinations are easily identifiable since they are the areas in which the Foundations of Programming Languages knowledge area overlaps with other knowledge areas. Such overlaps have been identified throughout this knowledge area.

Programming Language Concepts (Introduction) Course to include the following:

- [FPL-OOP](#): Object-Oriented Programming (6 hours)
- [FPL-Functional](#): Functional Programming (7 hours)
- [FPL-Logic](#): Logic Programming (3 hours)
- [FPL-Scripting](#): Shell Scripting (2 hours)
- [FPL-Event-Driven](#): Event-Driven and Reactive Programming (4 hours)
- [FPL-Parallel](#): Parallel and Distributed Computing (5 hours)
- [FPL-Types](#): Type Systems (6 hours)
- [FPL-Systems](#): Systems Execution and Memory Model (3 hours)
- [FPL-Translation](#): Language Translation and Execution (5 hours)
- [FPL-Abstraction](#): Program Abstraction and Representation (3 hours)
- [FPL-SEP](#): Society, Ethics, and the Profession (1 hour)

Prerequisites:

- Discrete Mathematics – Boolean algebra, proof techniques, digital logic, sets and set operations, mapping, functions and relations, states and invariants, graphs and relations, trees, counting, recurrence relations, finite state machine, regular grammar. (See also: [MSF-Discrete](#)).

Programming Language Implementation (Advanced) Course to include the following:

- [FPL-Types](#): Type Systems (3 hours)
- [FPL-Translation](#): Language Translation and Execution (2 hours)
- [FPL-Syntax](#): Syntax Analysis(3 hours)
- [FPL-Semantics](#): Compiler Semantic Analysis (5 hours)
- [FPL-Analysis](#): Program Analysis and Analyzers (5 hours)
- [FPL-Code](#): Code Generation(5 hours)
- [FPL-Run-Time](#): Run-time Systems (4 hours)
- [FPL-Constructs](#): Advanced Programming Constructs (4 hours)
- [FPL-Pragmatics](#): Language Pragmatics (3 hours)
- [FPL-Formalism](#): Formal Semantics (5 hours)
- [FPL-Methodologies](#): Formal Development Methodologies (5 hours)

Prerequisites:

- Discrete mathematics – Boolean algebra, proof techniques, digital logic, sets and set operations, mapping, functions and relations, states and invariants, graphs and relations, trees, counting, recurrence relations, finite state machine, regular grammar. (See also: [MSF-Discrete](#)).
- Logic – propositional logic (negations, conjunctions, disjunctions, conditionals, biconditionals), first-order logic, logical reasoning (induction, deduction, abduction). (See also: [MSF-Discrete](#)).
- Introductory programming course. (See also: [SDF-Fundamentals](#)).
- Programming proficiency in programming concepts such as: (See also: [SDF-Fundamentals](#)):

- Type declarations such as basic data types, records, indexed data elements such as arrays and vectors, and class/subclass declarations, types of variables.
- Scope rules of variables,
- Selection and iteration concepts, function and procedure calls, methods, object creation.
- Data structure concepts such as: (See also: [SDF-DataStructures](#)):
 - Abstract data types, sequence and string, stack, queues, trees, dictionaries. (See also: [SDF-Data-Structures](#))
 - Pointer-based data structures such as linked lists, trees and shared memory locations. (See also: [SDF-Data-Structures](#), [AL-Foundational](#))
 - Hashing and hash tables. (See also: [SDF-Data-Structures](#), [AL-Foundational](#))
- System fundamentals and computer architecture concepts such as (See also: [SF-Foundations](#)):
 - Digital circuits design, clocks, bus. (See also: [OS-Principles](#))
 - registers, cache, RAM and secondary memory. (See also: [OS-Memory](#))
 - CPU and GPU. (See also: [AR-Heterogeneity](#))
- Basic knowledge of operating system concepts such as:
 - Interrupts, threads and interrupt-based/thread-based programming. (See also: [OS-Concurrency](#))
 - Scheduling, including prioritization. (See also: [OS-Scheduling](#))
 - Memory fragmentation. (See also: [OS-Memory](#))
 - Latency.

Committee

Chair: Michael Oudshoorn, High Point University, High Point, NC, USA

Members:

- Annette Bieniusa, TU Kaiserslautern, Kaiserslautern, Germany
- Brijesh Dongol, University of Surrey, Guildford, UK
- Michelle Kuttel, University of Cape Town, Cape Town, South Africa
- Doug Lea, State University of New York at Oswego, Oswego, NY, USA
- James Noble, Victoria University of Wellington, Wellington, New Zealand
- Mark Marron, Microsoft Research, Seattle, WA, USA and University of Kentucky, Lexington, KY, USA
- Peter-Michael Osera, Grinnell College, Grinnell, IA, USA
- Michelle Mills Strout, University of Arizona, Tucson, AZ, USA

Contributors:

- Alan Dearle, University of St. Andrews, St. Andrews, Scotland

Graphics and Interactive Techniques (GIT)

Preamble

Computer graphics is the term used to describe the computer generation and manipulation of images and can be viewed as the science of enabling visual communication through computation. Its application domains include animation, Computer Generated Imagery (CGI) and Visual Effects (VFX); engineering; machine learning; medical imaging; scientific, information, and knowledge visualization; simulators; special effects; user interfaces; and video games. Traditionally, graphics at the undergraduate level focused on rendering, linear algebra, physics, the graphics pipeline, interaction, and phenomenological approaches. Today's graphics courses increasingly include data science, physical computing, animation, and haptics. Thus the knowledge area (KA) expanded beyond core image-based computer graphics. At the advanced level, undergraduate institutions are more likely to offer one or several courses specializing in a specific graphics knowledge unit (KU) or topic: e.g., gaming, animation, visualization, tangible or physical computing, and immersive courses such as Augmented Reality (AR)/Virtual Reality (VR)/eXtended Reality (XR). There is considerable connection with other computer science knowledge areas (KAs): Algorithmic Foundations, Architecture and Organization, Artificial Intelligence; Human-Computer Interaction; Parallel and Distributed Computing; Specialized Platform Development; Software Engineering, and Society, Ethics, and the Profession.

In order for students to become adept at the use and generation of computer graphics and interactive techniques, many issues must be addressed, such as human perception and cognition, data and image file formats, display specifications and protocols, hardware interfaces, and application program interfaces (APIs). Unlike other knowledge areas, KUs within Graphics and Interactive Techniques may be included in a variety of elective courses. Alternatively, graphics topics may be introduced in an applied project in courses primarily covering human computer interaction, embedded systems, web development, introductory programming courses, etc. Undergraduate computer science students who study the KUs specified below through a balance of theory and applied instruction will be able to understand, evaluate, and/or implement the related graphics and interactive techniques as users and developers. Because technology changes rapidly, the Graphics and Interactive Techniques subcommittee attempted to avoid being overly prescriptive. Any examples of APIs, programs, and languages should be considered as appropriate examples in 2023. In effect, this is a snapshot in time.

Graphics as a KA has expanded and become pervasive since the CS2013 report. AR/VR/XR, artificial intelligence, computer vision, data science, machine learning, and interfaces driven by embedded sensors in everything from cars to coffee makers use graphics and interactive techniques. The now ubiquitous smartphone has made the majority of the world's population regular users and creators of graphics, digital images, and the interactive techniques to manipulate them. Animations, games, visualizations, and immersive applications that ran on desktops in 2013, now can run on mobile devices. The amount of stored digital data grew exponentially since 2013, and both data and visualizations are now published by myriad sources including news media and scientific organizations. Revenue from mobile video games now exceeds that of music and movies combined [39]. CGI and VFX are employed in almost all films, animations, TV productions, advertising, and business graphics.

The number of people who create graphics have skyrocketed, as have the number of applications and generative tools used to produce graphics.

It is critical that students and faculty confront the ethical issues, questions, and conundrums that have arisen and will continue to arise in and because of applications in computer graphics. Today's headlines unfortunately already provide examples of inequity and/or wrong-doing in autonomous navigation, deepfakes, computational photography, generative images, and facial recognition.

Overview of Knowledge Units

The following knowledge units are included in Graphics and Interactive Techniques. Descriptions are included below where they are not explicitly evident from the title. Graphics as a KA is unique in that many of its KUs can and are taught as stand-alone courses where implementation projects are critical to student mastery. Apart from Applied Rendering and Techniques which scaffolds the typical undergraduate interactive computer graphics course, the other KU's are more specialized. To be consistent with the design of CS2023 those KUs are limited to two weeks of instruction, corresponding roughly to 6 hours of instruction. This limitation allows a two-week KU to be added to a course. Due to that time restriction, we list their Core Topic skill levels in most of the KU's as "Explain." However, if one of the thematic KUs is implemented as a full term course, our expectation is that the skill levels will rise to "Apply" or "Develop" which is reflected in many of the practical illustrative learning outcomes. This is not meant to be prescriptive but to encourage customization. How to implement a KU is left to the discretion of the instructor. If given a two-week constraint to teach one of the thematic KUs, many of us would choose to limit the topics and include an applied project. Our hope is that an inclusive list of topics will help faculty design a course that best meets their and their department's pedagogical goals.

- **GIT-Fundamentals: Fundamental Concepts.** For nearly every computer scientist and software developer, understanding of how humans interact with machines is essential.
- **GIT-Visualization: Visualization.** Visualization seeks to determine and present underlying correlated structures and relationships in data sets from a wide variety of application areas. The prime objective is to communicate the information in a way which enhances understanding.
- **GIT-Rendering: Applied Rendering and Techniques.** This unit includes basic rendering and fundamental graphics techniques that nearly every undergraduate course in graphics will cover and that are essential for further study in most graphics-related courses.
- **GIT-Modeling: Geometric Modeling.** Graphics must be encoded in computer memory, often in the form of a mathematical specification of shape and form.
- **GIT-Shading: Shading and Advanced Rendering.** This unit contains more in-depth coverage of rendering topics.
- **GIT-Animation: Computer Animation.** Computer Animation is concerned with the generation of moving imagery.
- **GIT-Simulation: Simulation.** Simulation has strong ties to Computational Science. However, in the graphics domain, simulation techniques are re-purposed to a different end. Rather than creating predictive models, the goal instead is to achieve a mixture of physical plausibility and artistic intention. To illustrate, the goals of "model surface tension in a liquid" and "produce a

“crown splash” are related, but different. Depending on the simulation goals, covered topics may vary as shown.

- Particle systems
 - Integration methods (Forward Euler, Midpoint, Leapfrog)
- Rigid Body Dynamics
 - Particle systems
 - Collision Detection
 - Triangle/point
 - Edge/edge
- Cloth
 - Particle systems
 - Mass/spring networks
 - Collision Detection
- Particle-Based Water
 - Integration methods
 - Smoother Particle Hydrodynamics (SPH) Kernels
 - Signed Distance Function-Based Collisions
- Grid-Based Smoke and Fire
 - Semi-Lagrangian Advection
 - Pressure Projection
- Grid and Particle-Based Water
 - Particle-Based Water
- Grid-Based Smoke and Fire
 - Semi-Lagrangian Advection
 - Pressure Projection
- Grid and Particle-Based Water
 - Particle-Based Water
 - Grid-Based Smoke, and Fire
- **GIT-Immersion: Immersion.** Immersion includes Augmented Reality (AR), Virtual Reality (VR), and Mixed Reality (MR).
- **GIT-Interaction: Interaction.** Interactive computer graphics is a requisite part of real-time applications ranging from the utilitarian-like word processors to virtual and/or augmented reality applications.
- **GIT-Image: Image Processing.** Image Processing consists of the analysis and processing of images for multiple purposes, but most frequently to improve image quality and to manipulate imagery. It lies at the cornerstone of computer vision.
- **GIT-Physical: Tangible/Physical Computing.** Tangible/Physical Computing refers to microcontroller-based interactive systems that detect and respond to sensor input.
- **GIT-SEP: Society, Ethics and the Profession.**

Changes since CS2013

In an effort to align CS2013's Graphics and Visualizations areas with ACM Special Interest Group on Graphic and Interactive Techniques (SIGGRAPH) and to reflect the natural expansion of the field to include haptic and physical computing in addition to images, we have renamed it Graphics and Interactive Techniques (GIT). To capture the expanded footprint of the KA, the following KUs have been added to the original list consisting of Fundamental Concepts, Visualization, Basic Rendering (renamed Rendering), Geometric Modeling, Advanced Rendering (renamed Shading), and Computer Animation:

- Immersion (MR, AR, VR)
- Interaction
- Image Processing
- Tangible/Physical Computing
- Simulation

Core Hours

Knowledge Unit	CS Core	KA Core
Fundamental Concepts	4	3
Visualization		6
Applied Rendering and Techniques		15
Geometric Modeling		6
Shading and Advanced Rendering		6
Computer Animation		6
Simulation		6
Immersion (MR, AR, VR)		6
Interaction		4
Image Processing		6
Tangible/Physical Computing		6
Society, Ethics, and the Profession	Included in SEP hours	
Total	4	70

Knowledge Units

GIT-Fundamentals: Fundamental Concepts

CS Core:

1. Uses of computer graphics and interactive techniques and their potential risks and abuses.
 - a. Entertainment, business, and scientific applications: e.g., visual effects, generative imagery, computer vision, machine learning, user interfaces, video editing, games and game engines, computer-aided design and manufacturing, data visualization, and virtual/augmented/mixed reality.
 - b. Intellectual property, deep fakes, facial recognition, privacy. (See also: [SEP-DEIA](#), [SEP-Privacy](#), [SEP-IP](#), [SEP-Professional-Ethics](#))
2. Graphic output
 - a. Displays (e.g., LCD)
 - b. Printers
 - c. Analog film
 - d. Concepts
 - i. Resolution (e.g., pixels, dots)
 - ii. Aspect ratio
 - iii. Frame rate
3. Human vision system
 - a. Tristimulus reception (RGB)
 - b. Eye as a camera (projection)
 - c. Persistence of vision (frame rate, motion blur)
 - d. Contrast (detection, Mach banding, dithering/aliasing)
 - e. Non-linear response (dynamic range, tone mapping)
 - f. Binocular vision (stereo)
 - g. Accessibility (color deficiency, strobing, monocular vision, etc.) (See also: [SEP-DEIA](#), [HCI-User](#))
4. Standard image formats
 - a. Raster
 - i. Lossless (e.g., TIF)
 - ii. Lossy (e.g., JPG, GIF, etc.)
 - b. Vector (e.g., SVG, Adobe Illustrator)
5. Digitization of analog data
 - a. Rasterization
 - b. Resolution
 - c. Sampling and quantization
6. Color models: additive (RGB), subtractive (CMYK), and color perception (HSV)
7. Tradeoffs between storing image data and re-computing image data
8. Spatialization: coordinate systems, absolute and relative positioning
9. Animation as a sequence of still images

KA Core:

10. Applied interactive graphics (e.g., processing, python,)
11. Display characteristics (protocols and ports)

Illustrative Learning Outcomes:

CS Core:

1. Identify common uses of digital presentation to humans (e.g., computer graphics, sound).
2. Describe how analog signals can be reasonably represented by discrete samples, for example, how images can be represented by pixels.
3. Compute the memory requirement for storing a color image given its resolution.
4. Create a graphic depicting how the limits of human perception affect choices about the digital representation of analog signals.
5. Indicate when and why you should use each of the following common file formats: JPG, PNG, MP3, MP4, and GIF?
6. Describe color models and their use in graphics display devices.
7. Compute the memory requirements for a multi-second movie (lasting n seconds) displaying at a specific framerate (f frames per second) at a specified resolution (r pixels per frame)
8. Compare and contrast digital video to analog video.
9. Describe the basic process of producing continuous motion from a sequence of discrete frames (sometimes called “flicker fusion”).
10. Describe a possible visual misrepresentation that could result from digitally sampling an analog world.
11. Compute memory space requirements based on resolution and color coding.
12. Compute time requirements based on refresh rates and rasterization techniques.

KA Core:

13. Design a user interface and an alternative for persons with color perception deficiency.
14. Construct a simple graphical user interface using a graphics library.

GIT-Visualization: Visualization

KA Core:

1. Scientific Data Visualization and Information Visualization
2. Visualization techniques
 - a. Statistical visualization (e.g., scatterplots, bar graphs, histograms, line graphs, pie charts, trees and graphs)
 - b. Text visualization
 - c. Geospatial visualization
 - d. 2D/3D scalar fields
 - e. Vector fields
 - f. Direct volume rendering
3. Visualization pipeline
 - a. Structuring data
 - b. Mapping data to visual representations (e.g., scales, grammar of graphics)
 - c. View transformations (e.g., pan, zoom, filter, select)

4. Common data formats (e.g., HDF, netCDF, geotiff, GeoJSON, shape files, raw binary, JSON, CSV, plain text)
5. High-dimensional data handling techniques
 - a. Statistical (e.g., averaging, clustering, filtering)
 - b. Perceptual (e.g., multi-dimensional vis, parallel coordinates, trellis plots)
6. Perceptual and cognitive foundations that drive visual abstractions.
 - a. Human optical system
 - b. Color theory
 - c. Gestalt theories
7. Design and evaluation of visualizations
 - a. Purpose (e.g., analysis, communication, aesthetics)
 - b. Accessibility
 - c. Appropriateness of encodings
 - d. Misleading visualizations

Illustrative Learning Outcomes:

KA Core:

1. Compare and contrast data visualization and information visualization.
2. Deploy basic algorithms for visualization.
3. Compare the tradeoffs of visualization algorithms in terms of accuracy and performance.
4. Design a suitable visualization for a particular combination of data characteristics, application tasks, and audience.
5. Analyze the effectiveness of a given visualization for a particular task.
6. Design a process to evaluate the utility of a visualization algorithm or system.
7. Identify a variety of applications of visualization including representations of scientific, medical, and mathematical data; flow visualization; and spatial analysis.

GIT-Rendering: Applied Rendering and Techniques

KA Core: (See also: [SPD-Game](#))

1. Object and scene modeling
 - a. Object representations: polygonal, parametric, etc.
 - b. Modeling transformations: affine and coordinate-system transformations
 - c. Scene representations: scene graphs
2. Camera and projection modeling
 - a. Pinhole cameras, similar triangles, and projection model
 - b. Camera models
 - c. Projective geometry
3. Radiometry and light models
 - a. Radiometry
 - b. Rendering equation
 - c. Rendering in nature – emission and scattering, etc.
4. Rendering
 - a. Simple triangle rasterization
 - b. Rendering with a shader-based API

- c. Visibility and occlusion, including solutions to this problem (e.g., depth buffering, Painter's algorithm, and ray tracing)
- d. Texture mapping, including minification and magnification (e.g., trilinear MIP mapping)
- e. Application of spatial data structures to rendering.
- f. Ray tracing
- g. Sampling and anti-aliasing

Illustrative Learning Outcomes:

KA Core:

1. Describe and illustrate the light transport problem (i.e., light is emitted, scatters around the scene, and is measured by the eye).
2. Describe the basic rendering pipeline.
3. Compare and contrast how forward and backwards rendering factor into the graphics pipeline.
4. Create a program to display 2D shapes in a window.
5. Create a program to display 3D models.
6. Produce linear perspective from similar triangles by converting points (x, y, z) to points $(x/z, y/z, 1)$.
7. Compute two-dimensional and 3-dimensional points by applying affine transformations.
8. Indicate the changes required to extend 2D transformation operations to handle transformations in 3D.
9. Define texture mapping, sampling, and anti-aliasing, and describe examples of each.
10. Compare ray tracing and rasterization for the visibility problem.
11. Construct a program that performs transformation and clipping operations on simple two-dimensional shapes.
12. Implement a simple real-time renderer using a rasterization API (e.g., OpenGL, WebGL) using vertex buffers and shaders.
13. Compare and contrast the different rendering techniques.
14. Compare and contrast the difference in transforming the camera vs the models.

GIT-Modeling: Geometric Modeling

KA Core:

1. Basic geometric operations such as intersection calculation and proximity tests on 2D objects
2. Surface representation/model
 - a. Tessellation
 - b. Mesh representation, mesh fairing, and mesh generation techniques such as Delaunay triangulation, and marching cubes/tetrahedrons
 - c. Parametric polynomial curves and surfaces
 - d. Implicit representation of curves and surfaces
 - e. Spatial subdivision techniques
3. Volumetric representation/model
 - a. Volumes, voxels, and point-based representations.
 - b. Signed Distance Fields
 - c. Sparse Volumes, i.e., VDB
 - d. Constructive Solid Geometry (CSG) representation

4. Procedural representation/model
 - a. Fractals
 - b. L-Systems
5. Multi-resolution modeling (See also: [SPD-Game](#))
6. Reconstruction, e.g., 3D scanning, photogrammetry, etc.

Illustrative Learning Outcomes:

KA Core:

1. Contrast representing curves and surfaces in both implicit and parametric forms.
2. Create simple polyhedral models by surface tessellation.
3. Create a mesh representation from an implicit surface.
4. Create a fractal model or terrain using a procedural method.
5. Create a mesh from data points acquired with a laser scanner.
6. Create CSG models from simple primitives, such as cubes and quadric surfaces.
7. Contrast modeling approaches with respect to space and time complexity and quality of image.

GIT-Shading: Shading and Advanced Rendering

KA Core:

1. Solutions and approximations to the rendering equation, for example:
 - a. Distribution ray tracing and path tracing
 - b. Photon mapping
 - c. Bidirectional path tracing
 - d. Metropolis light transport
2. Time (motion blur), lens position (focus), and continuous frequency (color) and their impact on rendering
3. Shadow mapping
4. Occlusion culling
5. Bidirectional Scattering Distribution function (BSDF) theory and microfacets
6. Subsurface scattering
7. Area light sources
8. Hierarchical depth buffering
9. Image-based rendering
10. Non-photorealistic rendering
11. Realtime rendering
12. GPU architecture (See also: [AR-Heterogeneity](#))
13. Human visual systems including adaptation to light, sensitivity to noise, and flicker fusion (See also: [HCI-Accessibility](#), [SEP-DEIA](#))

Illustrative Learning Outcomes:

KA Core:

1. Describe how an algorithm estimates a solution to the rendering equation.
2. Discuss the properties of a rendering algorithm (e.g., complete, consistent, and unbiased).
3. Analyze the bandwidth and computation demands of a simple shading algorithm.

4. Implement a non-trivial shading algorithm (e.g., toon shading, cascaded shadow maps) under a rasterization API.
5. State how a particular artistic technique might be implemented in a renderer.
6. Describe how one might recognize the shading techniques used to create a particular image.
7. Write a program that implements any of the specified graphics techniques using a primitive graphics system at the individual pixel level.
8. Write a ray tracer for scenes using a simple (e.g., Phong's) Bidirectional Reflection Distribution Function (BRDF) plus reflection and refraction.

GIT-Animation: Computer Animation

KA Core:

1. Principles of Animation: Squash and Stretch, Timing, Anticipation, Staging, Follow Through and Overlapping Action, Straight Ahead Action and Pose-to-Pose Action, Slow In and Out, Arcs, Exaggeration, and Appeal
2. Types of animation
 - a. 2- and 3-dimensional animation
 - b. Motion graphics
 - c. Motion capture
 - d. Motion graphics
 - e. Stop animation
3. Key-frame animation
 - a. Keyframe Interpolation Methods: Lerp/Slerp/Spline
4. Forward and inverse kinematics (See also: [SPD-Robot](#), [AI-Robotics](#))
5. Skinning algorithms
 - a. Capturing
 - b. Linear blend, dual quaternion
 - c. Rigging
 - d. Blend shapes
 - e. Pose space deformation
6. Motion capture
 - a. Set up and fundamentals
 - b. Blending motion capture clips
 - c. Blending motion capture and keyframe animation
 - d. Ethical considerations (See also: [SEP-DEIA](#), [SEP-Privacy](#))
 - i. Avoidance of "default" captures - there is no typical human walk cycle.
 - ii. Accessibility

Illustrative Learning Outcomes:

KA Core:

1. Using a simple open-source character model and rig, describe visually why each of the principles of animation is fundamental to realistic animation.
2. Compute the location and orientation of model parts using a forward kinematic approach.
3. Compute the orientation of articulated parts of a model from a location and orientation using an inverse kinematic approach.

4. Compare the tradeoffs in different representations of rotations.
5. Write a script that implements the spline interpolation method for producing in-between positions and orientations.
6. Deploy off-the-shelf animation software to construct, rig, and animate simple organic forms.

GIT-Simulation: Simulation

KA Core:

1. Collision detection and response
 - a. Signed Distance Fields
 - b. Sphere/sphere
 - c. Triangle/point
 - d. Edge/edge
2. Procedural animation using noise
3. Particle systems
 - a. Integration methods (e.g., forward Euler, midpoint, leapfrog)
 - b. Mass/spring networks
 - c. Position-based dynamics
 - d. Rules (e.g., boids, crowds)
 - e. Rigid bodies
4. Grid-based fluids
 - a. Semi-Lagrangian advection
 - b. Pressure projection
5. Heightfields
 - a. Terrain: transport, erosion
 - b. Water: ripple, shallow water.
6. Rule-based systems (e.g., L-systems, space-colonizing systems, Game of Life)

Illustrative Learning Outcomes:

KA Core:

1. Implement algorithms for physical modeling of particle dynamics using simple Newtonian mechanics (e.g., Witkin & Kass, snakes and worms, symplectic Euler, Stormer/Verlet, or midpoint Euler methods)
2. Contrast the basic ideas behind fluid simulation methods for modeling ballistic trajectories (e.g., for splashes, dust, fire, or smoke).
3. Implement a smoke solver with user interaction.

GIT-Immersion: Immersion

KA Core: (See also: [SPD-Game](#), [SPD-Mobile](#), [HCI-Design](#))

1. Immersion levels (i.e., Virtual Reality (VR), Augmented Reality (AR), and Mixed Reality (MR))
2. Definitions of and distinctions between immersion and presence
3. 360 Video
4. Stereoscopic display
 - a. Head-mounted displays

- b. Stereo glasses
- 5. Viewer tracking
 - a. Inside out and outside In
 - b. Head/Body/Hand/tracking
- 6. Time-critical rendering to achieve optimal Motion To Photon (MTP) latency
 - a. Multiple Levels Of Details (LOD)
 - b. Image-based VR
 - c. Branching movies
- 7. Distributed VR, collaboration over computer network
- 8. Presence and factors that impact level of immersion
- 9. 3D interaction
- 10. Applications in medicine, simulation, training, and visualization
- 11. Safety in immersive applications
 - a. Motion sickness
 - b. VR obscures the real world, which increases the potential for falls and physical accidents

Illustrative Learning Outcomes:

KA Core:

1. Create a stereoscopic image.
2. Design and write an AR or VR application.
3. Summarize the pros and cons of different types of viewer tracking.
4. Compare and contrast the differences between geometry- and image-based virtual reality.
5. Analyze the design issues of user action synchronization and data consistency in a networked environment.
6. Create the specifications for an augmented reality application to be used by surgeons in the operating room.
7. Assess an immersive application's accessibility (See also: [HCI-Accessibility](#), [SEP-DEIA](#))
8. Identify the most important technical characteristics of a VR system/application that should be controlled to avoid motion sickness and explain why.

GIT-Interaction: Interaction

KA Core:

1. Event Driven Programming (See also: [FPL-Event-Driven](#))
 - a. Mouse or touch events
 - b. Keyboard events
 - c. Voice input
 - d. Sensors
 - e. Message passing communication
 - f. Network events
2. Graphical User Interface (Single Channel)
 - a. Window
 - b. Icons
 - c. Menus
 - d. Pointing Devices

3. Accessibility (See also: [SEP-DEIA](#))

Non-core:

4. Gestural Interfaces (See also: [SPD-Game](#))
 - a. Touch screen gestures
 - b. Hand and body gestures
5. Haptic Interfaces
 - a. External actuators
 - b. Gloves
 - c. Exoskeletons
6. Multimodal Interfaces
7. Head-worn Interfaces
 - a. Brain-computer interfaces, e.g., ElectroEncephaloGraphy (EEG) electrodes and Multi-Electrode Arrays (MEAs)
 - b. Headsets with embedded eye tracking
 - c. AR glasses
8. Natural Language Interfaces (See also: [AI-NLP](#))

Illustrative Learning Outcomes:

KA Core:

1. Create a simple game that responds to single channel mouse and keyboard events.
2. Create a mobile app that responds to touch events.
3. Design and create an application that responds to different event triggers.

None-core:

4. Assess the consistency or lack of consistency in cross-platform touch screen gestures.
5. Design and create an application that provides haptic feedback.
6. Write a program that is controlled by gestures.

GIT-Image: Image Processing

KA Core: (See also: [AI-Vision](#))

1. Morphological operations
 - a. Connected components
 - b. Dilation
 - c. Erosion
 - d. Computing region properties (area, perimeter, centroid, etc.)
2. Color histograms
 - a. Representation
 - b. Contrast enhancement through normalization
3. Image enhancement
 - a. Convolution
 - b. Blur (e.g., Gaussian)
 - c. Sharpen (e.g., Laplacian)
 - d. Frequency filtering (e.g., low-pass, high-pass)

4. Image restoration
 - a. Noise, degradation
 - b. Inpainting and other completion algorithms
 - c. Wiener filter
5. Image coding
 - a. Redundancy
 - b. Compression (e.g., Huffman coding)
 - c. Discrete Cosine Transform (DCT), wavelet transform, Fourier transforms (See also: [SPD-Interactive](#))
 - d. Nyquist Theorem
 - e. Watermarks
6. Connections to deep learning (e.g., Convolutional Neural Networks) (See also: [AI-ML](#))

Illustrative Learning Outcomes:

KA Core:

1. Write a program that uses dilation and erosion to smooth the edges of a binary image.
2. Manipulating the hue of an image.
3. Write a program that applies a high-pass filter to an image. (The advanced variation would be to filter an image using a high-pass filter in the frequency domain.)
4. Write a program that restores missing parts of an image using an in-paint algorithm (e.g., Poisson image editing)
5. Assess the results of selectively filtering an image in the frequency domain.

GIT-Physical: Tangible/Physical Computing

KA Core:

1. Interaction with the physical world (See also: [SPD-Embedded](#))
 - a. Acquisition of data from sensors
 - b. Driving external actuators
2. Connection to physical artifacts
 - a. Computer-Aided Design (CAD)
 - b. Computer-Aided Manufacturing (CAM)
 - c. Fabrication (See also: [HCI-Design](#))
 - i. Prototyping
 - ii. Additive (3D printing)
 - iii. Subtractive (Computer Numerical Control (CNC) milling)
 - iv. Forming (vacuum forming)
3. Internet of Things (See also: [SPD-Interactive](#))
 - a. Network connectivity
 - b. Wireless communication

Illustrative Learning Outcomes:

KA Core:

1. Construct a simple virtual switch or application button and use it to turn on an LED.
2. Construct a simple system to move a servo in response to sensor data.

3. Create a circuit and accompanying microcontroller code that uses a light sensor to vary a property of something else (e.g., color or brightness of an LED or graphic, position of an external actuator).
4. Create a circuit with a variable resistor and write a microcontroller program that reads and responds to the resistor's changing values.
5. Create a 3D form in a CAD package.
 - a. Show how affine transformations are achieved in the CAD program.
 - b. Show an example of instances of an object.
 - c. Create a fabrication plan. Provide a cost estimate for materials and time. How will you fabricate it?
 - d. Fabricate it. How closely did your actual fabrication process match your plan? Where did it differ?
6. Write the G- and M-Code to construct a 3D maze, and use a CAD/CAM package to check your work.
7. Decide and defend your decision to use Ethernet, WiFi, Bluetooth, RFID/NFC, or something else for Internet connectivity when designing an IoT pill dispenser. Create an IoT pill dispenser.
8. Distinguish between the different types of fabrication and describe when you would use each.

GIT-SEP: Society, Ethics, and the Profession

KA Core:

1. Accessibility in immersive, interactive, and physical computing applications (See also: [SEP-DEIA](#))
 - a. Accessible to people with mobility impairments
 - b. Accessible to people with vision and/or hearing impairments
2. Ethics/privacy in graphics applications. (See also: [SEP-Privacy](#), [SEP-Professional-Ethics](#), and [SEP-Security](#))
 - a. Acquisition of private data (room scans, body proportions, active cameras, etc.)
 - b. Can't look away from immersive applications easily.
 - c. Danger to self/surroundings while immersed
 - d. Ethical pitfalls of facial recognition
 - e. Misleading visualizations
 - i. Due to incorrect data because of exaggeration, hole filling, smoothing, data cleanup, etc.
 - ii. Even correct data can mislead (e.g., aliasing can cause back moving or stopped fan blades)
 - f. Privacy regarding health and other personal information
 - g. Bias in image processing
 - i. Deep fakes
 - ii. Applications that misidentify people based on skin color or hairstyle
3. Intellectual Property law as it relates to computer graphics and interactive techniques (See also: [SEP-IP](#))
 - a. images used to train generative AI
 - b. images produced by generative AI
4. Current and past contributors to the field (See also: [SEP-DEIA](#))

Illustrative Learning Outcomes:

KA Core:

1. Discuss the security issues inherent in location tags.
2. Describe the ethical pitfalls of facial recognition. Can facial recognition be used ethically? If so, how?
3. Discuss the copyright issues of using watermarked images to train a neural network.

Professional Dispositions

- **Self-directed:** Graphics hardware and software evolves rapidly. Students need to understand the importance of being a life-long learner.
- **Collaborative:** Graphics developers typically work in diverse teams composed of people with disparate subject matter expertise. Students should understand the value of being a good team member, and their teamwork skills should be cultivated and evaluated with constructive feedback.
- **Effective communicator:** Communication is critical. Students technical communication—verbal, written, and in code—should be practiced and evaluated.
- **Creative:** Creative problem solving lies at the core of computer graphics.

Mathematics Requirements

Required:

1. Coordinate geometry
2. Trigonometry
3. [MSF-Linear](#)*
 - a. Points (coordinate systems & homogeneous coordinates), vectors, and matrices
 - b. Vector operations: addition, scaling, dot and cross products
 - c. Matrix operations: addition, multiplication, determinants
 - d. Affine transformations
4. [MSF-Calculus](#)*
 - a. Continuity

*Note, if students enroll in a graphics class without linear algebra or calculus, graphics faculty can teach what is needed. To wit, many graphics textbooks cover the requisite mathematics in the appendix.

Desirable:

1. [MSF-Linear](#)
 - a. Eigenvectors and Eigen decomposition
 - b. Gaussian elimination and lower upper factorization
 - c. Singular value decomposition
2. [MSF-Calculus](#)
 - a. Quaternions
 - b. Differentiation
 - c. Vector calculus
 - d. Tensors
 - e. Differential geometry
3. [MSF-Probability](#)

4. [MSF-Statistics](#)
5. [MSF-Discrete](#)
 - a. Numerical methods for simulation

Necessary and Desirable Data Structures:

1. Data Structures necessary for this knowledge area ([AL-Foundational](#), [SDF-Algorithms](#), [SDF-Data-Structures](#))
 - a. Directed Acyclic Graphs
 - b. Tuples (points / vectors / matrices of fixed dimension)
 - c. Dense 1D, 2D, 3D arrays.
2. Data Structures desirable for this knowledge area ([AL-Foundational](#), [SDF-Algorithms](#), [SDF-Data-Structures](#), [SDF-Practices](#))
 - a. Array Structures and Structure of Arrays
 - b. Trees (e.g., K-trees, quadtrees, Huffman Trees)

Course Packaging Suggestions

Interactive Computer Graphics to include the following:

- [GIT-Rendering](#) 20 hours
- [GIT-Modeling](#) 6 hours
- [GIT-Interaction](#) 4 hours
- [SEP-Professional-Ethics](#), [SEP-DEIA](#) 3 hours

Prerequisites:

- [AL-Foundational](#)
- [AL-Strategies](#)
- [SDF-Algorithms](#)
- [SDF-Data-Structures](#)
- [SDF-Practices](#)
- [MSF-Linear](#) as a prerequisite or cover relevant topics in class

Course objectives: Students should understand and be able to create basic computer graphics using an API. They should know how to position and orient models, the camera, and distant and local lights.

Note: depending on the instructor, this course can be customized to include topics from another graphics knowledge unit, for example a two week unit on image processing or advanced rendering.

Media Computation to include the following:

- [GIT-Fundamentals](#) (4 hours)
- [GIT-Rendering](#) (6 hours)
- [GIT-Interaction](#) (3 hours)
- [SDF-Fundamentals](#) (10 hours)
- [AL-Foundational](#) (7 hours)
- [HCI-User](#) (5 hours)
- [GIT-SEP](#) (4 hours)

Course objectives: In this introductory programming class, students should be able to explain, evaluate, and apply algorithms and arrays that use and produce digital media.

User-Centered Development to include the following:

- [GIT-Fundamentals](#) (4 hours)
- [GIT-Rendering](#) (6 hours)
- [GIT-Interaction](#) (3 hours)
- [HCI-User](#) (8 hours)
- [HCI-Accessibility](#) (3 hours)
- [HCI-SEP](#) (4 hours)
- [SE-Construction](#) (4 hours)
- [SPD-Web](#), [SPD-Game](#), [SPD-Mobile](#) (8 hours)

Students should be able to develop applications that are usable and useful for people. Graphical user interface (GUI) designs will be implemented and analyzed using rapid prototyping.

Tangible Computing to include the following:

- [GIT-Physical](#) (14 hours)
- [GIT-Interaction](#) (4 hours)
- [SPD-Embedded](#) (10 hours)
- [HCI-User](#) (3 hours)
- [HCI-Design](#) (3 hours)
- [SEP-Privacy](#) and [SEP-DEIA](#) (3 hours)

Prerequisites:

- [AL-Foundational](#)

Course objectives: Students should be able to use human-centered design to build circuits and program a networked microcontroller. Additionally, they will learn to work with real time sensors and understand polarity, Ohm's law, and how to work with electronics safely.

Image Processing to include the following:

- [GIT-Image](#) (20 hours)
- [GIT-Interaction](#) (4 hours)
- [SEP-Privacy](#), [SEP-DEIA](#) and [SEP-IP](#) (3 hours)

Prerequisites:

- [AL-Foundational](#)
- [AL-Strategies](#)
- [SDF-Algorithms](#)
- [SDF-Data-Structures](#)
- [SDF-Practices](#)

Course objectives: Students should understand and be able to appropriately acquire, process, display, and save digital images.

Data Visualization to include the following:

- [GIT-Visualization](#) (20 hours)
- [GIT-Interaction](#) (4 hours)
- [GIT-Fundamentals](#) (4 hours)

- [HCI-User](#) (3 hours)
- [HCI-Design](#) (3 hours)
- [SEP-Privacy](#), [SEP-DEIA](#), and [SEP-Professional-Ethics](#) (3 hours)

Prerequisites:

- [AL-Foundational](#)
- [AL-Strategies](#)
- [SDF-Algorithms](#)
- [SDF-Data-Structures](#)
- [SDF-Practices](#)
- [MSF-Probability](#)
- [MSF-Statistics](#)

Course objectives: Students should understand how to select a dataset; ensure the data are accurate and appropriate; design, develop and test a usable visualization program that depicts the data; and be able to read and evaluate existing visualizations.

Simulation to include the following:

- [GIT-Simulation](#) (10 hours)
- [GIT-Rendering](#) (15 hours)
- [GIT-Shading](#): (6 hours)
- [SEP-Professional-Ethics](#) (3 hours)

Prerequisites:

- [AL-Foundational](#)
- [AL-Strategies](#)
- [SDF-Algorithms](#)
- [SDF-Data-Structures](#)
- [SDF-Practices](#)
- [MSF-Linear](#) as a prerequisite or cover relevant topics in class
- [MSF-Probability](#)

Course objectives: Students should understand and be able to create directable simulations, both of physical and non-physical systems.

Introduction to AR and VR to include the following:

- [GIT-Immersive](#) (15 hours)
- [GIT-Fundamentals](#) (4 hours)
- [GIT-Interactive](#) (8 hours)
- [SEP-Privacy](#), [SEP-DEIA](#), and [SEP-Professional-Ethics](#) (3 hours)

Course objectives: Students should understand and be able to develop VR and AR applications.

Computer Animation to include the following:

- [GIT-Animation](#) (30 hours)
- [SEP-Privacy](#), [SEP-DEIA](#), and [SEP-Professional-Ethics](#) (3 hours)

Prerequisites:

- Interactive Computer Graphics course

Course objectives: Students should understand and be able to create short animations employing the principles of animation.

Lighting and Shading to include the following:

- [GIT-Shading](#) (12 hours)
- [GIT-Modeling](#) (6 hours)
- [GIT-Interaction](#) (4 hours)
- [SEP-IP](#), [SEP-DEIA](#), and [SEP-Professional-Ethics](#) (3 hours)

Prerequisites:

- Interactive Computer Graphics course

Course objectives: Students should be able to create realistic and non-photorealistic lighting and understand the underlying theory of shading and lighting.

Committee

Chair: Susan Reiser, University of North Carolina Asheville, Asheville, NC, USA

Members:

- Erik Brunvand, University of Utah, Salt Lake City, UT, USA
- Kel Elkins, NASA/GSFC Scientific Visualization Studio, Greenbelt, MD, USA
- Jeff Lait, SideFX, Toronto, Canada
- Amruth Kumar, Ramapo College, Mahwah, NJ, USA
- Paul Mihail, Valdosta State University, Valdosta, GA, USA
- Tabitha Peck, Davidson College, Davidson, NC, USA
- Ken Schmidt, NOAA NCEI, Asheville, NC, USA
- Dave Shreiner, UnityTechnologies & Sonoma State University, San Francisco, CA, USA

Contributors:

- Ginger Alford, Southern Methodist University, University Park, TX, USA
- Christopher Andrews, Middlebury College, Middlebury, VT, USA
- A. J. Christensen, NASA/GSFC Scientific Visualization Studio – SSAI, Champaign, IL, USA
- Roger Eastman, University of Maryland, College Park, MD, USA
- Ted Kim, Yale University, New Haven, CT, USA
- Barbara Mones, University of Washington, Seattle, WA, USA
- Greg Shirah, NASA/GSFC Scientific Visualization Studio, Greenbelt, MD, USA
- Beatriz Sousa Santos, University of Aveiro, Portugal
- Anthony Steed, University College, London, UK

Human-Computer Interaction (HCI)

Preamble

Computational systems not only enable users to solve problems, but also foster social connectedness and support a broad variety of human endeavors. Thus, these systems should work well with their users and solve problems in ways that respect individual dignity, social justice, and human values and creativity. Human-computer interaction (HCI) addresses those issues from an interdisciplinary perspective that includes computer science, psychology, business strategy, and design principles.

Each user is different and, from the perspective of HCI, the design of every system that interacts with people should anticipate and respect that diversity. This includes not only accessibility, but also cultural and societal norms, neural diversity, modality, and the responses the system elicits in its users. An effective computational system should evoke trust while it treats its users fairly, respects their privacy, provides security, and abides by ethical principles.

These goals require design-centric engineering that begins with intention and with the understanding that design is an iterative process, one that requires repeated evaluation of its usability and its impact on its users. Moreover, technology evokes user responses, not only by its output, but also by the modalities with which it senses and communicates. This knowledge area heightens the awareness of these issues and should influence every computer scientist.

Changes since CS2013

Driven by this broadened perspective, the HCI knowledge area has revised the CS2013 document in several ways:

- Knowledge units have been renamed and reformulated to reflect current practice and to anticipate future technological development.
- There is increased emphasis on the nature of diversity and the centrality of design focused on the user.
- Modality (e.g., text, speech) is still emphasized given its key role throughout HCI, but with a reduced emphasis on particular modalities in favor of a more timely and empathetic approach.
- The curriculum reflects the importance of understanding and evaluating the impacts and implications of a computational system on its users, including issues in ethics, fairness, trust, and explainability.
- Given its extensive interconnections with other knowledge areas, we believe HCI is itself a cross-cutting knowledge area with connections to Artificial Intelligence, Graphics and Interactive Techniques, Software Development Fundamentals, Software Engineering, and, Society, Ethics, and the Profession.

Core Hours

Knowledge Unit	CS Core	KA Core
Understanding the User	2	5
Accountability and Responsibility in Design	2	2
Accessibility and Inclusive Design	2	2
Evaluating the Design	1	2
System Design	1	5
Society, Ethics, and the Profession	Included in SEP hours	
Total Hours	8	16

The hours shared with SEP include overlapping topics and are counted there.

Knowledge Units

HCI-User: Understanding the User: Individual goals and interactions with others

CS Core:

1. User-centered design and evaluation methods. (See also: [SEP-Context](#), [SEP-Ethical-Analysis](#), [SEP-Professional-Ethics](#))
 - a. “You are not the users”
 - b. User needs-finding
 - c. Formative studies
 - d. Interviews
 - e. Surveys
 - f. Usability tests

KA Core:

2. User-centered design methodology. (See also: [SE-Tools](#))
 - a. Personas/persona spectrum
 - b. User stories/storytelling and techniques for gathering stories
 - c. Empathy maps
 - d. Needs assessment (techniques for uncovering needs and gathering requirements - e.g., interviews, surveys, ethnographic and contextual enquiry) (See also: [SE-Requirements](#))
 - e. Journey maps
 - f. Evaluating the design (See also: [HCI-Evaluation](#))
 - g. Interfacing with stakeholders, as a team
 - h. Risks associated with physical, distributed, hybrid and virtual teams
3. Physical and cognitive characteristics of the user
 - a. Physical capabilities that inform interaction design (e.g., color perception, ergonomics)
 - b. Cognitive models that inform interaction design (e.g., attention, perception and recognition, movement, memory)
 - c. Topics in social/behavioral psychology (e.g., cognitive biases, change blindness)
4. Designing for diverse user populations. (See also: [SEP-DEIA](#), [HCI-Accessibility](#))

- a. How differences (e.g., in race, ability, age, gender, culture, experience, and education) impact user experiences and needs
 - b. Internationalization
 - c. Designing for users from other cultures
 - d. Cross-cultural design
 - e. Challenges to effective design evaluation. (e.g., sampling, generalization; disability and disabled experiences)
 - f. Universal design
5. Collaboration and communication (See also: [AI-SEP](#), [SE-Teamwork](#), [SEP-Communication](#), [SPD-Game](#))
- a. Understanding the user in a multi-user context
 - b. Synchronous group communication (e.g., chat rooms, conferencing, online games)
 - c. Asynchronous group communication (e.g., email, forums, social networks)
 - d. Social media, social computing, and social network analysis
 - e. Online collaboration
 - f. Social coordination and online communities
 - g. Avatars, characters, and virtual worlds

Non-core:

- 6. Multi-user systems

Illustrative Learning Outcomes:

CS Core:

- 1. Conduct a user-centered design process that is integrated into a project.

KA Core:

- 2. Compare and contrast the needs of users with those of designers.
- 3. Identify the representative users of a design and discuss who else could be impacted by it.
- 4. Describe empathy and evaluation as elements of the design process.
- 5. Carry out and document an analysis of users and their needs.
- 6. Construct a user story from a needs assessment.
- 7. Redesign an existing solution to a population whose needs differ from those of the initial target population.
- 8. Contrast the different needs-finding methods for a given design problem.
- 9. Reflect on whether your design would benefit from low-tech or no-tech components.

Non-core:

- 10. Recognize the implications of designing for a multi-user system/context.

HCI-Accountability: Accountability and Responsibility in Design

CS Core: (See also: [SEP-Context](#))

- 1. Design impact
 - a. Sustainability (See also: [SEP-Sustainability](#))

- b. Inclusivity (See also: [SEP-DEIA](#))
 - c. Safety, security and privacy (See also: [SEP-Security](#), [SEC-Foundations](#))
 - d. Harm and disparate impact (See also: [SEP-DEIA](#))
- 2. Ethics in design methods and solutions (See also: [SEP-Ethical-Analysis](#), [SEP-Context](#), [SEP-Intellectual Property](#))
 - a. The role of artificial intelligence (See also: [AI-SEP](#))
 - b. Responsibilities for considering stakeholder impact and human factors (See also: [SEP-Professional-Ethics](#).)
 - c. Role of design to meet user needs
- 3. Requirements in design (See also: [SEP-Professional-Ethics](#))
 - a. Ownership responsibility
 - b. legal frameworks, compliance requirements
 - c. consideration beyond immediate user needs, including via iterative reconstruction of problem analysis and “digital well-being” features

KA Core:

- 4. Value-sensitive design (See also: [SEP-Ethical-Analysis](#), [SEP-Context](#), [SEP-Communication](#))
 - a. identify direct and indirect stakeholders
 - b. Determine and include diverse stakeholder values and value systems.
- 5. Persuasion through design (See also: [SEP-Communication](#))
 - a. Assess the persuasive content of a design
 - b. Employ persuasion as a design goal
 - c. Distinguish persuasive interfaces from manipulative interfaces

Illustrative Learning Outcomes:

CS Core:

- 1. Identify and critique the potential impacts of a design on society and relevant communities to address such concerns as sustainability, inclusivity, safety, security, privacy, harm, and disparate impact.

KA Core:

- 2. Identify the potential human factor elements in a design.
- 3. Identify and understand direct and indirect stakeholders.
- 4. Develop scenarios that consider the entire lifespan of a design, beyond the immediately planned uses that anticipate direct and indirect stakeholders.
- 5. Identify and critique the potential factors in a design that impact direct and indirect stakeholders and broader society (e.g., transparency, sustainability of the system, trust, artificial intelligence).
- 6. Assess the persuasive content of a design and its intent relative to user interests.
- 7. Critique the outcomes of a design given its intent.
- 8. Understand the impact of design decisions.

HCI-Accessibility: Accessibility and Inclusive Design

CS Core:

1. Background (See also: [SEP-DEIA](#), [SEP-Security](#))
 - a. Societal and legal support for and obligations to people with disabilities
 - b. Accessible design benefits everyone
2. Techniques
 - a. Accessibility standards (e.g., Web Content Accessibility Guidelines) (See also: [SPD-Web](#))
3. Technologies (See also: [SE-Tools](#))
 - a. Features and products that enable accessibility and support inclusive development by designers and engineers
4. IDFs (Inclusive Design Frameworks) (See also: [SEP-DEIA](#))
 - a. Recognizing differences
5. Universal design

KA Core:

6. Background
 - a. Demographics and populations (permanent, temporary and situational disability)
 - b. International perspectives on disability (See also: [SEP-DEIA](#))
 - c. Attitudes towards people with disabilities (See also: [SEP-DEIA](#))
7. Techniques
 - a. UX (user experience) design and research
 - b. Software engineering practices that enable inclusion and accessibility. (See also: [SEP-DEIA](#))
8. Technologies
 - a. Examples of accessibility-enabling features, such as conformance to screen readers
9. Inclusive Design Frameworks
 - a. Creating inclusive processes such as participatory design
 - b. Designing for larger impact

Non-core:

10. Background (See also: [SEP-DEIA](#))
 - a. Unlearning and questioning
 - b. Disability studies
11. Technologies: the Return On Investment (ROI) of inclusion
12. Inclusive Design Frameworks: user-sensitive inclusive design (See also: [SEP-DEIA](#))
13. Critical approaches to HCI (e.g., inclusivity) (See also: [SEP-DEIA](#))

Illustrative Learning Outcomes:

CS Core:

1. Identify accessibility challenges faced by people with different disabilities, and specify the associated accessible and assistive technologies that address them. (See also: AI-Agents, AI-Robotics)
2. Identify appropriate inclusive design approaches, such as universal design and ability-based design.
3. Identify and demonstrate understanding of software accessibility guidelines.
4. Demonstrate recognition of laws and regulations applicable to accessible design.

KA Core:

5. Apply inclusive frameworks to design, such as universal design and usability and ability-based design, and demonstrate accessible design of visual, voice-based, and touch-based UIs.
6. Demonstrate understanding of laws and regulations applicable to accessible design.
7. Demonstrate understanding of what is appropriate and inappropriate high level of skill during interaction with individuals from diverse populations.
8. Analyze web pages and mobile apps for current standards of accessibility.

Non-core:

9. Biases towards disability, race, and gender have historically, either intentionally or unintentionally, informed technology design.
 - a. Find examples.
 - b. Consider how those experiences (learnings?) might inform design.
10. Conceptualize user experience research to identify user needs and generate design insights.

HCI-Evaluation: Evaluating the Design**CS Core:**

1. Methods for evaluation with users
 - a. Formative (e.g. needs-finding and exploratory analysis) and summative assessment (e.g. functionality and usability testing)
 - b. Elements to evaluate (e.g., utility, efficiency, learnability, user satisfaction, affective elements such as pleasure and engagement)
 - c. Understanding ethical approval requirements before engaging in user research ((See also: [SE-Tools](#), [SEP-Ethical-Analysis](#), [SEP-Security](#), [SEP-Privacy](#), [SEP-Professional-Ethics](#))

KA Core:

2. Methods for evaluation with users (See also: [SE-Validation](#))
 - a. Qualitative methods (qualitative coding and thematic analysis)
 - b. Quantitative methods (statistical tests)
 - c. Mixed methods (e.g., observation, think-aloud, interview, survey, experiment)
 - d. Presentation requirements (e.g., reports, personas)
 - e. User-centered testing
 - f. Heuristic evaluation
 - g. Challenges and shortcomings to effective evaluation (e.g., sampling, generalization)
3. Study planning
 - a. How to set study goals
 - b. Hypothesis design
 - c. Approvals from Institutional Research Boards and ethics committees (See also: [SEP-Ethical-Analysis](#), [SEP-Security](#), [SEP-Privacy](#))
 - d. How to pre-register a study
 - e. Within-subjects vs between-subjects design
4. Implications and impacts of design with respect to the environment, material, society, security, privacy, ethics, and broader impacts. (See also: [SEC-Foundations](#))
 - a. The environment

- b. Material
- c. Society
- d. Security
- e. Privacy
- f. Ethics
- g. Broader impacts

Non-core:

- 5. Techniques and tools for quantitative analysis
 - a. Statistical packages
 - b. Visualization tools
 - c. Statistical tests (e.g., ANOVA, t-tests, post-hoc analysis, parametric vs non-parametric tests)
 - d. Data exploration and visual analytics; how to calculate effect size.
- 6. Data management
 - a. Data storage and data sharing (open science)
 - b. Sensitivity and identifiability.

Illustrative Learning Outcomes:

CS Core:

- 1. Discuss the differences between formative and summative assessment and their role in evaluating design

KA Core:

- 2. Select appropriate formative or summative evaluation methods at different points throughout the development of a design
- 3. Discuss the benefits of using both qualitative and quantitative methods for evaluation
- 4. Evaluate the implications and broader impacts of a given design
- 5. Plan a usability evaluation for a given user interface, and justify its study goals, hypothesis design, and study design
- 6. Conduct a usability evaluation of a given user interface and draw defensible conclusions given the study design

Non-core:

- 7. Select and run appropriate statistical tests on provided study data to test for significance in the results
- 8. Pre-register a study design, with planned statistical tests

HCI-Design: System Design

CS Core:

- 1. Prototyping techniques and tools
 - a. Low-fidelity prototyping
 - b. Rapid prototyping
 - c. Throw-away prototyping
 - d. Granularity of prototyping

2. Design patterns
 - a. Iterative design
 - b. Universal design (See also: [SEP-DEIA](#))
 - c. Interaction design (e.g., data-driven design, event-driven design)
3. Design constraints
 - a. Platforms (See also: [SPD-Game](#))
 - b. Devices
 - c. Resources
 - d. Balance among usability, security and privacy (See also: [SEC-Foundations](#))

KA Core:

4. Design patterns and guidelines
 - a. Software architecture patterns
 - b. Cross-platform design
 - c. Synchronization considerations
5. Design processes (See also: [SEP:Communication](#))
 - a. Participatory design
 - b. Co-design
 - c. Double-diamond
 - d. Convergence and divergence
6. Interaction techniques (See also: [GIT-Interaction](#))
 - a. Input and output vectors (e.g., gesture, pose, touch, voice, force)
 - b. Graphical user interfaces
 - c. Controllers
 - d. Haptics
 - e. Hardware design
 - f. Error handling
7. Visual UI design (See also: [GIT-Visualization](#))
 - a. Color
 - b. Layout
 - c. Gestalt principles

Non-core:

8. Immersive environments (See also: [GIT-Immersion](#))
 - a. Virtual reality
 - b. Augmented reality, mixed reality
 - c. XR (XReality which encompasses them)
 - d. Spatial audio
9. 3D printing and fabrication
10. Asynchronous interaction models
11. Creativity support tools
12. Voice UI designs

Illustrative Learning Outcomes:

CS Core:

1. Propose system designs tailored to a specified appropriate mode of interaction.
2. Follow an iterative design and development process that incorporates
 - a. Understanding the user
 - b. Developing an increment
 - c. Evaluating the increment
 - d. Feeding those results into a subsequent iteration
3. Explain the impact of changing constraints and design tradeoffs (e.g., hardware, user, security.) on system design

KA Core:

4. Evaluate architectural design approaches in the context of project goals.
5. Identify synchronization challenges as part of the user experience in distributed environments.
6. Evaluate and compare the privacy implications behind different input techniques for a given scenario
7. Explain the rationale behind a UI design based on visual design principles

Non-core:

8. Evaluate the privacy implications within a VR/AR/MR scenario

HCI-SEP: Society, Ethics, and the Profession**CS Core:**

1. Universal and user-centered design (See also: [HCI-User](#), [SEP-DEIA](#))
2. Accountability (See also: [HCI-Accountability](#))
3. Accessibility and inclusive design (See also: [SEP-DEIA](#), [SEP-Security](#))
4. Evaluating the design (See also: [HCI-Evaluation](#))
5. System design (See also: [HCI-Design](#))

KA Core:

6. Participatory and inclusive design processes
7. Evaluating the design: Implications and impacts of design: with respect to the environment, material, society, security, privacy, ethics, and broader impacts (See also: [SEC-Foundations](#), SEP-Privacy)

Non-core:

8. VR/AR/MR scenarios

Illustrative Learning Outcomes:**CS Core:**

1. Conduct a user-centered design process that is integrated into a project.
2. Identify and critique the potential impacts of a design on society and relevant communities to address such concerns as sustainability, inclusivity, safety, security, privacy, harm, and disparate impact

KA Core:

2. Critique a recent example of a non-inclusive design choice, its societal implications, and propose potential design improvements
3. Evaluating the design: Identify the implications and broader impacts of a given design.

Non-core:

4. Evaluate the privacy implications within a VR/AR/MR scenario

Professional Dispositions

- **Adaptable:** An HCI practitioner should be adaptable to address dynamic changes in technology, user needs, and design challenges.
- **Meticulous:** An HCI practitioner should be meticulous to ensure that their products are both user-friendly and meet the objectives of the design project.
- **Empathetic:** An HCI practitioner must demonstrate understanding of the user's needs.
- **Team-oriented:** The successful HCI practitioner should focus on the success of the team.
- **Inventive:** An HCI practitioner should design solutions that are informed by past practice, the needs of the audience, and HCI fundamentals. Creativity is required to blend these into something that solves the problem appropriately and elegantly.

Mathematics Requirements

Required:

- Basic statistics ([MSF-Statistics](#)) to support the evaluation and interpretation of results, including central tendency, variability, frequency distribution

Course Packaging Suggestions

Introduction to HCI for CS majors and minors, to include the following:

- [HCI-User: Understanding the User](#) (7 hours)
- [HCI-Accountability: Accountability and Responsibility in Design](#): (2 hours)
- [HCI-Accessibility: Accessibility and Inclusive Design](#): (4 hours)
- [HCI-Evaluation: Evaluating the Design](#): (3 hours)
- [HCI-Design: System Design](#): (10 hours)
- [HCI-SEP: Society, Ethics, and the Profession](#): (2 hours)

Prerequisites:

- CS2

Course objectives: A student who completes this course should be able to describe user-centered design principles and apply them in the context of a small project.

Introduction to Data Visualization to include the following:

- [GIT-Visualization](#) (30 hours)
- [GIT-Rendering](#) (10 hours)

- [HCI-User: Understanding the User](#) (3 hours)
- [SEP-Privacy, SEP-Ethical-Analysis](#) (4 hours)

Prerequisites:

- CS2
- [MSF-Linear](#)

Course objectives: Students should understand how to select a dataset; ensure the data are accurate and appropriate; and design, develop and test a visualization program that depicts the data and is usable.

Advanced Course: Usability Testing

- [HCI-User](#) (5 hours)
- [HCI-Accountability](#) (3 hours)
- [HCI-Accessibility](#) (4 hours)
- [HCI-Evaluation](#) (20 hours)
- [HCI-Design](#) (3 hours)
- [HCI-SEP](#) (5 hours)

Prerequisites:

- Introductory/Foundation courses in HCI
- Research methods, [MSF-Statistics](#)

Course objectives: Students should be able to formally evaluate products including the design and execution of usability test tasks, recruitment of appropriate users, design of test tasks, design of the test environment, test plan development and implementation, analysis and interpretation of the results, and documentation and presentation of results and recommendations. Students should be able to select appropriate techniques, procedures and protocols to apply in various situations for usability testing with users. Students should also be able to design an appropriate evaluation plan, effectively conduct the usability test, collect data, and analyze results so that they can suggest improvements.

Committee

Chair: Susan L. Epstein, Hunter College and The Graduate Center of The City University of New York, NY, USA

Members:

- Sherif Aly, The American University of Cairo, Cairo, Egypt
- Jeremiah Blanchard, University of Florida, Gainesville, FL, USA
- Zoya Bylinskii, Adobe Research, Cambridge, MA, USA
- Paul Gestwicki, Ball State University, Muncie, IN, USA
- Susan Reiser, University of North Carolina at Asheville, Asheville, NC, USA
- Amanda M. Holland-Minkley, Washington and Jefferson College, Washington, PA, USA
- Ajit Narayanan, Google, Chennai, India
- Nathalie Riche, Microsoft, Redmond, WA, USA
- Kristen Shinohara, Rochester Institute of Technology, Rochester, NY, USA
- Olivier St-Cyr, University of Toronto, Toronto, Canada

Mathematical and Statistical Foundations (MSF)

Preamble

A strong mathematical foundation remains a bedrock of computer science education and infuses the practice of computing whether in developing algorithms, designing systems, modeling real-world phenomena, or computing with data. This Mathematical and Statistical Foundations (MSF) knowledge area – the successor to the ACM CS2013 [6] curriculum's "Discrete Structures" area – seeks to identify the mathematical and statistical material that undergirds modern computer science. The change of name corresponds to a realization both that the broader name better describes the combination of topics from the 2013 report and from those required for the recently growing areas of computer science, such as artificial intelligence, machine learning, data science, and quantum computing, many of which have continuous mathematics as their foundations.

The committee sought the following inputs to prepare their recommendations:

- A survey about mathematical preparation distributed to computer science faculty (nearly 600 respondents) across a variety of institutional types and in various countries;
- Math-related curricular views amongst data collected from ACM's survey of industry professionals (865 respondents);
- Mathematics requirements stated by all the knowledge areas in the report;
- Direct input from the computer science theory community; and
- Review of past curricular reports including recent ones on data science (e.g., Park City report [28]) and quantum computing education.

Changes since CS2013

The breadth of mathematics needed to address the mathematical needs of rapidly growing areas such as artificial intelligence, machine learning, robotics, data science, and quantum computing has grown beyond discrete structures. These areas call for a renewed focus on probability, statistics, and linear algebra, as supported by the faculty survey that asked respondents to rate various mathematical areas in their importance for both an industry career as well as for graduate school: the combined such ratings for probability, statistics, and linear algebra, for example, were 98%, 98% and 89% respectively, reflecting a strong consensus in the computer science academic community.

Core Hours

Acknowledging some tensions

Several challenges face computer science (CS) programs when weighing mathematics requirements: (1) many CS majors, perhaps aiming for a software career, are unenthusiastic about investing in mathematics; (2) institutions such as liberal-arts colleges often limit how many courses a major may require, while others may require common engineering courses that fill up the schedule; and (3) some programs adopt a more pre-professional curricular outlook while others emphasize a more foundational one. Thus, we are hesitant to recommend an all-encompassing set of mathematical topics as "every CS

degree must require.” Instead, we outline two sets of *core* requirements, a CS Core set suited to hours-limited majors and a more expansive set of CS Core plus KA Core to align with technically focused programs. The principle here is that, in light of the additional foundational mathematics needed for AI, data science and quantum computing, programs ought to consider as much as possible from the more expansive CS+KA version unless there are sound institutional reasons for alternative requirements.

Note: the hours in a row (example: linear algebra) add up to 40 (= 5 + 35), reflecting a standard course; shorter combined courses may be created, for example, by including probability in discrete mathematics (29 hours of discrete mathematics, 11 hours of probability).

Knowledge Units	CS Core	KA Core
Discrete Mathematics	29	11
Probability	11	29
Statistics	10	30
Linear Algebra	5	35
Calculus	0	40
Total	55	145

Rationale for recommended hours

CS Core: While some discrete mathematics courses include probability, we highlight its importance with a minimum number of hours (11) to reflect the strong consensus in the academic community based on the survey. Taken together, the total CS core across discrete mathematics and probability (40 hours) is typical of a one-term course. 15 hours are allotted to statistics and linear algebra for basic definitions so that, for example, students should at least be familiar with terms like *mean*, *standard deviation* and *vector*. These could be covered in regular CS courses. Many programs typically include a broader statistics requirement.

KA Core: The KA Core hours can be read as the remaining hours available to flesh out each topic into a standard 40-hour course. Note that the calculus hours roughly correspond to the typical Calculus-I course now standard across the world. Based on our survey, most programs already require Calculus-I. However, we have left out Calculus-II (an additional 40 hours) and leave it to programs to decide whether Calculus-II should be added to program requirements. Programs could choose to require a more rigorous calculus-based probability or statistics sequence, or non-calculus-based versions. Similarly, linear algebra can be taught as an applied course without a calculus prerequisite or as a more advanced course.

Knowledge Units

MSF-Discrete: Discrete Mathematics

CS Core:

1. Sets, relations, functions, cardinality
2. Recursive mathematical definitions
3. Proof techniques (induction, proof by contradiction)
4. Permutations, combinations, counting, pigeonhole principle
5. Modular arithmetic
6. Logic: truth tables, connectives (operators), inference rules, formulas, normal forms, simple predicate logic
7. Graphs: basic definitions
8. Order notation

Illustrative Learning Outcomes:

CS Core:

1. Sets, Relations, and Functions, Cardinality
 - a. Explain with examples the basic terminology of functions, relations, and sets.
 - b. Perform the operations associated with sets, functions, and relations.
 - c. Relate practical examples to the appropriate set, function, or relation model, and interpret the associated operations and terminology in context.
 - d. Calculate the size of a finite set, including making use of the sum and product rules and inclusion-exclusion principle.
 - e. Explain the difference between finite, countable, and uncountable sets.
2. Recursive mathematical definitions
 - a. Apply recursive definitions of sequences or structures (e.g., Fibonacci numbers, linked lists, parse trees, fractals).
 - b. Formulate inductive proofs of statements about recursive definitions.
 - c. Solve a variety of basic recurrence relations.
 - d. Analyze a problem to determine underlying recurrence relations.
 - e. Given a recursive/iterative code snippet, describe its underlying recurrence relation, hypothesize a closed form for the recurrence relation, and prove the hypothesis correct (usually, using induction).
3. Proof Techniques
 - a. Identify the proof technique used in a given proof.
 - b. Outline the basic structure of each proof technique (direct proof, proof by contradiction, and induction) described in this unit.
 - c. Apply each of the proof techniques (direct proof, proof by contradiction, and induction) correctly in the construction of a sound argument.
 - d. Determine which type of proof is best for a given problem.
 - e. Explain the parallels between ideas of mathematical and/or structural induction to recursion and recursively defined structures.
 - f. Explain the relationship between weak and strong induction and give examples of the appropriate use of each.
4. Permutations, combinations, counting

- a. Apply counting arguments, including sum and product rules, inclusion-exclusion principle and arithmetic/geometric progressions.
 - b. Apply the pigeonhole principle in the context of a formal proof.
 - c. Compute permutations and combinations of a set, and interpret the meaning in the context of the particular application.
 - d. Map real-world applications to appropriate counting formalisms, such as determining the number of ways to arrange people around a table, subject to constraints on the seating arrangement, or the number of ways to determine certain hands in cards (e.g., a full house).
5. Modular arithmetic
- a. Perform computations involving modular arithmetic.
 - b. Explain the notion of greatest common divisor, and apply Euclid's algorithm to compute it.
6. Logic
- a. Convert logical statements from informal language to propositional and predicate logic expressions.
 - b. Apply formal methods of symbolic propositional and predicate logic, such as calculating validity of formulae, computing normal forms, or negating a logical statement.
 - c. Use the rules of inference to construct proofs in propositional and predicate logic.
 - d. Describe how symbolic logic can be used to model real-life situations or applications, including those arising in computing contexts such as software analysis (e.g., program correctness), database queries, and algorithms.
 - e. Apply formal logic proofs and/or informal, but rigorous, logical reasoning to real problems, such as predicting the behavior of software or solving problems such as puzzles.
 - f. Describe the strengths and limitations of propositional and predicate logic.
 - g. Explain what it means for a proof in propositional (or predicate) logic to be valid.
7. Graphs
- a. Illustrate by example the basic terminology of graph theory, and some of the properties and special cases of types of graphs, including trees.
 - b. Demonstrate different traversal methods for trees and graphs, including pre-, post-, and in-order traversal of trees, along with breadth-first and depth-first search for graphs.
 - c. Model a variety of real-world problems in computer science using appropriate forms of graphs and trees, such as representing a network topology, the organization of a hierarchical file system, or a social network.
 - d. Show how concepts from graphs and trees appear in data structures, algorithms, proof techniques (structural induction), and counting.

KA Core:

The recommended topics are the same between CS core and KA-core, but with far more hours, the KA-core can cover these topics in depth and might include more computing-related applications.

MSF-Probability: Probability

CS Core:

1. Basic notions: sample spaces, events, probability, conditional probability, Bayes' rule
2. Discrete random variables and distributions
3. Continuous random variables and distributions

4. Expectation, variance, law of large numbers, central limit theorem
5. Conditional distributions and expectation
6. Applications to computing, the difference between probability and statistics (as subjects)

KA Core:

The recommended topics are the same between CS core and KA-core, but with far more hours, the KA-core can cover these topics in depth and might include more computing-related applications.

Illustrative Learning Outcomes:

CS Core:

1. Basic notions: sample spaces, events, probability, conditional probability, Bayes' rule
 - a. Translate a prose description of a probabilistic process into a formal setting of sample spaces, outcome probabilities, and events.
 - b. Calculate the probability of simple events.
 - c. Determine whether two events are independent.
 - d. Compute conditional probabilities, including through applying (and explaining) Bayes' Rule.
2. Discrete random variables and distributions
 - a. Define the concept of a random variable and indicator random variable.
 - b. Determine whether two random variables are independent.
 - c. Identify common discrete distributions (e.g., uniform, Bernoulli, binomial, geometric).
3. Continuous random variables and distributions
 - a. Identify common continuous distributions (e.g., uniform, normal, exponential).
 - b. Calculate probabilities using cumulative density functions.
4. Expectation, variance, law of large numbers, central limit theorem
 - a. Define the concept of expectation and variance of a random variable.
 - b. Compute the expected value and variance of simple or common discrete/continuous random variables.
 - c. Explain the relevance of the law of large numbers and central limit theorem to probability calculations.
5. Conditional distributions and expectation
 - a. Explain the distinction between a joint, marginal and conditional distributions.
 - b. Compute marginal and conditional distributions from a full distribution, for both discrete and continuous random variables.
 - c. Compute conditional expectations for both discrete and continuous random variables.
6. Applications to computing
 - a. Describe how probability can be used to model real-life situations or applications, such as predictive text, hash tables, and quantum computation.
 - b. Apply probabilistic processes to solving computational problems, such as through randomized algorithms or in security contexts.

MSF-Statistics: Statistics

CS Core:

1. Basic definitions and concepts: populations, samples, measures of central tendency, variance
2. Univariate data: point estimation, confidence intervals

KA Core:

3. Multivariate data: estimation, correlation, regression
4. Data transformation: dimension reduction, smoothing
5. Statistical models and algorithms
6. Hypothesis testing

Illustrative Learning Outcomes:**CS Core:**

1. Basic definitions and concepts: populations, samples, measures of central tendency, variance
 - a. Create and interpret frequency tables
 - b. Display data graphically and interpret graphs (e.g. histograms)
 - c. Recognize, describe and calculate means, medians, quantiles (location of data)
 - d. Recognize, describe and calculate variances, interquartile ranges (spread of data)
2. Univariate data: point estimation, confidence intervals
 - a. Formulate maximum likelihood estimation (in linear-Gaussian settings) as a least-squares problem
 - b. Calculate maximum likelihood estimates
 - c. Calculate maximum a posteriori estimates and make a connection with regularized least squares
 - d. Compute confidence intervals as a measure of uncertainty

KA Core:

3. Sampling, bias, adequacy of samples, Bayesian vs frequentist interpretations
4. Multivariate data: estimation, correlation, regression
 - a. Formulate the multivariate maximum likelihood estimation problem as a least-squares problem
 - b. Interpret the geometric properties of maximum likelihood estimates
 - c. Derive and calculate the maximum likelihood solution for linear regression
 - d. Derive and calculate the maximum a posteriori estimate for linear regression
 - e. Implement both maximum likelihood and maximum a posteriori estimates in the context of a polynomial regression problem
 - f. Formulate and understand the concept of data correlation (e.g., in 2D)
5. Data transformation: dimension reduction, smoothing
 - a. Formulate and derive Principal Component Analysis (PCA) as a least-squares problem
 - b. Geometrically interpret PCA (when solved as a least-squares problem)
 - c. Describe when PCA works well (one can relate back to correlated data)
 - d. Geometrically interpret the linear regression solution (maximum likelihood)
6. Statistical models and algorithms
 - a. Apply PCA to dimensionality reduction problems
 - b. Describe the trade-off between compression and reconstruction power
 - c. Apply linear regression to curve-fitting problems
 - d. Explain the concept of overfitting
 - e. Discuss and apply cross-validation in the context of overfitting and model selection (e.g., degree of polynomials in a regression context)

MSF-Linear: Linear Algebra

CS Core:

1. Vectors: definitions, vector operations, geometric interpretation, angles: Matrices: definition, matrix operations, meaning of $Ax=b$.

KA Core:

2. Matrices, matrix-vector equation, geometric interpretation, geometric transformations with matrices
3. Solving equations, row-reduction
4. Linear independence, span, basis
5. Orthogonality, projection, least-squares, orthogonal bases
6. Linear combinations of polynomials, Bezier curves
7. Eigenvectors and eigenvalues
8. Applications to computer science: Principal Components Analysis (PCA), Singular Value Decomposition (SVD), page-rank, graphics

Illustrative Learning Outcomes:

CS Core:

1. Vectors: definitions, vector operations, geometric interpretation, angles
 - a. Describe algebraic and geometric representations of vectors in \mathbb{R}^n and their operations, including addition, scalar multiplication and dot product
 - b. List properties of vectors in \mathbb{R}^n
 - c. Compute angles between vectors in \mathbb{R}^n

KA Core:

2. Matrices, matrix-vector equation, geometric interpretation, geometric transformations with matrices
 - a. Perform common matrix operations, such as addition, scalar multiplication, multiplication, and transposition
 - b. Relate a matrix to a homogeneous system of linear equations
 - c. Recognize when two matrices can be multiplied
 - d. Relate various matrix transformations to geometric illustrations
3. Solving equations, row-reduction
 - a. Formulate, solve, apply, and interpret properties of linear systems
 - b. Perform row operations on a matrix
 - c. Relate an augmented matrix to a system of linear equations
 - d. Solve linear systems of equations using the language of matrices
 - e. Translate word problems into linear equations
 - f. Perform Gaussian elimination
4. Linear independence, span, basis
 - a. Define subspace of a vector space
 - b. List examples of subspaces of a vector space
 - c. Recognize and use basic properties of subspaces and vector spaces
 - d. Determine whether or not particular subsets of a vector space are subspaces
 - e. Discuss the existence of a basis of an abstract vector space
 - f. Describe coordinates of a vector relative to a given basis

- g. Determine a basis and the dimension of a finite-dimensional space
- h. Discuss spanning sets for vectors in \mathbb{R}^n
- i. Discuss linear independence for vectors in \mathbb{R}^n
- j. Define the dimension of a vector space
- 5. Orthogonality, projection, least-squares, orthogonal bases
 - a. Explain the Gram-Schmidt orthogonalization process
 - b. Define orthogonal projections
 - c. Define orthogonal complements
 - d. Compute the orthogonal projection of a vector onto a subspace, given a basis for the subspace
 - e. Explain how orthogonal projections relate to least square approximations
- 6. Linear combinations of polynomials, Bezier curves
 - a. Identify polynomials as generalized vectors
 - b. Explain linear combinations of basic polynomials
 - c. Describe orthogonality for polynomials
 - d. Distinguish between basic polynomials and Bernstein polynomials
 - e. Apply Bernstein polynomials to Bezier curves
- 7. Eigenvectors and eigenvalues
 - a. Find the eigenvalues and eigenvectors of a matrix
 - b. Define eigenvalues and eigenvectors geometrically
 - c. Use characteristic polynomials to compute eigenvalues and eigenvectors
 - d. Use eigenspaces of matrices, when possible, to diagonalize a matrix
 - e. Perform diagonalization of matrices
 - f. Explain the significance of eigenvectors and eigenvalues
 - g. Find the characteristic polynomial of a matrix
 - h. Use eigenvectors to represent a linear transformation with respect to a particularly nice basis
- 8. Applications to computer science: PCA, SVD, page-rank, graphics
 - a. Explain the geometric properties of PCA
 - b. Relate PCA to dimensionality reduction
 - c. Relate PCA to solving least-squares problems
 - d. Relate PCA to solving eigenvector problems
 - e. Apply PCA to reducing the dimensionality of a high-dimensional dataset (e.g., images)
 - f. Explain the page-rank algorithm and understand how it relates to eigenvector problems
 - g. Explain the geometric differences between SVD and PCA
 - h. Apply SVD to a concrete example (e.g., movie rankings)

MSF-Calculus

KA Core:

1. Sequences, series, limits
2. Single-variable derivatives: definition, computation rules (chain rule etc.), derivatives of important functions, applications
3. Single-variable integration: definition, computation rules, integrals of important functions, fundamental theorem of calculus, definite vs indefinite, applications (including in probability)
4. Parametric and polar representations
5. Taylor series

6. Multivariate calculus: partial derivatives, gradient, chain-rule, vector valued functions,
7. Optimization: convexity, global vs local minima, gradient descent, constrained optimization and Lagrange multipliers.
8. Ordinary Differential Equations (ODEs): definition, Euler method, applications to simulation, Monte Carlo integration
9. CS applications: gradient descent for machine learning, forward and inverse kinematics, applications of calculus to probability

Note: the calculus topics listed above are aligned with computer science goals rather than with traditional calculus courses. For example, multivariate calculus is often a course by itself but computer science undergraduates only need parts of it for machine learning.

Illustrative Learning Outcomes:

KA Core:

1. Sequences, series, limits
 - a. Explain the difference between infinite sets and sequences
 - b. Explain the formal definition of a limit
 - c. Derive the limit for examples of sequences and series
 - d. Explain convergence and divergence
 - e. Apply L'Hospital's rule and other approaches to resolving limits
2. Single-variable derivatives: definition, computation rules (chain rule etc.), derivatives of important functions, applications
 - a. Explain a derivative in terms of limits
 - b. Explain derivatives as functions
 - c. Perform elementary derivative calculations from limits
 - d. Apply sum, product and quotient rules
 - e. Work through examples with important functions
3. Single-variable integration: definition, computation rules, integrals of important functions, fundamental theorem of calculus, definite vs indefinite, applications (including in probability)
 - a. Explain the definitions of definite and indefinite integrals
 - b. Apply integration rules to examples with important functions
 - c. Explore the use of the fundamental theorem of calculus
 - d. Apply integration to problems
4. Parametric and polar representations
 - a. Apply parametric representations of important curves
 - b. Apply polar representations
5. Taylor series
 - a. Derive Taylor series for some important functions
 - b. Apply the Taylor series to approximations
6. Multivariate calculus: partial derivatives, gradient, chain-rule, vector valued functions, applications to optimization, convexity, global vs local minima
 - a. Compute partial derivatives and gradients
 - b. Work through examples with vector-valued functions with gradient notation
 - c. Explain applications to optimization
7. ODEs: definition, Euler method, applications to simulation

- a. Apply the Euler method to integration
- b. Apply the Euler method to a single-variable differential equation
- c. Apply the Euler method to multiple variables in an ODE

Professional Dispositions

We focus on dispositions helpful to students learning mathematics:

- **Growth mindset.** Perhaps the most important of the dispositions, students should be persuaded that anyone can learn mathematics, certainly the subset foundational to CS, and that success is not dependent on innate ability.
- **Practice mindset.** Students should be educated about the nature of “doing” mathematics and learning through practice with problems as opposed to merely listening or observing demonstrations in the classroom.
- **Deferred gratification.** Most students are likely to learn at least some mathematics from mathematics departments unfamiliar with computing applications; computing departments should acclimate the students to the notion of waiting to see computing applications. Many of the new growth areas such as AI or quantum computing can serve as motivation.
- **Persistence.** Student perceptions are often driven by frustration with unable to solve hard problems that they see some peers tackle seemingly effortlessly; computing departments should help promote the notion that eventual success through persistence is what matters.
- **Skepticism and inquiry.** Students often look for “given formulas” as handed down by experts only to be memorized and used. Yet, a theoretical mindset and, more broadly, a scientific one, should feature skepticism and a curiosity about how formulas are established.

Mathematics Requirements

The most important topics expected from students entering a computing program typically correspond to pre-calculus courses in high school..

Required:

- Algebra and numeracy:
 - Numeracy: numbers, operations, types of numbers, fluency with arithmetic, exponent notation, rough orders of magnitude, fractions and decimals.
 - Algebra: rules of exponents, solving linear or quadratic equations with one or two variables, factoring, algebraic manipulation of expressions with multiple variables.
- Precalculus:
 - Coordinate geometry: distances between points, areas of common shapes
 - Functions: function notation, drawing and interpreting graphs of functions
 - Exponentials and logarithms: a general familiarity with the functions and their graphs
 - Trigonometry: familiarity with basic trigonometric functions and the unit circle

Course Packaging Suggestions

Every department faces constraints in delivering content, which precludes merely requiring a long list of courses covering every single desired topic. These constraints include content-area ownership, faculty size, student preparation, and limits on the number of departmental courses a curriculum can require. We list below some options for offering mathematical foundations, combinations of which might best fit any particular institution:

- **Traditional course offerings.** With this approach, a computer science department can require students to take courses provided by mathematics departments in any of the five broad mathematical areas listed above.
- **A “Continuous Structures” analog of Discrete Structures.** Many computer science departments now offer courses that prepare students mathematically for AI and machine learning. Such courses can combine just enough calculus, optimization, linear algebra and probability; yet others may split linear algebra into its own course. These courses have the advantage of motivating students with computing applications, and including programming as pedagogy for mathematical concepts.
- **Integration into application courses.** An application course, such as machine learning, can be spread across two courses, with the course sequence including the needed mathematical preparation taught just-in-time, or a single machine learning course can balance preparatory material with new topics. This may have the advantage of mitigating turf issues and helping students see applications immediately after encountering mathematics.
- **Specific course adaptations.** For nearly a century, physics and engineering needs have driven the structure of calculus, linear algebra, and probability. Computer science departments can collaborate with their colleagues in mathematics departments to restructure mathematics-offered sections in these areas that are driven by computer science applications. For example, calculus could be reorganized to fit the needs of computing programs into two calculus courses, leaving a later third calculus course for engineering and physics students.

Committee

Chair: Rahul Simha, The George Washington University, Washington DC, USA

Members:

- Richard Blumenthal, Regis University, Denver, CO, USA
- Marc Deisenroth, University College London, London, UK
- Mlkey Goldweber, Denison University, Granville, OH, USA
- David Liben-Nowell, Carleton College, Northfield, MN, USA
- Jodi Tims, Northeastern University, Boston, MA, USA

Networking and Communication (NC)

Preamble

Networking and communication play a central role in interconnected computer systems that are transforming the daily lives of billions of people. The public Internet provides connectivity for networked

applications that serve ever-increasing numbers of individuals and organizations around the world. Complementing the public sector, major proprietary networks leverage their global footprints to support cost-effective distributed computing, storage, and content delivery. Advances in satellite networks expand connectivity to rural areas. Device-to-device communication underlies the emerging Internet of Things.

This knowledge area deals with key concepts in networking and communication, as well as their representative instantiations in the Internet and other computer networks. Besides the basic principles of switching and layering, the area at its core provides knowledge on naming, addressing, reliability, error control, flow control, congestion control, domain hierarchy, routing, forwarding, modulation, encoding, framing, and access control. The area also covers knowledge units in network security and mobility, such as security threats, countermeasures, device-to-device communication, and multi-hop wireless networking. In addition to the fundamental principles, the area includes their specific realization of the Internet as well as hands-on skills in the implementation of networking and communication concepts. Finally, the area comprises emerging topics such as network virtualization and quantum networking.

As the main learning outcome, learners develop a thorough understanding of the role and operation of networking and communication in networked computer systems. They learn how network structure and communication protocols affect the behavior of distributed applications. The area educates on not only key principles but also their specific instantiations in the Internet and equips the student with hands-on implementation skills. While computer-system, networking, and communication technologies are advancing at a fast pace, the gained fundamental knowledge enables the student to readily apply the concepts in new technological settings.

Changes since CS2013

Compared to the 2013 curricula, the knowledge area broadens its core focus to expand on reliability support, routing, forwarding, and single-hop communication. Due to the enhanced core, learners acquire a deeper understanding of the impact that networking and communication have on the behavior of distributed applications. Reflecting the increased importance of network security, the area adds a respective knowledge unit as a new elective. To track the advancing frontiers in networking and communication knowledge, the social networking knowledge unit was removed and an emerging knowledge unit on topics, such as middleboxes, software-defined networks, and quantum networking, was added. Other changes consist of redistributing all the topics from the old unit on resource allocation among other units to resolve overlap between knowledge units in the 2013 curricula.

Core Hours

Knowledge Units	CS Core	KA Core
Fundamentals	2.5 + 0.25 (SEP) + 0.25 (SF)	
Networked Applications	3.5 + 0.25 (SEP) + 0.25 (PDC)	

Reliability Support		5.75 + 0.25 (SF)
Routing And Forwarding		4
Single-Hop Communication		3
Mobility Support		4
Network Security		2.25 + 0.5 (SEC) + 0.25 (SEP)
Emerging Topics		4
Total	7	24

Knowledge Units

NC-Fundamentals: Fundamentals

CS Core:

1. Importance of networking in contemporary computing, and associated challenges. (See also: [SEP-Context](#), [SEP-Privacy](#))
2. Organization of the Internet (e.g. users, Internet Service Providers, autonomous systems, content providers, content delivery networks).
3. Switching techniques (e.g., circuit and packet).
4. Layers and their roles (application, transport, network, datalink, and physical).
5. Layering principles (e.g. encapsulation and hourglass model). (See also: [SF-Foundations](#))
6. Network elements (e.g. routers, switches, hubs, access points, and hosts).
7. Basic queueing concepts (e.g. relationship with latency, congestion, service levels, etc.).

Illustrative Learning Outcomes:

CS Core:

1. Articulate the organization of the Internet.
2. List and define the appropriate network terminology.
3. Describe the layered structure of a typical networked architecture.
4. Identify the different types of complexity in a network (edges, core, etc.).

NC-Applications: Networked Applications

CS Core:

1. Naming and address schemes (e.g. DNS, and Uniform Resource Identifiers).
2. Distributed application paradigms (e.g. client/server, peer-to-peer, cloud, edge, and fog). (See also: [PDC-Communication](#), [PDC-Coordination](#))
3. Diversity of networked application demands (e.g. latency, bandwidth, and loss tolerance) (See also: [PDC-Communication](#), [SEP-Sustainability](#), [SEP-Context](#)).
4. Coverage of application-layer protocols (e.g. HTTP).

5. Interactions with TCP, UDP, and Socket APIs (See also: [PDC-Programs](#)).

Illustrative Learning Outcomes:

CS Core:

1. Define the principles of naming, addressing, resource location.
2. Analyze the needs of specific networked application demands.
3. Describe the details of one application layer protocol.
4. Implement a simple client-server socket-based application.

NC-Reliability: Reliability Support

KA Core:

1. Unreliable delivery (e.g. UDP).
2. Principles of reliability (e.g. delivery without loss, duplication, or out of order) (See also: [SF-Reliability](#)).
3. Error control (e.g. retransmission, error correction).
4. Flow control (e.g. stop and wait, window based).
5. Congestion control (e.g. implicit and explicit congestion notification).
6. TCP and performance issues (e.g. Tahoe, Reno, Vegas, Cubic).

Illustrative Learning Outcomes:

KA Core:

1. Describe the operation of reliable delivery protocols.
2. List the factors that affect the performance of reliable delivery protocols.
3. Describe some TCP reliability design issues.
4. Design and implement a simple reliable protocol.

NC-Routing: Routing and Forwarding

KA Core:

1. Routing paradigms and hierarchy (e.g. intra/inter domain, centralized and decentralized, source routing, virtual circuits, QoS).
2. Forwarding methods (e.g. forwarding tables and matching algorithms).
3. IP and Scalability issues (e.g. NAT, CIDR, BGP, different versions of IP).

Illustrative Learning Outcomes:

KA Core:

1. Describe various routing paradigms and hierarchies.
2. Describe how packets are forwarded in an IP network.
3. Describe how the Internet tackles scalability challenges. .

NC-SingleHop: Single Hop Communication

KA Core:

1. Introduction to modulation, bandwidth, and communication media.
2. Encoding and Framing.

3. Medium Access Control (MAC) (e.g. random access and scheduled access).
4. Ethernet and WiFi.
5. Switching (e.g. spanning trees, VLANs).
6. Local Area Network Topologies (e.g. data center, campus networks).

Illustrative Learning Outcomes:

KA Core:

1. Describe some basic aspects of modulation, bandwidth, and communication media.
2. Describe in detail on a MAC protocol.
3. Demonstrate understanding of encoding and framing solution tradeoffs.
4. Describe details of the implementation of Ethernet.
5. Describe how switching works.
6. Describe one kind of a LAN topology.

NC-Security: Network Security

KA Core:

1. General intro about security (Threats, vulnerabilities, and countermeasures). (See also: [SEP-Security](#), [SEC-Foundations](#), [SEC-Engineering](#))
2. Network specific threats and attack types (e.g., denial of service, spoofing, sniffing and traffic redirection, man-in-the-middle, message integrity attacks, routing attacks, ransomware, and traffic analysis). (See also: [SEC-Foundations](#), [SEC-Engineering](#))
3. Countermeasures. (See also: [SEC-Foundations](#), [SEC-Crypto](#), [SEC-Engineering](#))
 - a. Cryptography (e.g. SSL, TLS, symmetric/asymmetric).
 - b. Architectures for secure networks (e.g., secure channels, secure routing protocols, secure DNS, VPNs, DMZ, Zero Trust Network Access, hyper network security, anonymous communication protocols, isolation).
 - c. Network monitoring, intrusion detection, firewalls, spoofing and DoS protection, honeypots, tracebacks, BGP Sec, RPKI.

Illustrative Learning Outcomes:

KA Core:

1. Describe some of the threat models of network security.
2. Describe specific network-based countermeasures.
3. Analyze various aspects of network security from a case study.

NC-Mobility: Mobility

KA Core:

1. Principles of cellular communication (e.g. 4G, 5G).
2. Principles of Wireless LANs (mainly 802.11).
3. Device to device communication (e.g. IoT communication).
4. Multi-hop wireless networks (e.g. ad hoc networks, opportunistic, delay tolerant).

Illustrative Learning Outcomes:

KA Core:

1. Describe some aspects of cellular communication such as registration.
2. Describe how 802.11 supports mobile users.
3. Describe practical uses of device to device communication, as well as multihop.
4. Describe one type of mobile network such as ad hoc.

NC-Emerging: Emerging Topics**KA Core:**

1. Middleboxes (e.g. advances in usage of AI, intent-based networking, filtering, deep packet inspection, load balancing, NAT, CDN).
2. Network Virtualization (e.g. SDN, Data Center Networks).
3. Quantum Networking (e.g. Intro to the domain, teleportation, security, Quantum Internet).
4. Satellite, mmWave, Visible Light.

Illustrative Learning Outcomes:**KA Core:**

1. Describe the value of advances in middleboxes in networks.
2. Describe the importance of Software Defined Networks.
3. Describe some of the added value achieved by using Quantum Networking.

Professional Dispositions

- **Meticulous:** Students must be particular about the specifics of understanding and creating networking protocols.
- **Collaborative:** Students must work together to develop multiple components that interact together, and respond to failures and threats.
- **Proactive:** Students must be able to predict failures, threats, and how to deal with them while avoiding reactive modes of operation only.
- **Professional:** Students must comply with the needs of the community and their expectations from a networked environment, and the demands of regulatory bodies.
- **Responsive:** Students must act swiftly to changes in requirements in network configurations and changing user requirements.
- **Adaptive:** Students need to reconfigure systems under varying modes of operation.

Mathematics Requirements

Required:

- [MSF-Probability](#).
- [MSF-Statistics](#).
- [MSF-Discrete](#).
- [MSF-Linear](#) Simple queuing theory concepts.

Course Packaging Suggestions

Coverage of the concepts of networking including but not limited to types of applications used by the network, reliability, routing and forwarding, single hop communication, security, and other emerging topics.

Note: both courses cover the same KU's but with different allocation of hours for each KU.

Course objectives: By the end of this course, learners should be able to understand many of the fundamental concepts associated with networking, learn about many types of networked applications, and develop at least one, understand basic routing and forwarding, single hop communications, and deal with some issues pertaining to mobility, security, and emerging areas, all with embedded social, ethical, and issues pertaining to the profession.

Introductory Course:

- [NC-Fundamentals](#) (8 hours)
- [NC-Applications](#) (12 hours)
- [NC-Reliability](#) (6 hours)
- [NC-Routing](#) (4 hours)
- [NC-SingleHop](#) (3 hours)
- [NC-Mobility](#) (3 hours)
- [NC-Security](#) (3 hours)
- [SEP-Context](#) (1 hour)
- [NC-Emerging](#) (2 hours)

Course objectives: By the end of this course, learners would have obtained a refresher about some of the fundamental issues of networking, networked applications, reliability, and routing and forwarding, and indulged in more details into single hop communications, mobility, security, and more emerging topics in the area, all with embedded social, ethical, and issues pertaining to the profession.

Advanced Course:

- [NC-Fundamentals](#) (3 hours)
- [NC-Applications](#) (4 hours)
- [NC-Reliability](#) (7 hours)
- [NC-Routing](#) (6 hours)
- [NC-SingleHop](#) (5 hours)
- [NC-Mobility](#) (5 hours)
- [NC-Security](#) (5 hours)
- [SEP-Privacy](#), [SEP-Security](#), [SEP-Sustainability](#) (2 hours)
- [NC-Emerging](#) (5 hours)

Committee

Chair: Sherif G. Aly, The American University in Cairo, Cairo, Egypt

Members:

- Khaled Harras, Carnegie Mellon University, Pittsburgh, PA, USA
- Moustafa Youssef, The American University in Cairo, Cairo, Egypt
- Sergey Gorinsky, IMDEA Networks Institute, Madrid, Spain
- Qiao Xiang, Xiamen University, Xiamen, China

Contributors:

- Alex (Xi) Chen: Huawei, Montreal, Canada

Operating Systems (OS)

Preamble

The operating system is a collection of services needed to safely interface the hardware with applications. Core topics focus on the mechanisms and policies needed to virtualize computation, memory, and Input/Output (I/O). Overarching themes that are reused at many levels in computer systems are well illustrated in operating systems (e.g., polling vs interrupts, caching, flexibility vs costs, scheduling approaches to processes, page replacement, etc.). The Operating Systems knowledge area contains the key underlying concepts for other knowledge areas – trust boundaries, concurrency, persistence, and safe extensibility.

Changes since CS2013

Changes from CS2013 include:

- Renamed File Systems knowledge unit to File Systems API and Implementation knowledge unit,
- Moved topics from the previous Performance and Evaluation knowledge unit to the Systems Fundamentals ([SF](#)) knowledge area,
- Moved some topics from File Systems API and Implementation and Device Management to the Advanced File Systems knowledge unit, and
- Added topics on systems programming and the creation of platform-specific executables to the Foundations of Programming Languages ([FPL](#)) knowledge area.

Core Hours

Knowledge Units	CS Core	KA Core
Role and Purpose of Operating Systems	2	
Principles of Operating System	2	
Concurrency	2	1
Protection and Safety	2	1
Scheduling		2
Process Model		2
Memory Management		0.5 +1.5 (AR)
Device Management		0.5+0.5 (AR)

Knowledge Units	CS Core	KA Core
Role and Purpose of Operating Systems	2	
Principles of Operating System	2	
Concurrency	2	1
File Systems API and Implementation		2
Advanced File Systems		1
Virtualization		1
Real-time and Embedded Systems		1
Fault Tolerance		1
Society, Ethics, and the Profession	Included in SEP hours	
Total	8	14 (+ 2 counted in AR)

Knowledge Units

OS-Purpose: Role and Purpose of Operating Systems

CS Core:

1. Operating systems mediate between general purpose hardware and application-specific software.
2. Universal operating system functions (e.g., process, user and device interfaces, persistence of data).
3. Extended and/or specialized operating system functions (e.g., embedded systems, server types such as file, web, multimedia, boot loaders and boot security).
4. Design issues (e.g., efficiency, robustness, flexibility, portability, security, compatibility, power, safety, tradeoffs between error checking and performance, flexibility and performance, and security and performance). (See also: [SEC-Engineering](#))
5. Influences of security, networking, multimedia, parallel and distributed computing.
6. Overarching concern of security/protection: Neglecting to consider security at every layer creates an opportunity to inappropriately access resources.

Example concepts:

- a. Unauthorized access to files on an unencrypted drive can be achieved by moving the media to another computer,
- b. Operating systems enforced security can be defeated by infiltrating the boot layer before the operating system is loaded,
- c. Process isolation can be subverted by inadequate authorization checking at API boundaries,

- d. Vulnerabilities in system firmware can provide attack vectors that bypass the operating system entirely,
 - e. Improper isolation of virtual machine memory, computing, and hardware can expose the host system to attacks from guest systems, and
 - f. The operating system may need to mitigate exploitation of hardware and firmware vulnerabilities, leading to potential performance reductions (e.g., Spectre and Meltdown mitigations).
7. Exposure of operating systems functions in shells and systems programming. (See also: [FPL-Scripting](#))

Illustrative Learning Outcomes:

CS Core:

- 1. Understand the objectives and functions of modern operating systems.
- 2. Evaluate the design issues in different usage scenarios (e.g., real time OS, mobile, server).
- 3. Understand the functions of a contemporary operating system with respect to convenience, efficiency, and the ability to evolve.
- 4. Understand how evolution and stability are desirable and mutually antagonistic in operating systems function.

OS-Principles: Principles of Operating System

CS Core:

- 1. Operating system software design and approaches (e.g., monolithic, layered, modular, micro-kernel, unikernel).
- 2. Abstractions, processes, and resources.
- 3. Concept of system calls and links to application program interfaces (e.g., Win32, Java, Posix). (See also: [AR-Assembly](#))
- 4. The evolution of the link between hardware architecture and the operating system functions.
- 5. Protection of resources means protecting some machine instructions/functions. (See also: [AR-Assembly](#))

Example concepts:

- a. Applications cannot arbitrarily access memory locations or file storage device addresses, and
 - b. Protection of coprocessors and network devices.
6. Leveraging interrupts from hardware level: service routines and implementations. (See also: [AR-Assembly](#))

Example concepts:

- a. Timer interrupts for implementing time slices, and
 - b. I/O interrupts for putting blocking threads to sleep without polling.
7. Concept of user/system state and protection, transition to kernel mode using system calls. (See also: [AR-Assembly](#))
8. Mechanism for invoking system calls, the corresponding mode and context switch and return from interrupt. (See also: [AR-Assembly](#))
9. Performance costs of context switches and associated cache flushes when performing process switches in Spectre-mitigated environments.

Illustrative Learning Outcomes:

CS Core:

1. Understand how the application of software design approaches to operating systems design/implementation (e.g., layered, modular, etc.) affects the robustness and maintainability of an operating system.
2. Categorize system calls by purpose.
3. Understand dynamics of invoking a system call (passing parameters, mode change, etc.).
4. Evaluate whether a function can be implemented in the application layer or can only be accomplished by system calls.
5. Apply OS techniques for isolation, protection and throughput across OS functions (e.g., starvation similarities in process scheduling, disk request scheduling, semaphores, etc.) and beyond.
6. Understand how the separation into kernel and user mode affects safety and performance.
7. Understand the advantages and disadvantages of using interrupt processing in enabling multiprogramming.
8. Analyze for potential vectors of attack via the operating systems and the security features designed to guard against them.

OS-Concurrency: Concurrency

CS Core:

1. Thread abstraction relative to concurrency.
2. Race conditions, critical regions (role of interrupts if needed). (See also: [PDC-Programs](#))
3. Deadlocks and starvation. (See also: [PDC-Coordination](#))
4. Multiprocessor issues (spin-locks, reentrancy).
5. Multiprocess concurrency vs multithreading.

KA Core:

6. Thread creation, states, structures. (See also: [SF-Foundations](#))
7. Thread APIs.
8. Deadlocks and starvation (necessary conditions/mitigations). (See also: [PDC-Coordination](#))
9. Implementing thread safe code (semaphores, mutex locks, condition variables). (See also: [AR-Performance-Energy](#), [SF-Evaluation](#), [PDC-Evaluation](#))
10. Race conditions in shared memory. (See also: [PDC-Coordination](#))

Non-Core:

11. Managing atomic access to OS objects (e.g., big kernel lock vs many small locks vs lockless data structures like lists).

Illustrative Learning Outcomes:

CS Core:

1. Understand the advantages and disadvantages of concurrency as inseparable functions within the operating system framework.
2. Understand how architecture level implementation results in concurrency problems including race conditions.
3. Understand concurrency issues in multiprocessor systems.

KA Core:

4. Understand the range of mechanisms that can be employed at the operating system level to realize concurrent systems and describe the benefits of each.
5. Understand techniques for achieving synchronization in an operating system (e.g., describe how a semaphore can be implemented using OS primitives) including intra-concurrency control and use of hardware atomics.
6. Accurately analyze code to identify race conditions and appropriate solutions for addressing race conditions.

OS-Protection: Protection and Safety**CS Core:**

1. Overview of operating system security mechanisms. (See also: [SEC-Foundations](#))
2. Attacks and antagonism (scheduling, etc.). (See also: [SEC-Foundations](#))
3. Review of major vulnerabilities in real operating systems. (See also: [SEC-Foundations](#))
4. Operating systems mitigation strategies such as backups. (See also: [SF-Reliability](#))

KA Core:

5. Policy/mechanism separation. (See also: [SEC-Governance](#))
6. Security methods and devices. (See also: [SEC-Foundations](#))
Example concepts:
 - a. Rings of protection (history from Multics to virtualized x86), and
 - b. x86_64 rings -1 and -2 (hypervisor and ME/PSP).
7. Protection, access control, and authentication. (See also: [SEC-Foundations](#), [SEC-Crypto](#))

Illustrative Learning Outcomes:**CS Core:**

1. Understand the requirement for protection and security mechanisms in operating systems.
2. List and describe the attack vectors that leverage OS vulnerabilities.
3. Understand the mechanisms available in an OS to control access to resources.

KA Core:

4. Summarize the features and limitations of an operating system that impact protection and security.

OS-Scheduling: Scheduling**KA Core:**

1. Preemptive and non-preemptive scheduling.
2. Schedulers and policies (e.g., first come, first serve, shortest job first, priority, round robin, multilevel). (See also: [SF-Resource](#))
3. Concepts of Symmetric Multi-Processor (SMP) scheduling and cache coherence. (See also: [AR-Memory](#))
4. Timers (e.g., building many timers out of finite hardware timers). (See also: [AR-Assembly](#))
5. Fairness and starvation.

Non-Core:

6. Subtopics of operating systems such as energy-aware scheduling and real-time scheduling. (See also: [AR-Performance-Energy](#), [SPD-Embedded](#), [SPD-Mobile](#))
7. Cooperative scheduling, such as Linux futexes and userland scheduling.

Illustrative Learning Outcomes:**KA Core:**

1. Compare and contrast the common algorithms used for both preemptive and non-preemptive scheduling of tasks in operating systems, such as priority, performance comparison, and fair-share schemes.
2. Explain the relationships between scheduling algorithms and application domains.
3. Explain the distinctions among types of processor scheduler such as short-term, medium-term, long-term, and I/O.
4. Evaluate a problem or solution to determine appropriateness for asymmetric and/or symmetric multiprocessing.
5. Evaluate a problem or solution to determine appropriateness as a process vs threads.
6. List some contexts benefitting from preemption and deadline scheduling.

Non-Core:

7. Explain the ways that the logic embodied in scheduling algorithms are applicable to other operating systems mechanisms, such as first come first serve or priority to disk I/O, network scheduling, project scheduling, and problems beyond computing.

OS-Process: Process Model**KA Core:**

1. Processes and threads relative to virtualization - protected memory, process state, memory isolation, etc.
2. Memory footprint/segmentation (e.g., stack, heap, etc.). (See also: [AR-Assembly](#))
3. Creating and loading executables, shared libraries, and dynamic linking. (See also: [FPL-Translation](#))
4. Dispatching and context switching. (See also: [AR-Assembly](#))
5. Interprocess communication (e.g., shared memory, message passing, signals, environment variables). (See also: [PDC-Communication](#))

Illustrative Learning Outcomes:**KA Core:**

1. Understand how processes and threads use concurrency features to virtualize control.
2. Understand reasons for using interrupts, dispatching, and context switching to support concurrency and virtualization in an operating system.
3. Understand the different states that a task may pass through and the data structures needed to support the management of many tasks.
4. Understand the different ways of allocating memory to tasks, citing the relative merits of each.

5. Apply the appropriate interprocess communication mechanism for a specific purpose in a programmed software artifact.

OS-Memory: Memory Management

KA Core:

1. Review of physical memory, address translation and memory management hardware. (See also: [AR-Memory](#), [MSF-Discrete](#))
2. Impact of memory hierarchy including cache concept, cache lookup, and per-CPU caching on operating system mechanisms and policy. (See also: [AR-Memory](#), [SF-Performance](#))
3. Logical and physical addressing, address space virtualization. (See also: [AR-Memory](#), [MSF-Discrete](#))
4. Concepts of paging, page replacement, thrashing and allocation of pages and frames.
5. Allocation/deallocation/storage techniques (algorithms and data structure) performance and flexibility.
Example concept: Arenas, slab allocators, free lists, size classes, heterogeneously sized pages (huge pages).
6. Memory caching and cache coherence and the effect of flushing the cache to avoid speculative execution vulnerabilities. (See also: [AR-Organization](#), [AR-Memory](#), [SF-Performance](#))
7. Security mechanisms and concepts in memory management including sandboxing, protection, isolation, and relevant vectors of attack. (See also: [SEC-Foundations](#))

Non-Core:

8. Virtual memory: leveraging virtual memory hardware for OS services and efficiency.

Illustrative Learning Outcomes:

KA Core:

1. Explain memory hierarchy and cost-performance tradeoffs.
2. Summarize the principles of virtual memory as applied to caching and paging.
3. Evaluate the tradeoffs in terms of memory size (main memory, cache memory, auxiliary memory) and processor speed.
4. Describe the reason for and use of cache memory (performance and proximity, how caches complicate isolation and virtual machine abstraction).
5. Code/Develop efficient programs that consider the effects of page replacement and frame allocation on the performance of a process and the system in which it executes.

Non-Core:

6. Explain how hardware is utilized for efficient virtualization.

OS-Devices: Device management

KA Core:

1. Buffering strategies. (See also: [AR-IO](#))
2. Direct Memory Access (DMA) and polled I/O, Memory-mapped I/O. (See also: [AR-IO](#))
Example concept: DMA communication protocols (e.g., ring buffers etc.).

3. Historical and contextual - Persistent storage device management (e.g., magnetic, Solid State Device (SSD)). (See also: [SEP-History](#))

Non-Core:

4. Device interface abstractions, hardware abstraction layer.
5. Device driver purpose, abstraction, implementation and testing challenges.
6. High-level fault tolerance in device communication.

Illustrative Learning Outcomes:

KA Core:

1. Explain architecture level device control implementation and link relevant operating system mechanisms and policy (e.g., buffering strategies, direct memory access).
2. Explain OS device management layers and the architecture (e.g., device controller, device driver, device abstraction).
3. Explain the relationship between the physical hardware and the virtual devices maintained by the operating system.
4. Explain I/O data buffering and describe strategies for implementing it.
5. Describe the advantages and disadvantages of direct memory access and discuss the circumstances in which its use is warranted.

Non-Core:

6. Describe the complexity and best practices for the creation of device drivers

OS-Files: File Systems API and Implementation

KA Core:

1. Concept of a file including data, metadata, operations and access-mode.
2. File system mounting.
3. File access control.
4. File sharing.
5. Basic file allocation methods including linked, allocation table.
6. File system structures comprising file allocation including various directory structures and methods for uniquely identifying files (e.g., name, identified or metadata storage location).
7. Allocation/deallocation/storage techniques (algorithms and data structure) impact on performance and flexibility (i.e. Internal and external fragmentation and compaction).
8. Free space management such as using bit tables vs linking.
9. Implementation of directories to segment and track file location.

Illustrative Learning Outcomes:

KA Core:

1. Explain the choices to be made in designing file systems.
2. Evaluate different approaches to file organization, recognizing the strengths and weaknesses of each.
3. Apply software constructs appropriately given knowledge of the file system implementation.

OS-Advanced-Files: Advanced File systems

KA Core:

1. File systems: partitioning, mount/unmount, virtual file systems.
2. In-depth implementation techniques.
3. Memory-mapped files. (See also AR-IO)
4. Special-purpose file systems.
5. Naming, searching, access, backups
6. Journaling and log-structured file systems. (See also [SF-Reliability](#))

Non-Core: (including emerging topics)

1. Distributed file systems.
2. Encrypted file systems.
3. Fault tolerance.

Illustrative Learning Outcomes:

KA Core:

1. Explain how hardware developments have led to changes in the priorities for the design and the management of file systems.
2. Map file abstractions to a list of relevant devices and interfaces.
3. Identify and categorize different mount types.
4. Explain specific file systems requirements and the specialize file systems features that meet those requirements.
5. Explain the use of journaling and how log-structured file systems enhance fault tolerance.

Non-Core:

6. Explain purpose and complexity of distributed file systems.
7. List examples of distributed file systems protocols.
8. Explain mechanisms in file systems to improve fault tolerance.

OS-Virtualization: Virtualization

KA Core:

1. Using virtualization and isolation to achieve protection and predictable performance. (See also: [SF-Performance](#))
2. Advanced paging and virtual memory. (See also: [SF-Performance](#))
3. Virtual file systems and virtual devices.
4. Containers and their comparison to virtual machine.
5. Thrashing (e.g., Popek and Goldberg requirements for recursively virtualizable systems).

Non-core:

6. Types of virtualization (including hardware/software, OS, server, service, network). (See also: [SF-Performance](#))
7. Portable virtualization; emulation vs isolation. (See also: [SF-Performance](#))
8. Cost of virtualization. (See also: [SF-Performance](#))

9. Virtual machines and container escapes, dangers from a security perspective. (See also: [SEC-Engineering](#))
10. Hypervisors- hardware virtual machine extensions, hosts with kernel support, QEMU KVM

Illustrative Learning Outcomes:

KA Core:

1. Explain how hardware architecture provides support and efficiencies for virtualization.
2. Explain the difference between emulation and isolation.
3. Evaluate virtualization tradeoffs.

Non-Core:

4. Explain hypervisors and the need for them in conjunction with different types of hypervisors.

OS-Real-time: Real-time and Embedded Systems

KA Core:

1. Process and task scheduling.
2. Deadlines and real-time issues. (See also: [SPD-Embedded](#))
3. Low-latency vs "soft real-time" vs "hard real time". (See also: [SPD-Embedded](#), [FPL-Event-Driven](#))

Non-Core:

4. Memory/disk management requirements in a real-time environment.
5. Failures, risks, and recovery.
6. Special concerns in real-time systems (safety).

Illustrative Learning Outcomes:

KA Core:

1. Explain what makes a system a real-time system.
2. Explain latency and its sources in software systems and its characteristics.
3. Explain special concerns that real-time systems present, including risk, and how these concerns are addressed.

Non-Core:

4. Explain specific real time operating systems features and mechanisms.

OS-Faults: Fault tolerance

KA Core:

1. Reliable and available systems. (See also: [SF-Reliability](#))
2. Software and hardware approaches to address tolerance (RAID). (See also: SF-Reliability)

Non-Core:

3. Spatial and temporal redundancy. (See also: [SF-Reliability](#))
4. Methods used to implement fault tolerance. (See also: [SF-Reliability](#))

5. Error identification and correction mechanisms, checksums of volatile memory in RAM. (See also: [AR-Memory](#))
6. File system consistency check and recovery.
7. Journaling and log-structured file systems. (See also: [SF-Reliability](#))
8. Use-cases for fault-tolerance (databases, safety-critical). (See also: [SF-Reliability](#))
9. Examples of OS mechanisms for detection, recovery, restart to implement fault tolerance, use of these techniques for the OS's own services. (See also: [SF-Reliability](#))

Illustrative Learning Outcomes:

KA Core:

1. Explain how operating systems can facilitate fault tolerance, reliability, and availability.
2. Explain the range of methods for implementing fault tolerance in an operating system.
3. Explain how an operating system can continue functioning after a fault occurs.
4. Explain the performance and flexibility tradeoffs that impact using fault tolerance.

Non-Core:

5. Describe operating systems fault tolerance issues and mechanisms in detail.

OS-SEP: Society, Ethics, and the Profession

KA Core:

1. Open source in operating systems (See also: [SEP-IP](#))
Example concepts:
 - a. Identification of vulnerabilities in open source kernels,
 - b. Open source guest operating systems,
 - c. Open source host operating systems, and
 - d. Changes in monetization (paid vs free upgrades).
2. End-of-life issues with sunsetting operating systems.
Example concept: Privacy implications of using proprietary operating systems/operating environments, including telemetry, automated scanning of personal data, built-in advertising, and automatic cloud integration.

Illustrative Learning Outcomes:

KA Core:

1. Explain advantages and disadvantages of finding and addressing bugs in open source kernels,
2. Contextualize history and positive and negative impact of Linux as an open source product,
3. List complications with reliance on operating systems past end-of-life,
4. Understand differences in finding and addressing bugs for various operating systems payment models,

Professional Dispositions

- **Proactive:** Students must anticipate the security and performance implications of how operating systems components are used.

- **Meticulous:** Students must carefully analyze the implications of operating system mechanisms on any project.

Mathematics Requirements

Required:

- [MSF-Discrete](#)

Course Packaging Suggestions

Introductory Course to include the following:

- [OS-Purpose](#) (3 hours)
- [OS-Principles](#) (3 hours)
- [OS-Concurrency](#) (7 hours)
- [OS-Scheduling](#) (3 hours)
- [OS-Process](#) (3 hours)
- [OS-Memory](#) (4 hours)
- [OS-Protection](#) (4 hours)
- [OS-Devices](#) (2 hours)
- [OS-Files](#) (2 hours)
- [OS-Virtualization](#) (3 hours)
- [OS-Advanced-Files](#) (2 hours)
- [OS-Real-time](#) (1 hour)
- [OS-Faults](#) (1 hour)
- [OS-SEP](#) (4 hours)

Prerequisites:

- [AR-Assembly](#)
- [AR-Memory](#)
- [AR-Reliability](#)
- [AR-IO](#)
- [AR-Organization](#)
- [MSF-Discrete](#)

Course objectives: Students should understand the impact and implications of operating system resource management in terms of performance and security. They should understand and implement inter-process communication mechanisms safely. They should be able to differentiate between the use and evaluation of open source and/or proprietary operating systems. They should understand virtualization as a feature of safe modern operating system implementation.

Committee

Chair: Monica D. Anderson, University of Alabama, Tuscaloosa, AL, USA

Members:

- Qiao Xiang, Xiamen University, Xiamen, China
- Mikey Goldweber, Denison University, Granville, OH, USA
- Marcelo Pias, Federal University of Rio Grande (FURG), Rio Grande, RS, Brazil
- Avi Silberschatz, Yale University, New Haven, CT, USA
- Renzo Davoli, University of Bologna, Bologna, Italy

Parallel and Distributed Computing (PDC)

Preamble

Parallel and distributed programming arranges, coordinates, and controls multiple computations occurring at the same time across different places. The ubiquity of parallelism and distribution are inevitable consequences of increasing numbers of gates in processors, processors in computers, and computers everywhere that may be used to improve performance compared to sequential programs, while also coping with the intrinsic interconnectedness of the world, and the possibility that some components or connections fail or behave maliciously. Parallel and distributed programming removes the restrictions of sequential programming that require computational steps to occur in a serial order in a single place, revealing further distinctions, techniques, and analyses applying at each layer of computing systems.

In most conventional usage, “parallel” programming focuses on establishing and coordinating multiple activities that may occur at the same time, “distributed” programming focuses on establishing and coordinating activities that may occur in different places, and “concurrent” programming focuses on interactions of ongoing activities with each other and the environment. However, all three terms may apply in most contexts. Parallelism generally implies some form of distribution because multiple activities occurring without sequential ordering constraints happen in multiple physical places (unless they rely on context-switching or quantum effects). Conversely, actions in different places need not bear any particular sequential ordering with respect to each other in the absence of communication constraints.

Parallel, distributed and concurrent programming techniques form the core of High Performance Computing (HPC), distributed systems, and increasingly, nearly every computing application. The PDC knowledge area has evolved from a diverse set of advanced topics into a central body of knowledge and practice, permeating almost every other aspect of computing. Growth of the field has occurred irregularly across different subfields of computing, sometimes with different goals, terminology, and practices, masking the considerable overlap of basic ideas and skills that are the main focus of this knowledge area. Nearly every problem with a sequential solution also admits parallel and/or distributed solutions; additional problems and solutions arise only in the context of concurrency. Nearly every application domain of parallel and distributed computing is a well-developed area of study and/or engineering too large to enumerate.

Overview

This knowledge area is divided into five knowledge units, each with CS Core and KA Core topics that extend but do not overlap CS Core coverage that appears in other knowledge areas. The five knowledge units cover: The nature of parallel and distributed **Programs** and their execution; **Communication** (via channels, memory, or shared data stores), **Coordination** among parallel activities to achieve common outcomes; **Evaluation** with respect to specifications, and **Algorithms** across multiple application domains.

CS Core topics span approaches to parallel and distributed computing, but restrict coverage to those that apply to nearly all of them. Learning outcomes include developing small programs (in a choice of several styles) with multiple activities and analyzing basic properties. The topics and hours do not include coverage of particular languages, tools, frameworks, systems, and platforms needed as a basis for implementing and evaluating concepts and skills. The topics also avoid reliance on specifics that may vary widely (for example GPU programming vs cloud container deployment scripts), Prerequisites for CS Core coverage include:

- [SDF-Fundamentals](#): programs, executions, specifications, implementations, variables, arrays, sequential control flow, procedural abstraction and invocation, Input/Output.
- [SF-Overview](#): Layered systems, state machines, reliability.
- [AR-Assembly](#), [AR-Memory](#): von Neumann architecture, memory hierarchy.
- [MSF-Discrete](#): Discrete structures including directed graphs.

Additionally, Foundations of Programming Languages ([FPL](#)) may be treated as a prerequisite, depending on other curricular choices. The CS Core requires familiarity with languages and platforms that enable construction of parallel and distributed programs. Also, PDC includes definitions of safety, liveness, and related concepts that are covered with respect to language properties and semantics in FPL. Similarly, the PDC CS Core includes concepts that are further developed in the context of network protocols in Networking and Communication ([NC](#)), Operating Systems ([OS](#)) and Security ([SEC](#)), that could be covered in any order.

KA Core topics in each unit are of the form “one or more of the following” for *a la carte* topics extending associated core topics. Any selection of KA-core topics meeting the KA Core hour requirement constitutes fulfillment of the KA Core. This structure permits variation in coverage depending on the focus of any given course (see below for examples). Depth of coverage of any KA Core subtopic is expected to vary according to course goals. For example, shared-memory coordination is a central topic in multicore programming, but much less so in most heterogeneous systems, and conversely for bulk data transfer. Similarly, fault tolerance is central to the design of distributed information systems, but much less so in most data-parallel applications.

Changes since CS2013

This knowledge area has been refactored to focus on commonalities across different forms of parallel and distributed computing, also enabling more flexibility in KA Core coverage, with more guidance on coverage options.

Core Hours

Knowledge Units	CS Core hours	KA Core hours
Programs	2	2
Communication	2	6

Coordination	2	6
Evaluation	1	3
Algorithms	2	9
Society, Ethics, and the Profession	Included in SEP hours	
Total	9	26

Knowledge Units

PDC-Programs:Programs

CS Core:

1. Parallelism
 - a. Declarative parallelism: Determining which actions may, or must not, be performed in parallel, at the level of instructions, functions, closures, composite actions, sessions, tasks, and services is the main idea underlying PDC algorithms; failing to do so is the main source of errors. (See also: [PDC-Algorithms](#))
 - b. Defining order: for example, using happens-before relations or series/parallel directed acyclic graphs representing programs.
 - c. Independence: determining when ordering does not matter, in terms of commutativity, dependencies, preconditions.
 - d. Ensuring ordering among otherwise parallel actions when necessary, including locking, safe publication; and imposing communication – sending a message happens before receiving it; conversely relaxing when unnecessary.
2. Distribution
 - a. Defining places, as devices executing actions, including hardware components, remote hosts, may also include external, uncontrolled devices, hosts, and users. (See also: [AR-IO](#))
 - b. One device may time-slice or otherwise emulate multiple parallel actions by fewer processors by scheduling and virtualization. (See also: [OS-Scheduling](#))
 - c. Naming or identifying places (e.g., device IDs) and actions as parties (e.g., thread IDs).
 - d. Activities across places may communicate across media. (See also: [PDC-Communication](#))
3. Starting activities
 - a. Options that enable actions to be performed (eventually) at places range from hardwiring to configuration scripts; also establishing communication and resource management; these are expressed differently across languages and contexts, usually relying on automated provisioning and management by platforms. (See also: [SF-Resources](#))
 - b. Procedural: Enabling multiple actions to start at a given program point; for example, starting new threads, possibly scoping or otherwise organizing them in hierarchical groups.
 - c. Reactive: Enabling upon an event by installing an event handler, with less control of when actions begin or end, and may apply even on uniprocessors.
 - d. Dependent: Enabling upon completion of others; for example, sequencing sets of parallel actions (See also: [PDC-Coordination](#))

- e. Granularity: Execution cost of action bodies should outweigh the overhead of arranging them.
4. Execution Properties
 - a. Nondeterministic execution of unordered actions.
 - b. Consistency: Ensuring agreement among parties about values and predicates when necessary to avoid races, maintain safety and atomicity, or arrive at consensus
 - c. Fault tolerance: Handling failures in parties or communication, including (Byzantine) misbehavior due to untrusted parties and protocols, when necessary to maintain progress or availability. (See also: [SF-Reliability](#))
 - d. Tradeoffs are one focus of evaluation. (See also: [PDC-Evaluation](#))

KA Core:

5. One or more of the following mappings and mechanisms across layered systems:
 - a. CPU data- and instruction-level- parallelism. (See also: [AR-Organization](#))
 - b. SIMD and heterogeneous data parallelism. (See also: [AR-Heterogeneity](#))
 - c. Multicore scheduled concurrency, tasks, actors. (See also: [OS-Scheduling](#))
 - d. Clusters, clouds; elastic provisioning. (See also: [SPD-Common](#))
 - e. Networked distributed systems. (See also: [NC-Applications](#))
 - f. Emerging technologies such as quantum computing and molecular computing.

Illustrative Learning Outcomes

CS Core:

1. Graphically show (as a Directed Acyclic Graph (DAG)) how to parallelize a compound numerical expression; for example, $a = (b+c) * (d + e)$.
2. Explain why the concepts of consistency and fault tolerance do not arise in purely sequential programs.

KA Core:

3. Write a function that efficiently counts events such as networking packet receptions.
4. Write a filter/map/reduce program in multiple styles.
5. Write a service that creates a thread (or other procedural form of activation) to return a requested web page to each new client.

PDC-Communication: Communication

CS Core:

1. Media
 - a. Varieties: channels (message passing or I/O), shared memory, heterogeneous, data stores
 - b. Reliance on the availability and nature of underlying hardware, connectivity, and protocols; language support, emulation. (See also: [AR-IO](#))
2. Channels
 - a. Explicit (usually named) party-to-party communication media
 - b. APIs: Sockets, architectural, language-based, and toolkit constructs, such as Message Passing Interface (MPI), and layered constructs such as Remote Procedure Call (RPC). (See also: [NC-Fundamentals](#))
 - c. I/O channel APIs

3. Memory
 - a. Shared memory architectures in which parties directly communicate only with memory at given addresses, with extensions to heterogeneous memory supporting multiple memory stores with explicit data transfer across them; for example, GPU local and shared memory, Direct Memory Access (DMA).
 - b. Memory hierarchies: Multiple layers of sharing domains, scopes and caches; locality: latency, false-sharing.
 - c. Consistency properties: Bitwise atomicity limits, coherence, local ordering.
4. Data Stores
 - a. Cooperatively maintained data structures implementing maps and related ADTs
 - b. Varieties: Owned, shared, sharded, replicated, immutable, versioned

KA Core:

5. One or more of the following properties and extensions:
 - a. Topologies: Unicast, Multicast, Mailboxes, Switches; Routing via hardware and software interconnection networks.
 - b. Media concurrency properties: Ordering, consistency, idempotency, overlapping communication with computation.
 - c. Media performance: Latency, bandwidth (throughput) contention (congestion), responsiveness (liveness), reliability (error and drop rates), protocol-based progress (acks, timeouts, mediation).
 - d. Media security properties: integrity, privacy, authentication, authorization. (See also: [SEC-Secure Coding](#))
 - e. Data formats: Marshaling, validation, encryption, compression.
 - f. Channel policies: Endpoints, sessions, buffering, saturation response (waiting vs dropping), Rate control.
 - g. Multiplexing and demultiplexing many relatively slow I/O devices or parties; completion-based and scheduler-based techniques; async-await, select and polling APIs.
 - h. Formalization and analysis of channel communication; for example, CSP.
 - i. Applications of queuing theory to model and predict performance.
 - j. Memory models: sequential and release/acquire consistency.
 - k. Memory management; including reclamation of shared data; reference counts and alternatives.
 - l. Bulk data placement and transfer; reducing message traffic and improving locality; overlapping data transfer and computation; impact of data layout such as array-of-structs vs struct-of-arrays.
 - m. Emulating shared memory: distributed shared memory, Remote Direct Memory Access (RDMA).
 - n. Data store consistency: Atomicity, linearizability, transactionality, coherence, causal ordering, conflict resolution, eventual consistency, blockchains.
 - o. Faults, partitioning, and partial failures; voting; protocols such as Paxos and Raft.
 - p. Design tradeoffs among consistency, availability, partition (fault) tolerance; impossibility of meeting all at once.
 - q. Security and trust: Byzantine failures, proof of work and alternatives.

Illustrative Learning Outcomes

CS Core:

1. Explain the similarities and differences among: (1) Party A sends a message on channel X with contents 1 received by party B (2) A sets shared variable X to 1, read by B (3) A sets "X=1" in a distributed shared map accessed by B.

KA Core:

2. Write a program that distributes different segments of a data set to multiple workers, and collects results (for the simplest example, summing segments of an array).
3. Write a parallel program that requests data from multiple sites, and summarizes them using some form of reduction.
4. Compare the performance of buffered versus unbuffered versions of a producer-consumer program.
5. Determine whether a given communication scheme provides sufficient security properties for a given usage.
6. Give an example of an ordering of accesses among concurrent activities (e.g., program with a data race) that is not sequentially consistent.
7. Give an example of a scenario in which blocking message sends can deadlock.
8. Describe at least one design technique for avoiding liveness failures in programs using multiple locks..
9. Write a program that illustrates memory-access or message reordering.
10. Describe the relative merits of optimistic versus conservative concurrency control under different rates of contention among updates.
11. Give an example of a scenario in which an attempted optimistic update may never complete.
12. Modify a concurrent system to use a more scalable, reliable or available data store
13. Using an existing platform supporting replicated data stores, write a program that maintains a key-value mapping even when one or more hosts fail.

PDC-Coordination: Coordination

CS Core:

1. Dependencies
 - a. Initiation or progress of one activity may be dependent on other activities, so as to avoid race conditions, ensure termination, or meet other requirements.
 - b. Ensuring progress by avoiding dependency cycles, using monotonic conditions, removing inessential dependencies.
2. Control constructs and design patterns
 - a. Completion-based: Barriers, joins, including termination control.
 - b. Data-enabled: Queues, producer-consumer designs.
 - c. Condition-based: Polling, retrying, backoffs, helping, suspension, signaling, timeouts.
 - d. Reactive: enabling and triggering continuations.
3. Atomicity
 - a. Atomic instructions, enforced local access orderings.
 - b. Locks and mutual exclusion; lock granularity.

- c. Using locks in a given language; maintaining liveness without introducing races.
- d. Deadlock avoidance: Ordering, coarsening, randomized retries; backoffs, encapsulation via lock managers.
- e. Common errors: Failing to lock or unlock when necessary, holding locks while invoking unknown operations.
- f. Avoiding locks: replication, read-only, ownership, and non-blocking constructions.

KA Core:

- 4. One or more of the following properties and extensions:
 - a. Progress properties including lock-free, wait-free, fairness, priority scheduling; interactions with consistency, reliability.
 - b. Performance with respect to contention, granularity, convoying, scaling.
 - c. Non-blocking data structures and algorithms.
 - d. Ownership and resource control.
 - e. Lock variants and alternatives: sequence locks, read-write locks; Read-Copy-Update (RCU), reentrancy; tickets; controlling spinning versus blocking.
 - f. Transaction-based control: Optimistic and conservative.
 - g. Distributed locking: reliability.
 - h. Alternatives to barriers: Clocks; counters, virtual clocks; dataflow and continuations; futures and RPC; consensus-based, gathering results with reducers and collectors.
 - i. Speculation, selection, cancellation; observability and security consequences.
 - j. Resource control using semaphores and condition variables.
 - k. Control flow: Scheduling computations, series-parallel loops with (possibly elected) leaders, pipelines and streams, nested parallelism.
 - l. Exceptions and failures. Handlers, detection, timeouts, fault tolerance, voting.

Illustrative Learning Outcomes

CS Core:

- 1. Show how to avoid or repair a race error in a given program.
- 2. Show how to ensure that a program correctly terminates when all of a set of concurrent tasks have completed.

KA Core:

- 3. Write a function that efficiently counts events such as sensor inputs or networking packet receptions.
- 4. Write a filter/map/reduce program in multiple styles.
- 5. Write a program in which the termination of one set of parallel actions is followed by another.
- 6. Write a program that speculatively searches for a solution by multiple activities, terminating others when one is found.
- 7. Write a program in which a numerical exception (such as divide by zero) in one activity causes termination of others.
- 8. Write a program for multiple parties to agree upon the current time of day; discuss its limitations compared to protocols such as network transfer protocol (NTP).

9. Write a service that creates a thread (or other procedural form of activation) to return a requested web page to each new client.

PDC-Evaluation: Evaluation

CS Core:

1. Safety and liveness requirements in terms of temporal logic constructs to express “always” and “eventually”. (See also: [FPL-Parallel](#))
2. Identifying, testing for, and repairing violations, including common forms of errors such as failure to ensure necessary ordering (race errors), atomicity (including check-then-act errors), and termination (livelock).
3. Performance requirements metrics for throughput, responsiveness, latency, availability, energy consumption, scalability, resource usage, communication costs, waiting and rate control, fairness; service level agreements. (See also: [SF-Performance](#))
4. Performance impact of design and implementation choices, including granularity, overhead, consensus costs, and energy consumption. (See also: [SEP-Sustainability](#))
5. Estimating scalability limitations, for example using Amdahl’s Law or Universal Scalability Law. (See also: [SF-Evaluation](#))

KA Core:

6. One or more of the following methods and tools:
 - a. Extensions to formal sequential requirements such as linearizability.
 - b. Protocol, session, and transactional specifications.
 - c. Use of tools such as Unified Modelling Language (UML), Temporal Logic of Actions (TLA), program logics.
 - d. Security analysis: safety and liveness in the presence of hostile or buggy behaviors by other parties; required properties of communication mechanisms (for example lack of cross-layer leakage), input screening, rate limiting. (See also: [SEC-Foundations](#))
 - e. Static analysis applied to correctness, throughput, latency, resources, energy. (See also [SEP-Sustainability](#))
 - f. Directed Acyclic Graph (DAG) model analysis of algorithmic efficiency (work, span, critical paths)
 - g. Testing and debugging; tools such as race detectors, fuzzers, lock dependency checkers, unit/stress/torture tests, visualizations, continuous integration, continuous deployment, and test generators,
 - h. Measuring and comparing throughput, overhead, waiting, contention, communication, data movement, locality, resource usage, behavior in the presence of excessive numbers of events, clients, or threads. (See also [SF-Evaluation](#))
 - i. Application domain specific analyses and evaluation techniques.

Illustrative Learning Outcomes

CS Core:

1. Revise a specification to enable parallelism and distribution without violating other essential properties or features.
2. Explain how concurrent notions of safety and liveness extend their sequential counterparts.

3. Specify a set of invariants that must hold at each bulk-parallel step of a computation.
4. Write a test program that can reveal a data race error; for example, missing an update when two activities both try to increment a variable.
5. In a given context, explain the extent to which introducing parallelism in an otherwise sequential program would be expected to improve throughput and/or reduce latency, and how it may impact energy efficiency.
6. Show how scaling and efficiency change for sample problems without and with the assumption of problem size changing with the number of processors; further explain whether and how scalability would change under relaxations of sequential dependencies.

KA Core:

7. Specify and measure behavior when a service is requested by unexpectedly many clients
8. Identify and repair a performance problem due to sequential bottlenecks
9. Empirically compare throughput of two implementations of a common design (perhaps using an existing test harness framework).
10. Identify and repair a performance problem due to communication or data latency.
11. Identify and repair a performance problem due to communication or data latency.
12. Identify and repair a performance problem due to resource management overhead.
13. Identify and repair a reliability or availability problem.

PDC-Algorithms: Algorithms

CS Core:

1. Expressing and implementing algorithms in given languages and frameworks, to initiate activities (for example threads), use shared memory constructs, and channel, socket, and/or remote procedure call APIs. (See also: [FPL-Parallel](#)).
 - a. Data parallel examples including map/reduce.
 - b. Using channel, socket, and/or RPC APIs in a given language, with program control for sending (usually procedural) vs receiving. (usually reactive or RPC-based).
 - c. Using locks, barriers, and/or synchronizers to; maintain liveness without introducing races.
2. Survey of common application domains across multicore, reactive, data parallel, cluster, cloud, open distributed systems and frameworks (with reference to the following table):

Category	Typical Execution agents	Typical Communication mechanisms	Typical Algorithmic domains	Typical Engineering goals
Multicore	Threads	Shared memory, Atomics, locks	Resource management, data processing	Throughput, latency, energy
Reactive	Handlers, threads	I/O Channels	Services, real-time	Latency

Data parallel	GPU, SIMD, accelerators, hybrid	Heterogeneous memory	Linear algebra, graphics, data analysis	Throughput, energy
Cluster	Managed hosts	Sockets, channels	Simulation, data analysis	Throughput
Cloud	Provisioned hosts	Service APIs	Web applications	Scalability
Open distributed	Autonomous hosts	Sockets, Data stores	Fault tolerant data stores and services	Reliability

KA Core:

3. One of more of the following algorithmic domains. (See also: [AL-Strategies](#)):
 - a. Linear algebra: Vector and matrix operations, numerical precision/stability, applications in data analytics and machine learning
 - b. Data processing: sorting, searching and retrieval, concurrent data structures
 - c. Graphs, search, and combinatorics: Marking, edge-parallelization, bounding, speculation, network-based analytics
 - d. Modeling and simulation: differential equations; randomization, N-body problems, genetic algorithms
 - e. Computational logic: satisfiability (SAT), concurrent logic programming
 - f. Graphics and computational geometry: Transforms, rendering, ray-tracing
 - g. Resource management: Allocating, placing, recycling and scheduling processors, memory, channels, and hosts; exclusive vs shared resources; static, dynamic and elastic algorithms; Real-time constraints; Batching, prioritization, partitioning; decentralization via work-stealing and related techniques;
 - h. Services: Implementing web APIs, electronic currency, transaction systems, multiplayer games.

Illustrative Learning Outcomes

CS Core:

1. Implement a parallel/distributed component based on a known algorithm
2. Write a data-parallel program that for example computes the average of an array of numbers.
3. Write a producer-consumer program in which one component generates numbers, and another computes their average. Measure speedups when the numbers are small scalars versus large multi-precision values.
4. Extend an event-driven sequential program by establishing a new activity in an event handler (for example a new thread in a GUI action handler).
5. Improve the performance of a sequential component by introducing parallelism and/or distribution
6. Choose among different parallel/distributed designs for components of a given system.

KA Core:

7. Design, implement, analyze, and evaluate a component or application for X operating in a given context, where X is in one of the listed domains; for example, a genetic algorithm for factory floor design.
8. Critique the design and implementation of an existing component or application, or one developed by classmates.
9. Compare the performance and energy efficiency of multiple implementations of a similar design; for example, multicore versus clustered versus GPU.

Professional Dispositions

- **Meticulous:** Students' attention to detail is essential when applying constructs with non-obvious correctness conditions.
- **Persistent:** Students must be prepared to try alternative approaches when solutions are not self-evident.

Mathematics Requirements

Required:

- [MSF-Discrete](#) – Logic, discrete structures including directed graphs.

Desired:

- [MSF-Linear](#)
- [MSF-Calculus](#) – Differential equations

Course Packaging Suggestions

The CS Core requirements need not be provided by a single course. They may be included across courses primarily devoted to software development, programming languages, systems, data management, networking, computer architecture, and/or algorithms.

Alternatively, the CS Core provides a basis for courses focusing on parallel and/or distributed computing. At one extreme, it is possible to offer a single broadly constructed course covering all PDC KA Core topics to varying depths. At the other extreme, it is possible to infuse PDC KA Core coverage across the curriculum with courses that cover parallel and distributed approaches alongside sequential ones for nearly every topic in computing. More conventional choices include courses that focus on one or a few categories (such as multicore or cluster), and algorithmic domains (such as linear algebra, or resource management). Such courses may go into further depth than listed in one or more KUs, and include additional software development experience, but include only CS-Core-level coverage of other topics.

As an example, a course mainly focusing on multicores could extend CS Core topics as follows:

1. Programs: KA Core on threads, tasks, instruction-level parallelism
2. Communication: KA Core on multicore architectures, memory, concurrent data stores

3. Coordination: KA Core on blocking and non-blocking synchronization, speculation, cancellation, futures, and divide-and-conquer data parallelism
4. Evaluation: KA Core on performance analysis
5. Algorithms: project-based KA Core coverage of data processing and resource management.

More extensive examples and guidance for courses focusing on HPC are provided by the NSF/IEEE-TCPP Curriculum Initiative on Parallel and Distributed Computing [12].

Committee

Chair: Doug Lea, State University of New York at Oswego, Oswego, NY, USA

Members:

- Sherif Aly, American University of Cairo, Cairo, Egypt
- Michael Oudshoorn, High Point University, High Point, NC, USA
- Qiao Xiang, Xiamen University, Xiamen, China
- Dan Grossman, University of Washington, Seattle, WA, USA
- Sebastian Burckhardt, Microsoft Research, Redmond WA, USA
- Vivek Sarkar, Georgia Tech, Atlanta, GA, USA
- Maurice Herlihy, Brown University, Providence, RI, USA
- Sheikh Ghafoor, Tennessee Tech, Cookeville, TN, USA
- Chip Weems, University of Massachusetts, Amherst, MA, USA

Contributors:

- Paul McKenney, Meta, Beaverton, OR, USA
- Peter Buhr, University of Waterloo, Waterloo, Ontario, Canada

Software Development Fundamentals (SDF)

Preamble

Fluency in the process of software development is fundamental to the study of computer science. In order to use computers to solve problems most effectively, students must be competent at reading and writing programs. Beyond programming skills, however, they must be able to select and use appropriate data structures and algorithms, and use modern development and testing tools.

The SDF knowledge area brings together fundamental concepts and skills related to software development, focusing on concepts and skills that should be taught early in a computer science program, typically in the first year. This includes fundamental programming concepts and their effective use in writing programs, use of fundamental data structures which may be provided by the programming language, basics of programming practices for writing good quality programs, reading and understanding programs, and some understanding of the impact of algorithms on the performance of the programs. The 43 hours of material in this knowledge area may be augmented with core material from other knowledge areas as a student progresses to mid- and upper-level courses.

This knowledge area assumes a contemporary programming language with built-in support for common data types including associative data types like dictionaries/maps as the vehicle for introducing students to programming (e.g. Python, Java). However, this is not to discourage the use of older or lower-level languages for SDF – the knowledge units below can be suitably adapted for the actual language used.

The emergence of generative AI and Large Language Models (LLMs), which can generate programs for many programming tasks, will undoubtedly affect the programming profession and consequently the teaching of many CS topics. However, to be able to effectively use generative AI in programming tasks, a programmer must have a good understanding of programs, and hence must still learn the foundations of programming and develop basic programming skills - which is the aim of SDF. Consequently, we feel that the desired outcomes for SDF should remain the same, though different instructors may now give more emphasis to program understanding, documenting, specifications, analysis, and testing. (This is similar to teaching students multiplication and tables, addition, etc. even though calculators can do all this).

Changes since CS 2013

The main change from 2013 is a stronger emphasis on developing fundamental programming skills and effective use of in-built data structures (which many contemporary languages provide) for problem solving.

Overview

This Knowledge Area has five Knowledge Units. These are:

1. **SDF-Fundamentals**: Fundamental Programming Concepts and Practices: This knowledge unit aims to develop an understanding of basic concepts, and the ability to fluently use basic

language constructs as well as modularity constructs. It also aims to familiarize students with the concept of common libraries and frameworks, including those to facilitate API-based access to resources.

2. **[SDF-Data-Structures](#)**: Fundamental Data Structures: This knowledge unit aims to develop core concepts relating to Data Structures and associated operations. Students should understand the important data structures available in the programming language or as libraries, and how to use them effectively, including choosing appropriate data structures while designing solutions for a given problem.
3. **[SDF-Algorithms](#)**: Algorithms: This knowledge unit aims to develop the foundations of algorithms and their analysis. The KU should also empower students in selecting suitable algorithms for building modest-complexity applications.
4. **[SDF-Practices](#)**: Software Development Practices: This knowledge unit develops the core concepts relating to modern software development practices. It aims to develop student understanding and basic competencies in program testing, enhancing the readability of programs, and using modern methods and tools including some general-purpose IDE.
5. **[SDF-SEP](#)**: Society, Ethics, and the Profession: This knowledge unit aims to develop an initial understanding of some of the ethical issues related to programming, professional values programmers need to have, and the responsibility to society that programmers have. This knowledge unit is a part of the SEP Knowledge Area.

Core Hours

Knowledge Units	CS Core	KA Core
Fundamental Programming Concepts and Practices	20	
Fundamental Data Structures	6 + 6 (AL)	
Algorithms	3 + 3 (AL)	
Software Development Practices	5	
Society, Ethics and the Profession	Included in SEP hours	
Total	43	

Note: The CS Core hours includes 9 hours shared with AL.

Knowledge Units

[SDF-Fundamentals: Fundamental Programming Concepts and Practices](#)

CS Core:

1. Basic concepts such as variables, primitive data types, expressions and their evaluation.
2. How imperative programs work: state and state transitions on execution of statements, flow of control.
3. Basic constructs such as assignment statements, conditional and iterative statements, basic I/O.
4. Key modularity constructs such as functions (and methods and classes, if supported in the language) and related concepts like parameter passing, scope, abstraction, data encapsulation. (See also: [FPL-OOP](#))
5. Input and output using files and APIs.
6. Structured data types available in the chosen programming language like sequences (e.g., arrays, lists), associative containers (e.g., dictionaries, maps), others (e.g., sets, tuples) and when and how to use them. (See also: [AL-Foundational](#))
7. Libraries and frameworks provided by the language (when/where applicable).
8. Recursion.
9. Dealing with runtime errors in programs (e.g., exception handling).
10. Basic concepts of programming errors, testing, and debugging. (See also: [SE-Construction](#)) (See also: [SEC-Coding](#))
11. Documenting/commenting code at the program and module level. (See also: [SE-Construction](#))
12. Develop a security mindset. (See also: [SEC-Foundations](#))

Illustrative Learning Outcomes

CS Core:

In these learning outcomes, the term "develop" means "design, write, test and debug".

1. Develop programs that use the fundamental programming constructs: assignment and expressions, basic I/O, conditional and iterative statements.
2. Develop programs using functions with parameter passing.
3. Develop programs that effectively use the different structured data types provided in the language like arrays/lists, dictionaries, and sets.
4. Develop programs that use file I/O to provide data persistence across multiple executions.
5. Develop programs that use language-provided libraries and frameworks (where applicable).
6. Develop programs that use APIs to access or update data (e.g., from the web).
7. Develop programs that create simple classes and instantiate objects of those classes (if supported by the language).
8. Explain the concept of recursion, and identify when and how to use it effectively.
9. Develop recursive functions.
10. Develop programs that can handle runtime errors.
11. Read a given program and explain what it does.
12. Write comments for a program or a module specifying what it does.
13. Trace the flow of control during the execution of a program.
14. Use appropriate terminology to identify elements of a program (e.g., identifier, operator, operand).

SDF-Data-Structures: Fundamental Data Structures

CS Core: (See also: [AL-Foundational](#))

1. Standard abstract data types such as lists, stacks, queues, sets, and maps/dictionaries, and operations on them.
2. Selecting and using appropriate data structures.
3. Performance implications of choice of data structure(s).
4. Strings and string processing.

Illustrative Learning Outcomes

CS Core:

1. Write programs that use each of the key abstract data types provided in the language (e.g., arrays, tuples/records/structs, lists, stacks, queues, and associative data types like sets, dictionaries/maps).
2. Select the appropriate data structure for a given problem.
3. Explain how the performance of a program may change when using different data structures or operations.
4. Write programs that work with text by using string processing capabilities provided by the language.

SDF-Algorithms: Algorithms

CS Core: (See also: [AL-Foundational](#), [AL-Complexity](#))

1. Concept of algorithm and notion of algorithm efficiency.
2. Some common algorithms (e.g., sorting, searching, tree traversal, graph traversal).
3. Impact of algorithms on time-space efficiency of programs.

Illustrative Learning Outcomes

CS Core:

1. Explain the role of algorithms for writing programs.
2. Demonstrate how a problem may be solved by different algorithms, each with different properties.
3. Explain some common algorithms (e.g., sorting, searching, tree traversal, graph traversal).
4. Explain the impact on space/time performance of some algorithms.

SDF-Practices: Software Development Practices

CS Core: (See also: [SE-Construction](#))

1. Basic testing including test case design.
2. Use of a general-purpose IDE, including its debugger.
3. Programming style that improves readability.
4. Specifying functionality of a module in a natural language.

Illustrative Learning Outcomes

CS Core:

1. Develop tests for modules, and apply a variety of strategies to design test cases.
2. Explain some limitations of testing programs.
3. Build, execute and debug programs using a modern IDE and associated tools such as visual debuggers.
4. Apply basic programming style guidelines to aid readability of programs such as comments, indentation, proper naming of variables, etc.
5. Write specifications of a module as module comment describing its functionality.

SDF-SEP: Society, Ethics, and the Profession

CS Core:

1. Intellectual property rights of programmers for programs they develop.
2. Plagiarism & academic integrity.
3. Responsibility and liability of programmers regarding code they develop for solutions. (See also: SEC-Foundations)
4. Basic professional work ethics of programmers.

Illustrative Learning Outcomes

CS Core:

1. Explain/understand some of the intellectual property issues relating to programs
2. Explain/understand when code developed by others can be used and proper ways of disclosing their use
3. Explain/understand the responsibility of programmers when developing code for an overall solution (which may be developed by a team)
4. Explain/understand one or more codes of conduct applicable to programmers

Professional Dispositions

- **Self-Directed:** Students must seek out solutions to issues on their own (e.g., using technical forums, FAQs, discussions). Resolving issues is an important part of becoming proficient in programming.
- **Experimental:** Students must experiment with language features to understand them and to quickly prototype solutions. This helps in learning about programming language features.
- **Technical curiosity:** Students must develop interest in understanding how programs are executed, how programs and data are stored in memory, etc. This will help build better mental models of the underlying execution system on which programs run.
- **Adaptable:** Students must be willing to learn and use different tools and technologies that facilitate software development. Tools are commonly used while programming and new tools often emerge - using tools effectively and learning the use of new tools will help.
- **Persistent:** Students must continue efforts until, for example, a bug is identified, a program is made robust and handles all situations, etc. This will help as programming requires effort and ability to persevere till a program works satisfactorily.

- **Meticulous:** Students must pay attention to detail and use orderly processes while programming. The underlying machine is unforgiving and there is no room for even small errors in the programs as they can cause major failures.

Mathematics Requirements

As SDF focuses on the first year and is foundational, it assumes only basic mathematical knowledge that students acquire in school, in particular Sets, Relations, Functions, and Logic. (See also: [MSF-Discrete](#))

Course Packaging Suggestions

The SDF KA will generally be covered in introductory courses, often called CS1 and CS2. How much of the SDF KA can be covered in CS1 and how much is to be left for CS2 is likely to depend on the choice of programming language for CS1. For languages like Python or Java, CS1 can cover all the Programming Concepts and Development Methods KAs, and some of the Data Structures KA. It is desirable that they be further strengthened in CS2. The topics under algorithms KA and some topics under data structures KA can be covered in CS2. In case CS1 uses a language with fewer in-built data structures, then much of the Data Structures KA and some aspects of the programming KA may also need to be covered in CS2. With the former approach, the introductory course in programming can include the following:

1. [SDF-Fundamentals](#) (20 hours)
2. [SDF-Data-Structures](#) (12 hours)
3. [SDF-Algorithms](#) (6 hours)
4. [SDF-Practices](#) (5 hours)
5. [SDF-SEP](#)

Prerequisites: High school mathematics in particular Sets, Relations, Functions, and Logic. (See also: [MSF-Discrete](#))

Course objectives: Students should be able to:

- Design, code, test, and debug a modest sized program that effectively uses functional abstraction.
- Select and use the appropriate language provided data structure for a given problem (like: arrays, tuples/records/structs, lists, stacks, queues, and associative data types like sets, dictionaries/maps.)
- Design, code, test, and debug a modest-sized object-oriented program using classes and objects.
- Design, code, test, and debug a modest-sized program that uses language provided libraries and frameworks (including accessing data from the web through APIs).
- Read and explain given code including tracing the flow of control during execution.
- Write specifications of a program or a module in natural language explaining what it does.

- Build, execute and debug programs using a modern IDE and associated tools such as visual debuggers.
- Explain the key concepts relating to programming like parameter passing, recursion, runtime exceptions and exception handling.

Committee

Chair: Pankaj Jalote, Chair, IIT-Delhi, Delhi, India

Members:

- Brett A. Becker, University College Dublin, Dublin, Ireland
- Titus Winters, Google, New York City, NY, USA
- Andrew Luxton-Reilly, University of Auckland, Auckland, New Zealand
- Christian Servin, El Paso Community College, El Paso, TX, USA
- Karen Reid, University of Toronto, Toronto, Canada
- Adrienne Decker, University at Buffalo, Buffalo, NY, USA

Software Engineering (SE)

Preamble

As far back as the early 1970s, British computer scientist Brian Randell allegedly said, “Software engineering is the multi-person construction of multi-version programs.” This is an essential insight: while programming is the skill that governs our ability to write a program, software engineering is distinct in two dimensions: time and people.

First, a software engineering project is a team endeavor; being a solitary programming expert is insufficient. Skilled software engineers must demonstrate expertise in communication and collaboration. Programming may be an individual activity, but software engineering is a collaborative one, deeply tied to issues of professionalism, teamwork, and communication.

Second, a software engineering project is usually “multi-version.” It has an expected lifespan; it needs to function properly for months, years, or decades. Features may be added or removed to meet product requirements. The engineering team itself will likely change. The technological context will shift, as our computing platforms evolve, programming languages change, dependencies upgrade, etc. This exposure to matters of time and change is novel when compared to a programming project: it isn’t enough to build a thing that works, instead it must work and stay working. Many of the most challenging topics in tech share “time will lead to change” as a root cause: backward compatibility, version skew, dependency management, schema changes, protocol evolution.

Software engineering presents a particularly difficult challenge for learning in an academic setting. Given that the major differences between programming and software engineering are time and teamwork, it is hard to generate lessons that *require* successful teamwork and that faithfully present the

challenges of time. Additionally, some topics in software engineering will be more authentic and more relevant if and when our learners experience collaborative and long-term software engineering projects *in vivo* rather than in the classroom. Regardless of whether that happens as an internship, involvement in an open source project, or full-time engineering role, a month of full-time hands-on experience has more available hours than the average software engineering course.

Thus, a software engineering curriculum must focus on concepts needed by a majority of new-grad hires, and that either are novel for those who are trained primarily as programmers, or that are abstract concepts that may not get explicitly stated/shared on the job. Such topics include, but are not limited to:

- Testing
- Teamwork, collaboration
- Communication
- Design
- Maintenance and evolution
- Software engineering tools

Some such material is reasonably suited to a standard lecture or lecture+lab course. Discussing theoretical underpinnings of version control systems, or branching strategies in such systems, can be an effective way to familiarize students with those ideas. Similarly, a theoretical discussion can highlight the difference between static and dynamic analysis tools, or may motivate discussion of diamond dependency problems in dependency networks.

On the other hand, many of the fundamental topics of software engineering are best experienced in a hands-on fashion. Historically, project-oriented courses have been a common vehicle for such learning. We believe that such experience is valuable but also bears some interesting risks: students may form erroneous notions about the difficulty/complexity of collaboration if their only exposure is a single project with teams formed of other novice software engineers. It falls to instructors to decide on the right balance between theoretical material and hands-on projects - neither is a perfect vehicle for this challenging material. We strongly encourage instructors of project courses to aim for iteration and fast feedback - a few simple tasks repeated, as in an Agile-structured project, is better than singular high-friction introductions to many types of tasks. Projects with real-world industry partners and clients are also particularly encouraged. If long-running project courses are not an option, anything that can expose learners to the collaborative and long-term aspects of software engineering is valuable: adding features to an existing codebase, collaborating on distinct parts of a larger whole, pairing up to write an encoder and decoder, etc.

All evidence suggests that the role of software in our society will continue to grow for the foreseeable future. Additionally, the era of “two programmers in a garage” seems to have drawn to a close. Most important software these days is a team effort, building on existing code and leveraging existing functionality. The study of software engineering skills is a deeply important counterpoint to the everyday experience of computing students - we *must* impress on them the reality that few software projects are managed by writing from scratch as a solo endeavor. Communication, teamwork, planning, testing, and tooling are far more important as our students move on from the classroom and make their mark on the wider world.

Although most CS graduates will go on to an industry position that requires this material, the CS Core topics presented here are of value regardless of whether graduates go on to industry or academia.

Changes since CS 2013

This document shifts the focus of the Software Engineering knowledge area in a few ways compared to the goals of CS2013. The common reasoning behind most of these changes is to focus on material that learners would not pick up elsewhere in the curriculum, and that will be relevant *immediately* upon graduation, rather than at some future point in their careers.

- More explicit focus on the software workflow (version control, testing, code review, tooling)
- Less focus on team *leadership* and project management.
- More focus on team *participation*, communication, and collaboration

Overview

1. **SE-Teamwork:** Because of the nature of learning programming, most students in introductory SE have little or no exposure to the collaborative nature of SE. Practice (for instance in project work) may help, but lecture and discussion time spent on the value of clear, effective, and efficient communication and collaboration is essential for Software Engineering.
2. **SE-Tools:** Industry reliance on SE tools has exploded in the past generation, with version control becoming ubiquitous, testing frameworks growing in popularity, increased reliance on static and dynamic analysis in practice, and near-ubiquitous use of continuous integration systems. Increasingly powerful IDEs provide code searching and indexing capabilities, as well as small scale refactoring tools and integration with other SE tools. An understanding of the nature of these tools is broadly valuable - especially version control systems.
3. **SE-Requirements:** Knowing how to build something is of little help if we do not know what to build. Product Requirements (aka Requirements Engineering, Product Design, Product Requirements solicitation, Product Requirements Documents, etc.) introduces students to the processes surrounding the specification of the broad requirements governing development of a new product or feature.
4. **SE-Design:** While Product Requirements focuses on the user-facing functionality of a software system, Software Design focuses on the engineer-facing design of internal software components. This encompasses large design concerns such as software architecture, as well as small-scale design choices like API design.
5. **SE-Construction:** Software Construction focuses on practices that influence the direct production of software: use of tests, test driven development, coding style. More advanced topics extend into secure coding, dependency injection, work prioritization, etc.
6. **SE-Validation:** Software Verification and Validation focuses on how to improve the value of testing - understand the role of testing, failure modes, and differences between good tests and poor ones.
7. **SE-Refactoring:** Refactoring and Code Evolution focuses on refactoring and maintenance strategies, incorporating code health, use of tools, and backwards compatibility considerations.
8. **SE-Reliability:** Software Reliability aims to improve understanding of and attention to error cases, failure modes, redundancy, and reasoning about fault tolerance.

9. **SE-Formal:** Formal Methods provides mathematically rigorous mechanisms to apply to software, from specification to verification. (Prerequisites: Substantial dependence on core material from the Discrete Structures area, particularly knowledge units DS/Basic Logic and DS/Proof Techniques.)

Core Hours

Knowledge Units	CS Core	KA Core
Teamwork	2 + 3 (SEP)	2
Tools and Environments	1	3 + 1 (SDF)
Product Requirements	0 + 3 (SEP)	2
Software Design	1	4 + 2 (DM)
Software Construction	1 + 3 (SDF)	3 + 1 (SDF)
Software Verification and Validation	1	3
Refactoring and Code Evolution		2
Software Reliability		2
Formal Methods		
Total	6	21

Note: We have specifically highlighted Teamwork and Product Requirements as two Knowledge Units where SEP lessons are most directly obvious and applicable. Issues like impact on society, interaction with others, and social power disparities are pervasive in Software Engineering and should be woven into as many practical lessons as possible.

Knowledge Units

SE-Teamwork: Teamwork

CS Core:

1. Effective communication, including oral and written, as well as formal (email, docs, comments, presentations) and informal (team chat, meetings) (See also: [SEP-Communication](#))
2. Common causes of team conflict, and approaches for conflict resolution
3. Cooperative programming
 - a. Pair programming or Swarming
 - b. Code review
 - c. Collaboration through version control
4. Roles and responsibilities in a software team (See also: [SEP-Professional-Ethics](#))
 - a. Advantages of teamwork
 - b. Risks and complexity of such collaboration
5. Team processes
 - a. Responsibilities for tasks, effort estimation, meeting structure, work schedule
6. Importance of team diversity and inclusivity (See also: [SEP-Communication](#))

KA Core:

7. Interfacing with stakeholders, as a team
 - a. Management & other non-technical teams
 - b. Customers
 - c. Users
8. Risks associated with physical, distributed, hybrid and virtual teams
 - a. Including communication, perception, structure, points of failure, mitigation and recovery, etc.

Illustrative Learning Outcomes:**CS Core:**

1. Follow effective team communication practices.
2. Articulate the sources of, hazards of, and potential benefits of team conflict - especially focusing on the value of disagreeing about ideas or proposals without insulting people.
3. Facilitate a conflict resolution and problem solving strategy in a team setting.
4. Collaborate effectively in cooperative development/programming.
5. Propose and delegate necessary roles and responsibilities in a software development team.
6. Compose and follow an agenda for a team meeting.
7. Facilitate through involvement in a team project, the central elements of team building, establishing healthy team culture, and team management including creating and executing a team work plan.
8. Promote the importance of and benefits that diversity and inclusivity brings to a software development team

KA Core:

9. Reference the importance of, and strategies to, as a team, interface with stakeholders outside the team on both technical and non-technical levels.
10. Enumerate the risks associated with physical, distributed, hybrid and virtual teams and possible points of failure and how to mitigate against and recover/learn from failures.

SE-Tools: Tools and Environments**CS Core:**

1. Software configuration management and version control (See also: [SDF-Practices](#))
 - a. Configuration in version control, reproducible builds/configuration
 - b. Version control branching strategies. Development branches vs release branches. Trunk-based development.
 - c. Merging/rebasing strategies, when relevant.

KA Core:

2. Release management
3. Testing tools including static and dynamic analysis tools (See also: [SDF-Practices](#), [SEC-Coding](#))
4. Software process automation
 - a. Build systems - the value of fast, hermetic, reproducible builds, compare/contrast approaches to building a project
 - b. Continuous Integration (CI) - the use of automation and automated tests to do preliminary validation that the current head/trunk revision builds and passes (basic) tests

- c. Dependency management - updating external/upstream dependencies, package management, SemVer
- 5. Design and communication tools (docs, diagrams, common forms of design diagrams like UML)
- 6. Tool integration concepts and mechanisms (See also: [SDF-Practices](#))
- 7. Use of modern IDE facilities - debugging, refactoring, searching/indexing, ML-powered code assistants, etc. (See also: [SDF-Practices](#))

Illustrative Learning Outcomes:

CS Core:

1. Describe the difference between centralized and distributed software configuration management.
2. Describe how version control can be used to help manage software release management.
3. Identify configuration items and use a source code control tool in a small team-based project.

KA Core:

4. Describe how available static and dynamic test tools can be integrated into the software development environment.
5. Understand the use of CI systems as a ground-truth for the state of the team's shared code (build and test success).
6. Describe the issues that are important in selecting a set of tools for the development of a particular software system, including tools for requirements tracking, design modeling, implementation, build automation, and testing.
7. Demonstrate the capability to use software tools in support of the development of a software product of medium size.

SE-Requirements: Product Requirements

KA Core:

1. Describe functional requirements using, for example, use cases or user stories
 - a. Using at least one method of documenting and structuring functional requirements
 - b. Understanding how the method supports design and implementation
 - c. Strengths and weaknesses of using a particular approach
2. Properties of requirements including consistency, validity, completeness, and feasibility
3. Requirements elicitation
 - a. Sources of requirements, for example, users, administrators, or support personnel
 - b. Methods of requirement gathering, for example, surveys, interviews, or behavioral analysis
4. Non-functional requirements, for example, security, usability, or performance, aka Quality Attributes. (See also: [SEP-Sustainability](#))
5. Risk identification and management, including ethical considerations surrounding the proposed product (See also: [SEP-Professional-Ethics](#))
6. Communicating and/or formalizing requirement specifications

Non-core:

7. Prototyping
 - a. A tool for both eliciting and validating/confirming requirements

8. Product evolution
 - a. When requirements change, how to understand what effect that has and what changes need to be made
9. Effort estimation
 - a. Learning techniques for better estimating the effort required to complete a task
 - b. Practicing estimation and comparing to how long tasks actually take
 - c. Effort estimation is quite difficult, so students are likely to be way off in many cases, but seeing the process play out with their own work is valuable

Illustrative Learning Outcomes:

KA Core:

1. Compare different methods of eliciting requirements along multiple axes.
2. Identify differences between two methods of describing functional requirements (e.g., customer interviews, user studies, etc.) and the situations where each would be preferred.
3. Identify which behaviors are required, allowed, or barred from a given set of requirements and a list of candidate behaviors.
4. Collect a set of requirements for a simple software system.
5. Identify areas of a software system that need to be changed, given a description of the system and a set of new requirements to be implemented.
6. Identify the functional and non-functional requirements in a set of requirements.

Non-core:

7. Create a prototype of a software system to validate a set of requirements. (Building a mock-up, MVP, etc.)
8. Estimate the time to complete a set of tasks, then compare estimates to the actual time taken.
9. Determine an implementation sequence for a set of tasks, adhering to dependencies between them, with a goal to retire risk as early as possible.
10. Write a requirement specification for a simple software system.

SE-Design: Software Design

CS Core:

1. System design principles (See also: [SF-Reliability](#))
 - a. Levels of abstraction (e.g., architectural design and detailed design)
 - b. Separation of concerns
 - c. Information hiding
 - d. Coupling and cohesion
2. Software architecture (See also: [SF-Reliability](#))
 - a. Design paradigms
 - i. Top-down functional decomposition / layered design
 - ii. Data-oriented architecture
 - iii. Object-oriented analysis and design
 - iv. Event-driven design
 - b. Standard architectures (e.g., client-server and microservice architectures including REST discussions, n-layer, pipes-and-filters, Model View Controller)

- c. Identifying component boundaries and dependencies
- 3. Programming in the large vs programming in the small (See also: [SF-Reliability](#))
- 4. Code smells and other indications of code quality, distinct from correctness.(See also: [SEC-Engineering](#))

KA Core:

- 5. API design principles
 - a. Consistency
 - i. Consistent APIs are easier to learn and less error-prone
 - ii. Consistency is both internal (between different portions of the API) and external (following common API patterns)
 - b. Composability
 - c. Documenting contracts
 - i. API operations should describe their effect on the system, but not generally their implementation
 - ii. Preconditions, postconditions, and invariants
 - d. Expandability
 - e. Error reporting
 - i. Errors should be clear, predictable, and actionable
 - ii. Input that does not match the contract should produce an error
 - iii. Errors that can be reliably managed without reporting should be managed
- 6. Identifying and codifying data invariants and time invariants
- 7. Structural and behavioral models of software designs
- 8. Data design (See also: [DM-Modeling](#))
 - a. Data structures
 - b. Storage systems
- 9. Requirement traceability
 - a. Understanding which requirements are satisfied by a design

Non-Core:

- 10. Design modeling, for instance with class diagrams, entity relationship diagrams, or sequence diagrams
- 11. Measurement and analysis of design quality
- 12. Principles of secure design and coding (See also: [SEC-Engineering](#))
 - a. Principle of least privilege
 - b. Principle of fail-safe defaults
 - c. Principle of psychological acceptability
- 13. Evaluating design tradeoffs (e.g., efficiency vs reliability, security vs usability)

Illustrative Learning Outcomes:

CS Core:

- 1. Identify the standard software architecture of a given high-level design.
- 2. Select and use an appropriate design paradigm to design a simple software system and explain how system design principles have been applied in this design.

3. Adapt a flawed system design to better follow principles such as separation of concerns or information hiding.
4. Identify the dependencies among a set of software components in an architectural design.

KA Core:

5. Design an API for a single component of a large software system, including identifying and documenting each operation's invariants, contract, and error conditions.
6. Evaluate an API description in terms of consistency, composability, and expandability.
7. Expand an existing design to include a new piece of functionality.
8. Design a set of data structures to implement a provided API surface.
9. Identify which requirements are satisfied by a provided software design.

Non-Core:

10. Translate a natural language software design into class diagrams.
11. Adapt a flawed system design to better follow the principles of least privilege and fail-safe defaults.
12. Contrast two software designs across different qualities, such as efficiency or usability.

SE-Construction: Software Construction

CS Core:

1. Practical small-scale testing (See also: [SDF-Practices](#))
 - a. Unit testing
 - b. Test-driven development - This is particularly valuable for students psychologically, as it is far easier to engage constructively with the challenge of identifying challenging inputs for a given API (edge cases, corner cases) a priori. If they implement first, the instinct is often to avoid trying to crash their new creation, while a test-first approach gives them the intellectual satisfaction of spotting the problem cases and then watching as more tests pass during the development process.
2. Documentation (See also: [SDF-Practices](#))
 - a. Interface documentation - Describe interface requirements, potentially including (formal or informal) contracts, pre and post conditions, invariants.
 - b. Implementation documentation should focus on tricky and non-obvious pieces of code, whether because the code is using advanced language features or the behavior of the code is complex. (Do not add comments that re-state common/obvious operations and simple language features.)
 - i. Clarify dataflow, computation, etc., focusing on what the code is
 - ii. Identify subtle/tricky pieces of code and refactor to be self-explanatory if possible, or provide appropriate comments to clarify.

KA Core:

3. Coding style (See also: [SDF-Practices](#))
 - a. Style guides
 - b. Commenting
 - c. Naming
4. "Best Practices" for coding: techniques, idioms/patterns, mechanisms for building quality programs (See also: [SEC-Coding](#), [SDF-Practices](#))

- a. Defensive coding practices
 - b. Secure coding practices and principles
 - c. Using exception handling mechanisms to make programs more robust, fault-tolerant
5. Debugging (See also: [SDF-Practices](#))
 6. Logging
 7. Use of libraries and frameworks developed by others (See also: [SDF-Practices](#))

Non-Core:

8. Larger-scale testing
 - a. Test doubles (stubs, mocks, fakes)
 - b. Dependency injection
9. Work sequencing, including dependency identification, milestones, and risk retirement
 - a. Dependency identification: Identifying the dependencies between different tasks
 - b. Milestones: A collection of tasks that serve as a marker of progress when completed. Ideally, the milestone encompasses a useful unit of functionality.
 - c. Risk retirement: Identifying what elements of a project are risky and prioritizing completing tasks that address those risks
10. Potential security problems in programs (See also: [SEC-Coding](#))
 - a. Buffer and other types of overflows
 - b. Race conditions
 - c. Improper initialization, including choice of privileges
 - d. Input validation
11. Documentation (autogenerated)
12. Development context: “green field” vs existing code base
 - a. Change impact analysis
 - b. Change actualization
13. Release management
14. DevOps practices

Illustrative Learning Outcomes:

CS Core:

1. Write appropriate unit tests for a small component (several functions, a single type, etc.).
2. Write appropriate interface and (if needed) implementation comments for a small component.

KA Core:

3. Describe techniques, coding idioms and mechanisms for implementing designs to achieve desired properties such as reliability, efficiency, and robustness.
4. Write robust code using exception handling mechanisms.
5. Describe secure coding and defensive coding practices.
6. Select and use a defined coding standard in a small software project.
7. Compare and contrast integration strategies including top-down, bottom-up, and sandwich integration.
8. Describe the process of analyzing and implementing changes to code base developed for a specific project.

9. Describe the process of analyzing and implementing changes to a large existing code base.

Non-Core:

10. Rewrite a simple program to remove common vulnerabilities, such as buffer overflows, integer overflows and race conditions.
11. Write a software component that performs some non-trivial task and is resilient to input and run-time errors.

SE-Validation: Software Verification and Validation

CS Core:

1. Verification and validation concepts
 - a. Verification: Are we building the thing right?
 - b. Validation: Did we build the right thing?
2. Why testing matters
 - a. Does the component remain functional as the code evolves?
3. Testing objectives
 - a. Usability
 - b. Reliability
 - c. Conformance to specification
 - d. Performance
 - e. Security
4. Test kinds
 - a. Unit
 - b. Integration
 - c. Validation
 - d. System
5. Stylistic differences between tests and production code
 - a. DAMP vs DRY - more duplication is warranted in test code.

KA Core:

6. Test planning and generation
 - a. Test case generation, from formal models, specifications, etc.
 - b. Test coverage
 - i. Test matrices
 - ii. Code coverage (how much of the code is tested)
 - iii. Environment coverage (how many hardware architectures, OSes, browsers, etc. are tested)
 - c. Test data and inputs
7. Test development
 - a. Test-driven development
 - b. Object oriented testing, mocking, and dependency injection
 - c. Black-box and white-box testing techniques
 - d. Test tooling, including code coverage, static analysis, and fuzzing
8. Verification and validation in the development cycle

- a. Code reviews
 - b. Test automation, including automation of tooling
 - c. Pre-commit and post-commit testing
 - d. Trade-offs between test coverage and throughput/latency of testing
 - e. Defect tracking and prioritization
 - i. Reproducibility of reported defects
9. Domain specific verification and validation challenges
- a. Performance testing and benchmarking
 - b. Asynchrony, parallelism, and concurrency
 - c. Safety-critical
 - d. Numeric

Non-Core:

10. Verification and validation tooling and automation
- a. Static analysis
 - b. Code coverage
 - c. Fuzzing
 - d. Dynamic analysis and fault containment (sanitizers, etc.)
 - e. Fault logging and fault tracking
11. Test planning and generation
- a. Fault estimation and testing termination including defect seeding
 - b. Use of random and pseudo random numbers in testing
12. Performance testing and benchmarking
- a. Throughput and latency
 - b. Degradation under load (stress testing, FIFO vs LIFO handling of requests)
 - c. Speedup and scaling
 - i. Amadhl's law
 - ii. Gustafson's law
 - iii. Soft and weak scaling
 - d. Identifying and measuring figures of merits
 - e. Common performance bottlenecks
 - i. Compute-bound
 - ii. Memory-bandwidth bound
 - iii. Latency-bound
 - f. Statistical methods and best practices for benchmarking
 - i. Estimation of uncertainty
 - ii. Confidence intervals
 - g. Analysis and presentation (graphs, etc.)
 - h. Timing techniques
13. Testing asynchronous, parallel, and concurrent systems
14. Verification and validation of non-code artifacts (documentation, training materials)

Illustrative Learning Outcomes:

CS Core:

1. Explain why testing is important.

2. Distinguish between program validation and verification.
3. Describe different objectives of testing.
4. Compare and contrast the different types and levels of testing (regression, unit, integration, systems, and acceptance).

KA Core:

5. Describe techniques for creating a test plan and generating test cases.
6. Create a test plan for a medium-size code segment which includes a test matrix and generation of test data and inputs.
7. Implement a test plan for a medium-size code segment.
8. Identify the fundamental principles of test-driven development methods and explain the role of automated testing in these methods.
9. Discuss issues involving the testing of object-oriented software.
10. Describe mocking and dependency injection and their application.
11. Undertake, as part of a team activity, a code review of a medium-size code segment.
12. Describe the role that tools can play in the validation of software.
13. Automate testing in a small software project.
14. Explain the roles, pros, and cons of pre-commit and post-commit testing.
15. Discuss the tradeoffs between test coverage and test throughput/latency and how this can impact verification.
16. Use a defect tracking tool to manage software defects in a small software project.
17. Discuss the limitations of testing in certain domains.

Non-Core:

18. Describe and compare different tools for verification and validation.
19. Automate the use of different tools in a small software project.
20. Explain how and when random numbers should be used in testing.
21. Describe approaches for fault estimation.
22. Estimate the number of faults in a small software application based on fault density and fault seeding.
23. Describe throughput and latency and provide examples of each.
24. Explain speedup and the different forms of scaling and how they are computed.
25. Describe common performance bottlenecks.
26. Describe statistical methods and best practices for benchmarking software.
27. Explain techniques for and challenges with measuring time when constructing a benchmark.
28. Identify the figures of merit, construct and run a benchmark, and statistically analyze and visualize the results for a small software project.
29. Describe techniques and issues with testing asynchronous, concurrent, and parallel software.
30. Create a test plan for a medium-size code segment which contains asynchronous, concurrent, and/or parallel code, including a test matrix and generation of test data and inputs.
31. Describe techniques for the verification and validation of non-code artifacts.

SE-Refactoring: Refactoring and Code Evolution

KA Core:

1. Hyrum's Law / The Law of Implicit Interfaces
2. Backward compatibility
 - a. Compatibility is not a property of a single entity, it's a property of a *relationship*.
 - b. Backward compatibility needs to be evaluated in terms of provider + consumer(s) or with a well-specified model of what forms of compatibility a provider aspires to / promises.
3. Refactoring
 - a. Standard refactoring patterns (rename, inline, outline, etc.)
 - b. Use of refactoring tools in IDE
 - c. Application of static-analysis tools (to identify code in need of refactoring, generate changes, etc.)
 - d. Value of refactoring as a remedy for technical debt
4. Versioning
 - a. Semantic Versioning (SemVer)
 - b. Trunk-based development

Non-Core:

5. "Large Scale" Refactoring - techniques when a refactoring change is too large to commit safely (large projects), or when it is impossible to synchronize change between provider + all consumers (multiple repositories, consumers with private code).
 - a. Express both old and new APIs so that they can co-exist
 - b. Minimize the size of *behavior* changes
 - c. Why these techniques are required, (e.g., "API consumers I can see" vs "consumers I can't see")

Illustrative Learning Outcomes:

KA-Core:

1. Identify both explicit and implicit behavior of an interface, and identify potential risks from Hyrum's Law
2. Consider inputs from static analysis tools and/or Software Design principles to identify code in need of refactoring.
3. Identify changes that can be broadly considered "backward compatible," potentially with explicit statements about what usage is or is not supported
4. Refactor the implementation of an interface to improve design, clarity, etc. with minimal/zero impact on existing users
5. Evaluate whether a proposed change is sufficiently safe given the versioning methodology in use for a given project

Non-Core:

6. Plan a complex multi-step refactoring to change default behavior of an API safely.

SE-Reliability: Software Reliability

KA Core:

1. Concept of reliability as probability of failure or mean time between failures, and faults as cause of failures
2. Identifying reliability requirements for different kinds of software
3. Software failures caused by defects/bugs, and so for high reliability the goal is to have minimum defects - by injecting fewer defects (better training, education, planning), and by removing most of the injected defects (testing, code review, etc.)
4. Software reliability, system reliability and failure behavior
5. Defect injection and removal cycle, and different approaches for defect removal
6. Compare the “error budget” approach to reliability with the “error-free” approach, and identify domains where each is relevant

Non-Core:

7. Software reliability models
8. Software fault tolerance techniques and models
 - a. Contextual differences in fault tolerance (e.g., crashing a flight critical system is strongly avoided, crashing a data processing system before corrupt data is written to storage is highly valuable)
9. Software reliability engineering practices - including reviews, testing, practical model checking
10. Identification of dependent and independent failure domains, and their impact on system reliability
11. Measurement-based analysis of software reliability - telemetry, monitoring and alerting, dashboards, release qualification metrics, etc.

Illustrative Learning Outcomes:

KA Core:

1. Describe how to determine the level of reliability required by a software system.
2. Explain the problems that exist in achieving very high levels of reliability.
3. Understand approaches to minimizing faults that can be applied at each stage of the software lifecycle.

Non-Core:

4. Demonstrate the ability to apply multiple methods to develop reliability estimates for a software system.
5. Identify methods that will lead to the realization of a software architecture that achieves a specified level of reliability.
6. Identify ways to apply redundancy to achieve fault tolerance.
7. Identify single-point-of-failure (SPF) dependencies in a system design.

SE-Formal: Formal Methods

Non-Core:

1. Formal specification of interfaces
 - a. Specification of pre- and post- conditions
 - b. Formal languages for writing and analyzing pre- and post-conditions.
2. Problem areas well served by formal methods

- a. Lock-free programming, data races
- b. Asynchronous and distributed systems, deadlock, livelock, etc.
3. Comparison to other tools and techniques for defect detection
 - a. Testing
 - b. Fuzzing
4. Formal approaches to software modeling and analysis
 - a. Model checkers
 - b. Model finders

Illustrative Learning Outcomes:

1. Describe the role formal specification and analysis techniques can play in the development of complex software and compare their use as validation and verification techniques with testing.
2. Apply formal specification and analysis techniques to software designs and programs with low complexity.
3. Explain the potential benefits and drawbacks of using formal specification languages.

Professional Dispositions

- **Collaborative:** Software engineering is increasingly described as a “team sport” - successful software engineers are able to work with others effectively. Humility, respect, and trust underpin the collaborative relationships that are essential to success in this field.
- **Professional:** Software engineering produces technology that has the chance to influence literally billions of people. Awareness of our role in society, strong ethical behavior, and commitment to respectful day-to-day behavior outside of one’s team are essential.
- **Communicative:** No single software engineer on a project is likely to know all of the project details. Successful software projects depend on engineers communicating clearly and regularly in order to coordinate effectively.
- **Meticulous:** Software engineering requires attention to detail and consistent behavior from everyone on the team. Success in this field is clearly influenced by a meticulous approach - comprehensive understanding, proper procedures, and a solid avoidance of cutting corners.
- **Responsible:** The collaborative aspects of software engineering also highlight the value of being responsible. Failing to take responsibility, failing to follow through, and failing to keep others informed are all classic causes of team friction and bad project outcomes.

Mathematics Requirements

Desirable:

- Introductory statistics (performance comparisons, evaluating experiments, interpreting survey results, etc.). (See also CS-Core requirements for [MSF-Statistics](#))

Course Packaging Suggestions

Advanced Course to include at least the following:

- [SE-Teamwork](#) (4 hours)
- [SE-Tools](#) (4 hours)
- [SE-Requirements](#) (2 hours)
- SE-Design (5 hours)
- [SE-Construction](#) (4 hours)
- [SE-Validation](#) (4 hours)
- [SE-Refactoring](#) (2 hours)
- [SE-Reliability](#) (2 hours)
- [SEP-Professional-Ethics](#) (7 hours)

Prerequisites:

- [SDF-Fundamentals](#)

Course objectives: Students should be able to perform good quality code review for colleagues (especially focusing on professional communication and teamwork needs), read and write unit tests, use basic software tools (IDEs, version control, static analysis tools) and perform basic activities expected of a new hire on a software team.

Committee

Chair: Titus Winters, Google, New York City, NY, USA

Members:

- Brett A. Becker, University College Dublin, Dublin, Ireland
- Adam Vartanian, Cord, London, UK
- Bryce Adelstein Lelbach, NVIDIA, New York City, NY, USA
- Patrick Servello, CIWRO, Norman, OK, USA
- Pankaj Jalote, IIIT-Delhi, Delhi, India
- Christian Servin, El Paso Community College, El Paso, TX, USA

Contributors:

- Hyrum Wright, Google, Pittsburgh, PA, USA
- Olivier Giroux, Apple, Cupertino, CA, USA
- Gennadiy Civil, Google, New York City, NY, USA

Security (SEC)

Preamble

Computing supports nearly every facet of modern critical infrastructure: transportation, communication, healthcare, education, energy generation and distribution, to name a few. With rampant attacks on and breaches of this infrastructure, computer science graduates have an important role in designing, implementing, and operating software systems that are robust, safe, and secure.

The Security (SEC) knowledge area (KA) focuses on developing a *security mindset* into the overall ethos of computer science graduates so that security is embedded in all of their work products. Computer science students need to learn about system vulnerabilities and understand threats against computer systems. The *Security* title choice was intentional to serve as a one-word umbrella term for this knowledge area, which also includes concepts to support privacy, cryptography, secure systems, secure data, and secure code.

The SEC KA relies on shared concepts pervasive in all the other areas of CS2023. It identifies seven crosscutting concepts of cybersecurity: *confidentiality*, *integrity*, *availability*, *risk assessment*, *systems thinking*, *adversarial thinking*, and *human-centered thinking*. The seventh concept, *human-centered thinking*, is additional to the six crosscutting concepts originally defined in the Cybersecurity Curricula 2017 (CSEC2017) [26]. This addition reinforces to students that humans are also a link in the overall chain of security, a theme that is also covered in KAs, such as [HCI](#). Principles of protecting systems (also in the [DM](#), [OS](#), [SDF](#), [SE](#) and [SF](#) knowledge areas) include security-by-design, privacy-by-design, defense-in-depth, and zero-trust.

Another concept is the notion of assurance, which is an attestation that security mechanisms need to comply with the security policies that have been defined for data, processes, and systems. Assurance is tied in with the concepts of verification and validation in the [SE](#) knowledge area. Considerations of data privacy and security are shared with the [DM](#) (technical aspects) and [SEP](#) KAs.

The SEC KA thus sits atop several of the other CS2023 KAs, while including additional concepts not present in those KAs. The specific dependence on other KAs is stated below, starting with the Core Hours table. CS2023 treats security as a crucial component of the skillset of any CS graduate, and the hours needed for security preparation come from all of the other 16 CS2023 KAs.

Changes since CS2013

The Security KA is an updated name for CS2013's Information Assurance and Security (IAS) KA. Since 2013, Information Assurance and Security has been rebranded as Cybersecurity, which has become a new computing discipline, with its own curricular guidelines (CSEC 2017) developed by a Joint Task Force of the ACM, IEEE Computer Society, AIS and IFIP in 2017.

Moreover, since 2013, other curricular recommendations for cybersecurity beyond CS2013 and CSEC 2017 have been made. In the US, the National Security Agency recognizes institutions as Centers of Academic Excellence (CAE) in Cyber Defense and/or Cyber Operations if their cybersecurity programs

meet the respective CAE curriculum requirements. Additionally, the National Initiative for Cybersecurity Education (NICE) of the US National Institute for Standards and Technologies (NIST) has developed and revised the Workforce Framework for Cybersecurity (NICE Workforce Framework), which identifies competencies (knowledge and skills) needed to perform tasks relevant to cybersecurity work roles. The European Cybersecurity Skills Framework (ECSF) includes a standard ontology to describe cybersecurity tasks and roles, as well as addressing the cybersecurity personnel shortage in EU member countries. Similarities and differences of these cybersecurity guidelines, viewed from the CS perspective, also informed the SEC knowledge area.

Building on CS2013's recognition of the pervasiveness of security in computer science, the CS2023 SEC knowledge area focuses on ensuring that students develop a *security mindset* so that they are prepared for the continual changes occurring in computing. One useful addition is the knowledge unit for security analysis, design, and engineering to support the concepts of security-by-design and privacy-by-design.

The importance of computer science in ensuring the protection of future computing systems and societal critical infrastructure will continue to grow. Consequently, it is imperative that faculty teaching computer science incorporate the latest advances in security and privacy approaches to keep their curriculum current.

Differences between CS2023 Security KA and Cybersecurity

CS2023's SEC knowledge area focuses on those aspects of security, privacy, and related concepts important for computer science students. In comparison, CSEC 2017 characterizes similarities and differences in the cybersecurity book of knowledge using the disciplinary lenses of computer science, computer engineering, software engineering, information systems, information technology, and other disciplines. In short, the major goal of the SEC knowledge area is to ensure computer science graduates are able to design and develop more secure code, ensure data security and privacy, and can apply a security mindset to their daily activities.

Protecting what happens *within* the perimeter of a networked computer system is a core competency of computer science graduates. Although the computer science and cybersecurity knowledge units have overlaps, the demands upon cybersecurity graduates typically are to protect the perimeter. CSEC 2017 defines cybersecurity as a highly interdisciplinary field of study that covers eight areas (data, software, component, connection, system, human, organizational, and societal security) and prepares its students for both technical and managerial work roles in cybersecurity.

The first five CSEC 2017 areas are technical and have overlaps with the CS2023 SEC KA, but the intent of coverage is substantively different as CS students bring to bear the core competencies described in all the 17 CS2023 knowledge areas. For instance, consider the SEC KA's Secure Coding knowledge unit. The CS student will need to view this knowledge unit from a computer science lens, as an extension of the material covered in the [SDF](#), [SE](#) and [PDC](#) KAs, while the Cybersecurity student will need to view software security in the overall context of diverse cybersecurity goals. These viewpoints are not totally distinct and have overlaps, but the lenses used to examine and present the content are

different. There are similar commonalities and differences among CS2023 SEC knowledge units and corresponding CSEC 2017 knowledge units.

Core Hours

Knowledge Units	CS Core	KA Core
Foundational Security	1 + 7 (DM , FPL , PDC , SDF , SE , OS)	7
Society, Ethics, and the Profession	1 + 4 (SEP)	2
Secure Coding	2 + 6 (FPL, SDF , SE)	5
Cryptography	1 + 8 (MSF)	4
Security Analysis, Design, and Engineering	1 + 4 (MSF , SE)	8
Digital Forensics	0	6
Security Governance	0	3
Total hours	6	35

Note that the Foundational Security knowledge unit here adds only 1 CS Core hour as it relies on CS Core hours from [SDF](#), [SE](#), and [OS](#). Similarly, the Society, Ethics, and the Profession knowledge unit, the Cryptography knowledge unit, Security Analysis, Design, and Engineering knowledge unit each add 1 CS Core hour while the Secure Coding knowledge unit adds 2 CS Core hours. In total, the SEC knowledge area adds only 6 CS Core hours to the curriculum.

The SEC knowledge area also includes several KA Core hours from the other knowledge areas, as noted in the many “see also” references in each knowledge unit below. If one looks at the shared hours in the core and reduces duplicative hours, the result is that approximately 28 hours of CS Core hours from the other knowledge areas are needed, either to provide the basis for the SEC knowledge area or to complement its content that is shown below. Of these, [MSF-Discrete](#), [MSF-Probability](#), and [MSF-Statistics](#) are likely to be relied upon extensively in all the SEC knowledge units, as are [SDF-Fundamentals](#), [SDF-Algorithms](#), and [SDF-Practices](#). The others are mentioned within each of the SEC knowledge units described below. As the same content of the different knowledge units might form the basis, the CS Core hours from the other knowledge areas overlap and the total of 28 hours eliminates obvious duplication.

Knowledge Units

SEC-Foundations: Foundational Security

CS Core:

1. Developing a security mindset incorporating crosscutting concepts: confidentiality, integrity, availability, risk assessment, systems thinking, adversarial thinking, human-centered thinking
2. Basic concepts of authentication and authorization/access control
3. Vulnerabilities, threats, attack surfaces, and attack vectors (See also: [OS-Protection](#))
4. Denial of Service (DoS) and Distributed Denial of Service (DDoS) (See also: [OS-Protection](#))
5. Principles and practices of protection, e.g., least privilege, open design, fail-safe defaults, defense in depth, and zero trust; and how they can be implemented (See also: [OS-Principles](#), [OS-Protection](#), [SE-Construction](#), [SEP-Security](#))
6. Optimization considerations between security, privacy, performance, and other design goals (See also: [SDF-Practices](#), [SE-Validation](#), [HCI-Design](#))
7. Impact of AI on security and privacy: using AI to bolster defenses as well as address increased adversarial capabilities due to AI (See also: [AI-SEP](#), [HCI-Design](#), [HCI-SEP](#))

KA Core:

8. Access control models (e.g., discretionary, mandatory, role-based, and attribute-based)
9. Security controls
10. Concepts of trust and trustworthiness
11. Applications of a security mindset: web, cloud, and mobile devices (See also: [SF-System Design](#), [SPD-Common](#))
12. Protecting embedded and cyber-physical systems (See also: [SPD-Embedded](#))
13. Principles of usable security and human-centered computing (See also: [HCI-Design](#), [SEP-Security](#))
14. Security and trust in AI/machine learning systems, e.g., fit for purpose, ethical operating boundaries, authoritative knowledge sources, verified training data, repeatable system evaluation tests, system attestation, independent validation/certification; unintended consequences from: adverse effect (See also: [AI-Introduction](#), [AI-ML](#), [AI-SEP](#), [SEP-Security](#))
15. Security risks in building and operating AI/machine learning systems. e.g., algorithm bias, knowledge corpus bias, training corpus bias, copyright violation (See also: [AI-Introduction](#), [AI-ML](#), [AI-SEP](#))
16. Hardware considerations in security, e.g., principles of secure hardware, secure processor architectures, cryptographic acceleration, compartmentalization, software-hardware interaction (See also: [AR-Assembly](#), [AR-Representation](#), [OS-Purpose](#))

Illustrative Learning Outcomes:

CS Core:

1. Evaluate a system for possible attacks that can be launched by an adversary
2. Design and develop approaches to protect a system from a set of identified threats

KA Core:

3. Describe how harm to user privacy can be avoided
4. Develop a system that incorporates various principles of security and privacy
5. Compare the different access control models in terms of functionality and performance
6. Show how an adversary could use machine learning algorithms to reduce the security of a system
7. Show how a developer could improve the security of a system using machine learning algorithms
8. Describe hardware (especially CPU) vulnerabilities that can impact software.

SEC-SEP: Society, Ethics, and the Profession

CS Core:

1. Principles and practices of privacy (See also: [SEP-Security](#))
2. Societal impacts on breakdowns in security and privacy (See also: [SEP-Context](#), [SEP-Privacy](#), [SEP-Security](#))
3. Applicability of laws and regulations on security and privacy (See also: [SEP-Security](#))
4. Professional ethical considerations when designing secure systems and maintaining privacy; ethical hacking (See also: [SEP-Professional-Ethics](#), [SEP-Privacy](#), [SEP-Security](#))

KA-Core:

5. Security by design (See also: [SF-Security](#), [SF-Design](#))
6. Privacy by design and privacy engineering (See also: [SEP-Privacy](#), [SEP-Security](#))
7. Security and privacy implications of malicious AI/machine learning actors, e.g., identifying deep fakes (See also: [AI-Introduction](#), [AI-ML](#), [SEP-Privacy](#), [SEP-Security](#))
8. Societal impacts of Internet of Things (IoT) devices and other emerging technologies on security and privacy (See also: [SEP-Privacy](#), [SEP-Security](#))

Illustrative Learning Outcomes:

CS Core:

1. Calculate the impact of a breakdown in security of a given system
2. Construct a system that conforms to security laws.
3. Apply a set of privacy regulations to design a system that protects privacy

KA Core:

4. Evaluate the legal ramifications of a system not corresponding to applicable laws and regulations
5. Construct a system that is designed to avoid harm to user privacy

SEC-Coding: Secure Coding

CS Core:

1. Common vulnerabilities and weaknesses
2. SQL injection and other injection attacks
3. Cross-site scripting techniques and mitigations
4. Input validation and data sanitization (See also: [OS-Protection](#), [SDF-Fundamentals](#), [SE-Validation](#))
5. Type safety and type-safe languages (See also: [FPL-Types](#), [FPL-Systems](#), [OS-Protection](#), [SDF-Fundamentals](#), [SE-Validation](#))
6. Buffer overflows, stack smashing, and integer overflows (See also: [AR-Assembly](#), [FPL-Systems](#), [OS-Protection](#))
7. Security issues due to race conditions (See also: [FPL-Parallel](#), [PDC-Evaluation](#))

KA Core:

8. Principles of noninterference and nondeducibility
9. Preventing information flow attacks
10. Offensive security techniques as a defense
11. AI-assisted malware detection techniques
12. Ransomware: creation, prevention, and mitigation
13. Secure use of third-party components (See also: [SE-Construction](#), [SE-Validation](#))

14. Malware: varieties, creation, reverse engineering, and defense against them (See also: [FPL-Systems](#), [FPL-Translation](#))
15. Assurance: testing (including fuzzing and penetration testing), verification, and validation (See also: [OS-Protection](#), [SDF-Fundamentals](#), [SE-Construction](#), [SE-Validation](#))
16. Static and dynamic analyses (See also: [FPL-Analysis](#), [MSF-Protection](#), [PDC-Evaluation](#), [SE-Validation](#))
17. Secure compilers and secure code generation (See also: [FPL-Runtime](#), [FPL-Translation](#))

Illustrative Learning Outcomes:

CS Core:

1. Identify underlying problems in given examples of an enumeration of common weaknesses, and explain how they can be circumvented
2. Apply input validation and data sanitization techniques to enhance security of a program
3. Describe how the selection of a programming language can impact the security of the system being constructed
4. Rewrite a program in a type-safe language (e.g., Java or Rust) originally written in an unsafe programming language (e.g., C/C++)
5. Evaluate a program for possible buffer overflow attacks and rewrite to prevent such attacks
6. Evaluate a set of related programs for possible race conditions and prevent an adversary from exploiting them
7. Evaluate and prevent SQL injections attacks on a database application
8. Evaluate and prevent cross-site scripting attacks against a website

KA Core:

9. Describe different kinds of malicious software
10. Construct a program that tests for all input handling errors
11. Explain the risks with misusing interfaces with third-party code and how to correctly use third-party code
12. Discuss the need to update software to fix security vulnerabilities and the lifecycle management of the fix
13. Construct a system that is protected from unauthorized information flows
14. Apply static and dynamic tools to identify programming faults
15. Evaluate a system for existence of malware and remove it
16. Implement preventive techniques to reduce the occurrence of ransomware

SEC-Crypto: Cryptography

CS Core:

1. Differences between algorithmic, applied, and mathematical views of cryptography.
2. Mathematical preliminaries: modular arithmetic, Euclidean algorithm, probabilistic independence, linear algebra basics, number theory, finite fields, complexity, asymptotic analysis (See also: [MSF-Discrete](#), [MSF-Linear](#))
3. Basic cryptography: symmetric key and public key cryptography (See also: [AL-Foundational](#), [MSF-Discrete](#))
4. Basic cryptographic building blocks, including symmetric encryption, asymmetric encryption, hashing, and message authentication (See also: [MSF-Discrete](#))

5. Classical cryptosystems, such as shift, substitution, transposition ciphers, code books, and machines (See also: [MSF-Discrete](#))
6. Kerckhoff's principle and use of vetted libraries (See also: [SE-Construction](#))
7. Usage of cryptography in real-world applications, e.g., electronic cash, secure channels between clients and servers, secure electronic mail, entity authentication, device pairing, steganography, and voting systems (See also: [NC-Security](#), [GIT-Image](#))

KA Core:

8. Additional mathematics: primality, factoring and elliptic curve cryptography (See also: [MSF-Discrete](#))
9. Private-key cryptosystems: substitution-permutation networks, linear cryptanalysis, differential cryptanalysis, DES, and AES (See also: [MSF-Discrete](#), [NC-Security](#))
10. Public-key cryptosystems: Diffie-Hellman and RSA (See also: [MSF-Discrete](#))
11. Data integrity and authentication: hashing, and digital signatures (See also: [MSF-Discrete](#), [DM-Security](#))
12. Cryptographic protocols: challenge-response authentication, zero-knowledge protocols, commitment, oblivious transfer, secure two- or multi-party computation, hash functions, secret sharing, and applications (See also: [MSF-Discrete](#))
13. Attacker capabilities: chosen-message attack (for signatures), birthday attacks, side channel attacks, and fault injection attacks (See also: [NC-Security](#))
14. Quantum cryptography; Post Quantum/Quantum resistant cryptography (See also: [AL-Foundational](#), [MSF-Discrete](#))
15. Blockchain and cryptocurrencies (See also: [MSF-Discrete](#), [PDF-Communication](#))

Illustrative Learning Outcomes:

CS Core:

1. Explain the role of cryptography in supporting security and privacy
2. Discuss the risks of inventing one's own cryptographic methods
3. Discuss the importance of prime numbers in cryptography and explain their use in cryptographic algorithms
4. Implement and cryptanalyze classical ciphers

KA Core:

5. Describe how crypto keys can be managed securely
6. Compare the space and time performance of a given set of cryptographic methods
7. Discuss how modern private-key cryptosystems work and ways to cryptanalyze them
8. Discuss how modern public-key cryptosystems work and ways to cryptanalyze them
9. Compare different cryptographic algorithms in terms of security
10. Explain key exchange protocols and show approaches to reduce their failure
11. Describe real-world applications of cryptographic primitives and protocols
12. Discuss how quantum cryptography works and the impact of quantum computing on cryptographic algorithms

SEC-Engineering: Security Analysis, Design, and Engineering

CS Core:

1. Security engineering goals: building systems that remain dependable despite errors, accidents, or malicious adversaries (See also: [SE-Construction](#), [SE-Validation](#), [SEP-Security](#))
2. Privacy engineering goals: building systems that design, implement, and deploy privacy features and controls (See also: [SEP-Privacy](#))
3. Problem analysis and situational analysis to address system security ([SE-Validation](#))
4. Engineering tradeoff analysis based on time, cost, risk tolerance, risk acceptance, return on investment, and so on (See also: [PDC-Evaluation](#), [SE-Validation](#))

KA Core:

5. Security design and engineering, including functional requirements, secure subsystems, information protection, security testing, security assessment, and evaluation (See also: [PDC-Evaluation](#), [SE-Requirements](#), [SE-Validation](#))
6. Security analysis, covering security requirements analysis; security controls analysis; threat analysis; and vulnerability analysis (See also: [FPL-Analysis](#), [PDC-Evaluation](#))
7. Security attack domains and attack surfaces, e.g., communications and networking, hardware, physical, social engineering, software, and supply chain (See also: [NC-Security](#))
8. Security attack modes, techniques, and tactics, e.g., authentication abuse; brute force; buffer manipulation; code injection; content insertion; denial of service; eavesdropping; function bypass; impersonation; integrity attack; interception; phishing; protocol analysis; privilege abuse; spoofing; and traffic injection (See also: [NC-Security](#), [OS-Protection](#), [SE-Validation](#))
9. Attestation of software products with respect to their specification and adaptiveness (See also: [SE-Requirements](#), [SE-Validation](#))
10. Design and development of cyber-physical systems
11. Considerations for trustworthy computing, e.g., tamper resistant packaging, trusted boot, trusted kernel, hardware root of trust, software signing and verification, hardware-based cryptography, virtualization, and containers (See also: [SE-Construction](#), [SE-Validation](#))

Illustrative Learning Outcomes:

CS Core:

1. Create a threat model for a system or system design
2. Apply situational analysis to develop secure solutions under a specified scenario
3. Evaluate a given scenario for tradeoff analysis for system performance, risk assessment, and costs

KA Core:

4. Design a set of technical security controls, countermeasures, and information protections to meet the security requirements and security objectives for a system
5. Evaluate the effectiveness of security functions, technical controls, and componentry for a system.
6. Identify and mitigate security vulnerabilities and weaknesses in a system
7. Evaluate and predict emergent behavior in areas such as Data Science, AI, and Machine Learning

SEC-Forensics: Digital Forensics

KA Core:

1. Basic principles and methodologies for digital forensics
2. System design for forensics

3. Forensics in different situations: operating systems, file systems, application forensics, web forensics, network forensics, mobile device forensics, use of database auditing (See also: [NC-Security](#))
4. Attacks on forensics and preventing such attacks
5. Incident handling processes
6. Rules of evidence – general concepts and differences between jurisdictions (See also: [SEP-Security](#))
7. Legal issues: digital evidence protection and management, chains of custody, reporting, serving as an expert witness (See also: [SEP-Security](#))

Illustrative Learning Outcomes:

KA Core:

1. Explain what a digital investigation is and how it can be implemented (See also: [SEP-Security](#))
2. Design and implement software to support forensics
3. Describe legal requirements for using seized data and its usage (See also: [SEP-Security](#))
4. Describe and implement an end-to-end chain of custody from initial digital evidence seizure to evidence disposal (See also: [SEP-Privacy](#); [SEP-Security](#))
5. Extract data from a hard drive to comply with the law (See also: [SEP-Security](#))
6. Discuss a person's professional responsibilities and liabilities when testifying as a forensics expert (See also: [SEP-Professional-Ethics](#))
7. Recover data based on a given search term from an imaged system
8. Reconstruct data and events from an application history, or a web artifact, or a cloud database, or a mobile device (See also: [SPD-Mobile](#), [SPD-Web](#))
9. Capture and analyze network traffic (See also: [NC-Security](#))
10. Develop approaches to address the challenges associated with mobile device forensics
11. Apply forensics tools to investigate security breaches
12. Identify and mitigate anti-forensic methods

SEC-Governance: Security Governance

KA Core:

1. Protecting critical assets from threats
2. Security governance: organizational objectives and general risk assessment
3. Security management: achieve and maintain appropriate levels of confidentiality, integrity, availability, accountability, authenticity, and reliability (See also: [SE-Validation](#))
4. Security policy: organizational policies, issue-specific policies, system-specific policies
5. Approaches to identifying and mitigating risks to computing infrastructure
6. Data lifecycle management policies: data collection, backups, and retention; cloud storage and services; breach disclosure. (See also: [DM-Security](#))

Illustrative Learning Outcomes:

KA Core:

1. Describe critical assets and how they can be protected
2. Differentiate between security governance, management, and controls, giving examples of each

3. Describe a technical control and implement it to mitigate against specific threats
4. Identify and assess risk of programs and database applications causing breaches
5. Design and implement appropriate backup strategies conforming to a given policy
6. Discuss a breach disclosure policy based on legal requirements and implement the policy
7. Identify the risks and benefits of outsourcing to the cloud

Professional Dispositions

- **Meticulous:** students need to pay careful attention to details to ensure the protection of real-world software systems.
- **Self-directed:** students must be ready to deal with the many novel and easily unforeseeable ways in which adversaries might launch attacks.
- **Collaborative:** students must be ready to collaborate with others, as collective knowledge and skills will be needed to prevent attacks, protect systems and data during attacks, and plan for the future after the immediate attack has been mitigated.
- **Responsible:** students need to show responsibility when designing, developing, deploying, and maintaining secure systems, as their enterprise and society is constantly at risk.
- **Accountable:** students need to know that as future professionals that they will be held accountable if a system or data breach were to occur, which should strengthen their resolve to prevent such breaches from occurring in the first place.

Mathematics Requirements

Required:

- [MSF-Discrete](#)
- [MSF-Probability](#)
- [MSF-Statistics](#)

Desired:

- [MSF-Linear](#)

Course Packaging Suggestions

There are two suggestions for course packaging, along with an additional suggestion for a more advanced course.

The first suggestion for course packaging is to infuse the CS Core hours of the SEC KA into appropriate places in other coursework that covers related security topics in the following knowledge units. As the CS Core Hours of the SEC KA are only 6 hours, coursework covering one or more of the following knowledge units could accommodate them.

- [AI-SEP](#)
- [AL-SEP](#)

- [AR-Assembly](#)
- [AR-Memory](#)
- [DM-Security](#)
- [FPL-Translation](#)
- [FPL-Run-Time](#)
- [FPL-Analysis](#)
- [FPL-Types](#)
- [HCI-Design](#)
- [HCI-Accountability](#)
- [HCI-SEP](#)
- [NC-Security](#)
- [OS-Protection](#)
- [PDC-Communication](#)
- [PDC-Coordination](#)
- [PDC-Evaluation](#)
- [SDF-Fundamentals](#)
- [SDF-Practices](#)
- [SE-Validation](#)
- [SEP-Privacy](#)
- [SEP-Security](#)
- [SF-Design](#)
- [SF-Security](#)
- [SPD-Common](#)
- [SPD-Mobile](#)
- [SPD-Web](#)

The second approach for course packaging is to create an additional full course focused on security that packages the following, building on the topics already covered in other knowledge areas.

Fundamentals of Computer Security:

- [SEC-Foundations](#) (6 hours)
- [SEC-SEP](#) (4 hours)
- [SEC-Coding](#) (7 hours)
- [SEC-Crypto](#) (5 hours)
- [SEC-Engineering](#) (4 hours)
- [SEC-Forensics](#) (2 hours)
- [SEC-Governance](#) (1 hour)
- [AI-SEP](#) (1 hour)
- [AR-Assembly](#) (1 hour)
- [AR-Memory](#) (1 hour)
- [DM-Security](#) (3 hours)
- [FPL-Translation](#) (1 hour)
- [FPL-Run-Time](#) (1 hour)
- [FPL-Analysis](#) (1 hour)

- [FPL-Types](#) (2 hours)
- [HCI-Design](#) (1 hour)
- [HCI-Accountability](#) (1 hour)
- [HCI-SEP](#) (1 hour)
- [NC-Security](#) (2 hours)
- [OS-Protection](#) (1 hour)
- [PDC-Communication](#) (1 hour)
- [PDC-Coordination](#) (1 hour)
- [PDC-Evaluation](#) (1 hour)
- [SDF-Fundamentals](#) (1 hour)
- [SDF-Practices](#) (1 hour)
- [SE-Validation](#): (2 hours)
- [SEP-Privacy](#) (1 hour)
- [SEP-Security](#) (2 hours)
- [SF-Design](#) (2 hours)
- [SF-Security](#) (2 hours)
- [SPD-Common](#) (2 hours)
- [SPD-Mobile](#) (2 hours)
- [SPD-Web](#): Web Platforms (2 hours)

The coverage exceeds 45 lecture hours, and so in a typical course, instructors would need to decide what topics to emphasize and what not to cover without losing the perspective that the course should help students develop a *security mindset*.

Prerequisites: Depends on the selected topics, but appropriate coursework covering MSF, SDF and SE knowledge units is needed.

Course objectives: Students should develop a security mindset and be ready to apply this mindset to securing data, software, systems, and applications.

A third suggested packaging is to create an advanced course that develops a security architect/engineer's view by including the following:

Security Engineering:

- [SEC-Foundations](#) (6 hours)
- [SEC-SEP](#) (4 hours)
- [SEC-Coding](#) (6 hours)
- [SEC-Crypto](#) (2 hours)
- [SEC-Engineering](#) (10 hours)
- [SEC-Forensics](#) (2 hours)
- [SEC-Governance](#) (1 hour)
- [DM-Security](#) (2 hours)
- [NC-Security](#) (3 hours)
- [OS-Protection](#) (2 hours)

- [PDC-Evaluation](#) (2 hours)
- [SDF-Fundamentals](#) (1 hour)
- [SDF-Practices](#) (1 hour)
- [SE-Validation](#) (2 hours)
- [SEP-Privacy](#) (1 hour)
- [SEP-Security](#) (1 hour)
- [SF-Design](#) (2 hours)
- [SF-Security](#) (2 hours)
- [SPD-Mobile](#) (2 hours)
- [SPD-Web](#) (2 hours)

The coverage for all topics is over 45 lecture hours, and so instructors would need to decide what topics to emphasize and what not to cover without losing the perspective that the course should help students develop the security engineer's mindset. Laboratory time related to data and network security, web platform, secure coding and validation would be valuable aspects of this course.

Prerequisites: Depends on the selected topics, either the first or second packaging suggested above would be recommended based on degree program needs.

Course objectives: Computer science students should develop the mindset of a security engineer and be ready to apply this mindset to problems in designing and evaluating the security of a range of computing systems and information services.

Committee

Chair: Rajendra K. Raj, Rochester Institute of Technology, Rochester, NY, USA

Members:

- Vijay Anand, University of Missouri – St. Louis, St. Louis, MO, USA
- Diana Burley, American University, Washington, DC, USA
- Sherif Hazem, Central Bank of Egypt, Cairo, Egypt
- Michele Maasberg, United States Naval Academy, Annapolis, MD, USA
- Bruce McMillin, Missouri University of Science and Technology, Rolla, MO, USA
- Sumita Mishra, Rochester Institute of Technology, Rochester, NY, USA
- Nicolas Sklavos, University of Patras, Patras, Greece
- Blair Taylor, Towson University, Towson, MD, USA
- Jim Whitmore, Dickinson College, Carlisle, PA, USA

Contributors:

- Markus Geissler, Cosumnes River College, Sacramento, CA, USA
-
- Michael Huang, Rider University, Lawrenceville, NJ, USA
- Tim Preuss, Minnesota State Community and Technical College, Moorhead, MN, USA
- Daniel Zappala, Brigham Young University, Provo, UT, USA

Society, Ethics, and the Profession (SEP)

Preamble

The ACM Code of Ethics and Professional Conduct states: “Computing professionals’ actions change the world. To act responsibly, they should reflect upon the wider impacts of their work, consistently supporting the public good.” The IEEE Code of Ethics starts by recognizing “... the importance of our technologies in affecting the quality of life throughout the world.” The AAAI Code of Professional Ethics and Conduct begins with “Computing professionals, and in particular, AI professionals’ actions change the world. To act responsibly, they should reflect upon the wider impacts of their work, consistently supporting the public good.”

While technical issues dominate the computing curriculum, they do not constitute a complete educational program in the broader context. It is more evident today than ever that students must also be exposed to the larger societal context of computing to develop an understanding of the critical and relevant social, ethical, legal, and professional issues and responsibilities at hand. This need to incorporate the study of these non-technical issues into the ACM curriculum was formally recognized in 1991, as articulated in the following excerpt from CS1991 [3]:

Undergraduates also need to understand the basic cultural, social, legal, and ethical issues inherent in the discipline of computing. They should understand where the discipline has been, where it is, and where it is heading. They should also understand their individual roles in this process, as well as appreciate the philosophical questions, technical problems, and aesthetic values that play an important part in the development of the discipline.

Students also need to develop the ability to ask serious questions about the social impact of computing and to evaluate proposed answers to those questions. Future practitioners must be able to anticipate the impact of introducing a given product into a given environment. Will that product enhance or degrade the quality of life? What will the impact be upon individuals, groups, and institutions?

Finally, students need to be aware of the basic legal rights of software and hardware vendors and users, and they also need to appreciate the ethical values that are the basis for those rights. Future practitioners must understand the responsibility that they will bear, and the possible consequences of failure. They must understand their own limitations as well as the limitations of their tools. All practitioners must make a long-term commitment to remaining current in their chosen specialties and in the discipline of computing as a whole.

Nonetheless, in recent years myriad high-profile issues affecting society at large have occurred leading to the conclusion that computer science professionals are not as prepared as they should be.

As technological advances (more specifically, how these advances are used by humans) continue to significantly impact the way we live and work, the critical importance of social and ethical issues and professional practice continues to increase in magnitude and consequence. The ways humans use

computing products and platforms, while hopefully providing opportunities, also introduce ever more challenging problems. A recent example is the emergence of generative AI, including large language models that generate code. A 2020 *Communications of the ACM* article [29] stated: "... because computing as a discipline is becoming progressively more entangled within the human and social lifeworld, computing as an academic discipline must move away from engineering-inspired curricular models and integrate the analytic lenses supplied by social science theories and methodologies."

In parallel to, and as part of, the heightened awareness of the social consequences computing has on the world, computing communities have become much more aware - and active - in areas of diversity, equity, inclusion and accessibility. These feature in statements and initiatives at ACM [33], IEEE [34], and AAAI [35] and in their codes of conduct [3-5]. All students deserve an inclusive, diverse, equitable and accessible learning environment. Computing students also have a unique duty to ensure that when put to practice, their skills, knowledge, and competencies are applied in ways that work for, and not against, the principles of diversity, equity, inclusion and accessibility. These principles are inherently a part of computing, and a new knowledge unit "Diversity, Equity, Inclusion and Accessibility" ([SEP-DEIA](#)) has been added to this knowledge area.

Computer science educators may opt to deliver the material in this knowledge area within the contexts of traditional technical and theoretical courses, in dedicated courses, and as part of capstone, project, and professional practice courses. The material in this knowledge area is best covered through a combination of all the above. It is too commonly held that many topics in this knowledge area may not readily lend themselves to being covered in other more traditional computer science courses. However, many of these topics naturally arise in traditional courses, or can be included with minimal effort. The benefits of exposing students to SEP topics within the context of those traditional courses is invaluable. Nonetheless institutional challenges will present barriers; for instance, some of these traditional courses may not be offered at a given institution and, in such cases, it is difficult to cover these topics appropriately without a dedicated SEP course. If social, ethical, and professional considerations are covered only in a dedicated course and not in the context of others, it could reinforce the false notion that technical processes are void of these important aspects, or that they are more isolated than they are in reality. Because of the broad relevance of these knowledge units, it is important that as many traditional courses as possible include aspects such as case studies, that analyze ethical, legal, social, and professional considerations in the context of the technical subject matter of those courses. Courses in areas such as software engineering, databases, computer graphics, computer networks, information assurance & security, and introduction to computing, all provide obvious context for analysis of such issues. However, an ethics-related module could be developed for almost any program. It would be explicitly against the spirit of these recommendations to have *only* a dedicated course within a specific computer science curriculum without great practical reason. Further, these topics should be covered in courses starting from year 1. Presenting them as advanced topics in later courses only creates an artificial perception that SEP topics are only important at a certain level or complexity. While it is true that the importance and consequence of SEP topics *increases* with level and complexity, introductory topics are not devoid of SEP topics. Further, many SEP topics are *best* presented early to lay a foundation for more intricate topics later in the curriculum.

Running through all the topics in this knowledge area is the need to speak to the computing practitioner's responsibility to proactively address issues through both ethical and technical actions. Today it is important not only for the topics in this knowledge area, but for students' knowledge in general, that the ethical issues discussed in any course should be directly related to - and arise naturally from - the subject matter of that course. Examples include a discussion in a database course of the SEP aspects of data aggregation or data mining; or a discussion in a software engineering course of the potential conflicts between obligations to the customer and users as well as all others affected by their work. Computing faculty who are unfamiliar with the content and/or pedagogy of applied ethics are urged to take advantage of the considerable resources from ACM, IEEE-CS, AAAI, SIGCAS (ACM Special Interest Group on Computers and Society), and other organizations. Additionally, it is the educator's responsibility to impress upon students that this area is just as important - in ways more important - than technical areas. The societal, ethical, and professional knowledge gained in studying topics in this knowledge area will be used throughout one's career and are transferable between projects, jobs, and often even industries, particularly as one's career progresses into project leadership and management.

The ACM Code of Ethics and Professional Conduct [30], the IEEE Code of Ethics [31], and the AAAI Code of Professional Ethics and Conduct [32] provide guidance that serve as the basis for the conduct of all computing professionals in their work. The ACM Code emphasizes that ethical reasoning is not an algorithm to be followed, and computer professionals are expected to consider how their work impacts the public good as the primary consideration. It falls to computing educators to highlight the domain-specific role of these topics for our students, but computer science programs should certainly be willing to lean on complementary courses from the humanities and social sciences.

Most computing educators are not also moral philosophers. Yet CS2023, along with past CS curricular recommendations, indicate the need for ethical analysis. CS2023 and prior curricular recommendations are quite clear on the required mathematical foundations that students are expected to gain which are often delivered by mathematics departments. Yet, the same is not true of moral philosophy. No one would expect a student to be able to provide a proof by induction until after having successfully completed a course in discrete mathematics. Yet, the parallel with respect to ethical analyses is somehow absent. We seemingly do (often) expect our students to perform ethical analysis without having the appropriate prerequisite knowledge from philosophy. Further, the application of ethical analysis also underlies every other knowledge unit in this knowledge area. We acknowledge that the knowledge unit Methods for Ethical Analysis ([SEP-Ethical-Analysis](#)) is the only one in this knowledge area that does not readily lend itself to being taught in the context of other CS2023 knowledge areas. Suggestions in terms of addressing this appear in the Course Packaging Suggestions.

The lack of prerequisite training in social, ethical, and professional topics has facilitated graduates operating with a certain ethical egoism (e.g., 'Here is what I believe/think/feel is right'). Regardless of how well intentioned, one might conclude that this is what brought us to a point in history where there are frequent occurrences of unintended consequences of technology, serious data breaches, and software failures causing economic, emotional and physical harms. Certainly, computing graduates who have learned how to apply the various ethical frameworks or lenses proposed through the ages would only serve to improve this situation. In retrospect, to ignore the lessons from moral philosophy,

which have been debated and refined for millennia - on what it means to act justly, or work for the common good - appears as hubris.

A computer science student must not graduate without understanding how society and ethics influence the computing profession. Nor should it be possible to complete a computer science degree without learning how computing professionals influence society, the ethical considerations involved in shaping that impact, and the student-turned-graduate's role in these relationships, as both computing professionals and members of society.

Changes Since CS2013

- The overall number of hours dedicated to this knowledge unit has doubled since CS2013 from 16 to 32 - CS2013 had 11 'Tier 1' hours and 5 'Tier 2' hours while CS2023 has 18 CS Core hours and 14 KA Core hours). However, many of these hours are best covered within the context of other knowledge areas as discussed in the introduction to CS2023. Additionally:
 - Several knowledge units that were 'elective only' in CS2013 now have CS Core and KA Core hours.
 - The number of hours in most knowledge units has increased.
- SEP has been re-titled from "Social Issues and Professional Practice" to "Society, Ethics, and the Profession". Our rationale follows.
 - Professional practice is an important part of being a professional. However, a solid understanding of the profession in which one is to become a professional is needed - indeed it is a prerequisite. Through this lens, professional conduct is part of a profession, along with its history, values, norms, etc. Additionally, "the profession" appears in the ACM tagline "Advancing Computing as a Science & Profession", and the ACM code [30]: "The ACM Code of Ethics and Professional Conduct ...expresses the conscience of the profession." In a similar way the AAAI Code [32] states that it "...expresses the conscience of the AI profession", and the IEEE code [31] states: "in accepting a personal obligation to our profession, its members and the communities we serve, do hereby commit ourselves to the highest ethical and professional conduct." While studying computing, students should not only learn how to become professionals, but also learn about the profession in which they will become professionals.
- Inclusion of the Diversity, Equity, Inclusion and Accessibility (DEIA) knowledge unit.
- Changed titles of two knowledge units
 - Professional Communication -> Communication
 - Analytical Tools -> Methods for Ethical Analysis

Core Hours

Knowledge Unit	CS Core	KA Core
Social Context	3	2
Methods for Ethical Analysis	2	1
Professional Ethics	2	2
Intellectual Property	1	1
Privacy and Civil Liberties	2	1
Communication	2	1
Sustainability	1	1
History	1	1
Economies of Computing	0	1
Security Policies, Laws, and Computer Crimes	2	1
Diversity, Equity, Inclusion and Accessibility	2	2
Total	18	14

Knowledge Units

SEP-Context: Social Context

Computers, the internet, and artificial intelligence - perhaps more than any other technologies - have transformed society over the past several decades, with dramatic increases in human productivity; an explosion of options for news, entertainment, and communication; and fundamental breakthroughs in almost every branch of science and engineering. It is imperative to recognize that this is not a one-way street. Society also affects computing, resulting in a complex socio-technical context that is constantly changing, requiring the perspective of history to put the present, as well as possible futures, into appropriate perspective.

Social Context provides the foundation for all other knowledge units in SEP, particularly Professional Ethics.

CS Core:

1. Social implications (e.g., political, and cultural ideologies) in a hyper-networked world where the capabilities and impact of social media, artificial intelligence and computing in general are rapidly evolving

2. Impact of computing applications (e.g., social media, artificial intelligence applications) on individual well-being, and safety of all kinds (e.g., physical, emotional, economic)
3. Consequences of involving computing technologies, particularly artificial intelligence, biometric technologies, and algorithmic decision-making systems, in civic life (e.g., facial recognition technology, biometric tags, resource distribution algorithms, policing software) and how human agency and oversight is crucial
4. How deficits in diversity and accessibility in computing affect society and what steps can be taken to improve equity in computing

KA Core:

5. Growth and control of the internet, data, computing, and artificial intelligence
6. Often referred to as the digital divide, differences in access to digital technology resources and its resulting ramifications for gender, class, ethnicity, geography, and/or developing countries, including consideration of responsibility to those who might be less wealthy, under threat, or who would struggle to have their voices heard.
7. Accessibility issues, including legal requirements such as Web Content Accessibility Guidelines (www.w3.org/TR/WCAG21)
8. Context-aware computing

Illustrative Learning Outcomes:

CS Core:

1. Describe different ways that computer technology (networks, mobile computing, artificial intelligence) mediates social interaction at the personal and collective levels.
2. Identify developers' assumptions and values embedded in hardware and software design, especially as they pertain to usability for diverse populations including under-served and those with disabilities.
3. Interpret the social context of a given design and its implementation.
4. Analyze the efficacy of a given design and implementation using empirical data.
5. Understand the implications of technology (e.g., social media) use, for different identities, cultures, and communities.

KA Core:

6. Describe the internet's role in facilitating communication between citizens, governments, and each other.
7. Analyze the effects of reliance on computing in the implementation of democracy (e.g., delivery of social services, electronic voting).
8. Describe the impact of a lack of appropriate representation of people from historically minoritized populations in the computing profession (e.g., industry culture, product diversity).
9. Discuss the implications of context awareness in ubiquitous computing systems.
10. Express how access to the internet and computing technologies affect different societies.
11. Identify why/how internet access can be viewed as a human right.

SEP-Ethical-Analysis: Methods for Ethical Analysis

Ethical theories and principles are the foundations of ethical analysis because they are the viewpoints which can provide guidance along the pathway to a decision. Each theory emphasizes different assumptions and methods for determining the ethicality of a given action. It is important for students to recognize that decisions in different contexts may require different ethical theories (including combinations) to arrive at ethically acceptable outcomes, and what constitutes 'acceptable' depends on a variety of factors such as cultural context. Applying methods for ethical analysis requires both an understanding of the underlying principles and assumptions guiding a given tool and an awareness of the social context for that decision. Traditional ethical frameworks (e.g. [36]) as provided by western philosophy can be useful, but they are not all-inclusive. Effort must be taken to include decolonial, indigenous, and historically marginalized ethical perspectives whenever possible. No theory will be universally applicable to all contexts, nor is any single ethical framework the 'best'. Engagement across various ethical schools of thought is important for students to develop the critical thinking needed in judiciously applying methods for ethical analysis of a given situation.

CS Core:

1. Avoiding fallacies and misrepresentation in argumentation
2. Ethical theories and decision-making (philosophical and social frameworks, e.g. [3])
3. Recognition of the role culture plays in our understanding, adoption, design, and use of computing technology
4. Why ethics is important in computing, and how ethics is similar to, and different from, laws and social norms

KA Core:

5. Professional checklists
6. Evaluation rubrics
7. Stakeholder analysis
8. Standpoint theory
9. Introduction to ethical frameworks (e.g., consequentialism such as utilitarianism, non-consequentialism such as duty, rights, or justice, agent-centered such as virtue or feminism, contractarianism, ethics of care) and their use for analyzing an ethical dilemma

Illustrative Learning Outcomes:

CS Core:

1. Describe how a given cultural context impacts decision making.
2. Express the use of example and analogy in ethical argument.
3. Analyze (and avoid) basic logical fallacies in an argument.
4. Analyze an argument to identify premises and conclusion.
5. Evaluate how and why ethics is so important in computing and how it relates to cultural norms, values, and law.
6. Justify a decision made on ethical grounds.

KA Core:

7. Distinguish all stakeholder positions in relation to their cultural context in a given situation.
8. Analyze the potential for introducing or perpetuating ethical debt (deferred consideration of ethical impacts or implications) in technical decisions.
9. Discuss the advantages and disadvantages of traditional ethical frameworks
10. Analyze ethical dilemmas related to the creation and use of technology from multiple perspectives using ethical frameworks

SEP-Professional-Ethics: Professional Ethics

Computer ethics is a branch of practical philosophy that deals with how computing professionals should make decisions regarding professional and social conduct. There are three primary influences: 1) The individual's own personal ethical code, 2) Any informal or formal regulation/decreed/etc. of ethical behavior existing in the workplace, applicable licensures, certifications, or laws, and 3) Exposure to formal codes of ethics and ethical frameworks.

CS Core:

1. Community values and the laws by which we live
2. The nature of being a professional including care, attention, discipline, fiduciary responsibility, and mentoring
3. Keeping up to date as a computing professional in terms of familiarity, tools, skills, legal and professional frameworks as well as the ability and responsibility to self-assess and progress in the computing field
4. Professional certification, codes of ethics, conduct, and practice, such as the ACM, IEEE, AAAI, and other international societies
5. Accountability, responsibility, and liability (e.g., software correctness, reliability and safety, warranty, negligence, strict liability, ethical approaches to security vulnerability disclosures) including whether a product/service should be built, not just doing so because it is technically possible.
6. Introduction to theories describing the human creation and use of technology including instrumentalism, sociology of technological systems, disability justice, neutrality thesis, pragmatism, and decolonial models, including developing and using technology to right wrongs and do good
7. Strategies for recognizing and reporting designs, systems, software, and professional conduct (or their outcomes) that may violate law or professional codes of ethics

KA Core:

8. The role of the computing professional and professional societies in public policy
9. Maintaining awareness of consequences
10. Ethical dissent and whistleblowing
11. The relationship between regional culture and ethical dilemmas
12. Dealing with harassment and discrimination
13. Forms of professional credentialing
14. Ergonomics and healthy computing environments

15. Time-to-market and cost considerations versus quality professional standards

Illustrative Learning Outcomes:

CS Core:

1. Identify ethical issues that arise in software design, development practices, and software deployment.
2. Discuss how to address ethical issues in specific situations.
3. Express the ethical responsibility of ensuring software correctness, reliability and safety including from where this responsibility arises (e.g., ACM/IEEE/AAAI Codes of Ethics, laws and regulations, organizational policies).
4. Describe the mechanisms that typically exist for a professional to keep up to date in ethical matters.
5. Describe the strengths and weaknesses of relevant professional codes as expressions of being a professional and guides to decision-making.
6. Analyze a global computing issue, observing the role of professionals and government officials in managing this problem.

KA Core:

8. Describe ways in which professionals and professional organizations may contribute to public policy.
9. Describe the consequences of inappropriate professional behavior.
10. Be familiar with whistleblowing and have access to knowledge to guide one through an incident.
11. Identify examples of how regional culture interplays with ethical dilemmas.
12. Describe forms of harassment and discrimination and avenues of assistance.
13. Assess various forms of professional credentialing.
14. State the relationship between ergonomics in computing environments and people's health.
15. Describe issues associated with industries' push to focus on time-to-market versus enforcing quality professional standards.

SEP-IP: Intellectual Property

Intellectual property refers to a range of intangible rights of ownership in any product of human intellect, such as a software program. Laws, which vary by locality, provide different methods for protecting these rights of ownership based on their type. Ideally, intellectual property laws balance the interests of creators and users of the property. There are essentially four types of intellectual property rights relevant to software: patents, copyrights, trade secrets and trademarks. Moreover, property rights are also often protected by user licenses. Each affords a different type of legal protection.

CS Core:

1. Intellectual property rights
2. Intangible digital intellectual property (IDIP)
3. Legal foundations for intellectual property protection
4. Common software licenses (e.g., MIT, GPL and its variants, Apache, Mozilla, Creative Commons)

5. Plagiarism and authorship

KA Core:

6. Philosophical foundations of intellectual property
7. Forms of intellectual property (e.g., copyrights, patents, trade secrets, trademarks) and the rights they protect
8. Limitations on copyright protections, including fair use and the first sale doctrine
9. Intellectual property laws and treaties that impact the enforcement of copyrights
10. Software piracy and technical methods for enforcing intellectual property rights, such as digital rights management and closed source software as a trade secret
11. Moral and legal foundations of the open source movement
12. Systems that use others' data (e.g., large language models)

Illustrative Learning Outcomes:

CS Core:

1. Describe and critique legislation and precedent aimed at digital copyright infringements.
2. Identify contemporary examples of intangible digital intellectual property.
3. Select an appropriate software license for a given project.
4. Defend legal and ethical uses of copyrighted materials.
5. Interpret the intent and implementation of software licensing.
6. Discuss whether a use of copyrighted material is likely to be fair use.
7. Analyze the ethical issues inherent in various plagiarism detection mechanisms.
8. Identify multiple forms of plagiarism beyond verbatim copying of text or software (e.g., intentional paraphrasing, authorship misrepresentation, and improper attribution).

KA Core:

9. Discuss the philosophical bases of intellectual property in an appropriate context (e.g., country, etc.).
10. Distinguish the conflicting issues involved in securing software patents.
11. Contrast the protections and obligations of copyright, patent, trade secret, and trademarks.
12. Describe the rationale for the legal protection of intellectual property in the appropriate context (e.g., country, etc.).
13. Analyze the use of copyrighted work under the concepts of fair use and the first sale doctrine.
14. Identify the goals of the open source movement and its impact on fields beyond computing, such as the right-to-repair movement.
15. Summarize the global nature of software piracy.
16. Criticize the use of technical measures of digital rights management (e.g., encryption, watermarking, copy restrictions, and region lockouts) from multiple stakeholder perspectives.
17. Discuss the nature of anti-circumvention laws in the context of copyright protection.

SEP-Privacy: Privacy and Civil Liberties

Electronic information sharing highlights the need to balance privacy protections with information access. The ease of digital access to, in addition to copying and distribution of, many types of data makes privacy rights and civil liberties more complex, especially given cultural and legal differences in these areas. Complicating matters further, privacy also has interpersonal, organizational, professional/business, and governance components. In addition, the interconnected nature of online communities raises challenges for managing expectations and protections for freedom of expression in various cultures and nations. Technology companies that provide platforms for user-generated content are under increasing pressure to perform governance tasks, potentially facing liability for their decisions.

CS Core:

1. Privacy implications of widespread data collection including but not limited to transactional databases, data warehouses, surveillance systems, cloud computing, and artificial intelligence
2. Conceptions of anonymity, pseudonymity, and identity
3. Technology-based solutions for privacy protection (e.g., end-to-end encryption and differential privacy)
4. Civil liberties, privacy rights, and cultural differences

KA Core:

5. Philosophical and legal conceptions of the nature of privacy including the right to privacy
6. Legal foundations of privacy protection in relevant jurisdictions (e.g., GDPR in the EU)
7. Privacy legislation in areas of practice (e.g., HIPAA in the US, AI Act in the EU)
8. Basic Principles of human-subjects research and principles beyond what the law requires (e.g., Belmont Report, UN Universal Declaration on Human Rights and how this relates to technology)
9. Freedom of expression and its limitations
10. User-generated content, content moderation, and liability

Illustrative Learning Outcomes:

CS Core:

1. Evaluate solutions to privacy threats in transactional databases and data warehouses.
2. Describe the role of data collection in the implementation of pervasive surveillance systems (e.g., RFID, face recognition, toll collection, mobile computing).
3. Distinguish the concepts and goals of anonymity and pseudonymity
4. Describe the ramifications of technology-based privacy protections, including differential privacy and end-to-end encryption
5. Identify cultural differences regarding the nature and necessity of privacy and other civil liberties.

KA Core:

6. Discuss the philosophical basis for the legal protection of personal privacy in an appropriate context (e.g., country, etc.).

7. Critique the intent, potential value, and implementation of various forms of privacy legislation and principles beyond what the law requires.
8. Identify strategies to enable appropriate freedom of expression.

SEP-Communication: Communication

Computing is an inherently collaborative and social discipline making communication an essential aspect of the profession. Much but not all of this communication occurs in a professional setting where communication styles, expectations, and norms differ from other contexts where similar technology might be used. Both professional and informal communication conveys information to various audiences who may have very different goals and needs for that information. Good communication is also necessary for transparency and trustworthiness. It is also important to note that computing professionals are not just communicators but are also listeners who must be able to hear and thoughtfully make use of feedback received from various stakeholders. Effective communication skills are not something one 'just knows' - they are developed and can be learned. Communication skills are best taught in context throughout the undergraduate curriculum.

CS Core:

1. Oral, written, and electronic team and group communication
2. Technical communication materials (e.g., source code, and documentation, tutorials, reference materials, API documentation)
3. Communicating with different stakeholders such as customers, leadership, or the general public
4. Team collaboration (including tools) and conflict resolution
5. Accessibility and inclusivity requirements for addressing professional audiences
6. Cultural competence in communication including considering the impact of difference in natural language

KA Core:

7. Trade-offs in competing factors that affect communication channels and choices
8. Communicating to solve problems or make recommendations in the workplace, such as raising ethical concerns or addressing accessibility issues

Illustrative Learning Outcomes:

CS Core:

1. Understand the importance of writing concise and accurate technical documents following well-defined standards for format and for including appropriate tables, figures, and references.
2. Analyze written technical documentation for technical accuracy, concision, lack of ambiguity, and awareness of audience.
3. Compose and deliver an audience aware, accessible, and organized formal presentations.
4. Plan interactions (e.g., virtual, face-to-face, shared documents) with others in ways that invite inclusive participation, model respectful consideration of others' contributions, and explicitly value diversity of ideas.

5. Identify and describe qualities of effective communication (e.g., virtual, face-to-face, intragroup, shared documents).
6. Understand how to communicate effectively and appropriately as a member of a team including conflict resolution techniques.
7. Discuss ways to influence performance and results in diverse and cross-cultural teams.

KA Core:

8. Assess personal strengths and weaknesses to work remotely as part of a team drawing from diverse backgrounds and experiences.
9. Choose an appropriate way to communicate delicate ethical concerns.

SEP-Sustainability: Sustainability

Sustainability is defined by the United Nations as “development that meets the needs of the present without compromising the ability of future generations to meet their own needs” [37]. Alternatively, it is the “balance between the environment, equity and economy” [38]. As computing extends into more and more aspects of human existence, we are already seeing estimates that double-digit percentages of global electricity usage are consumed by computing activities, which unchecked will likely grow. Further, electronics contribute individually to demand for rare earth elements, mineral extraction, and countless e-waste concerns. Students should gain a background that recognizes these global and environmental costs and their potential long-term effects on the environment and local communities.

CS Core:

1. Environmental, social, and cultural impacts of implementation decisions (e.g., sustainability goals, algorithmic bias/outcomes, economic viability, and resource consumption)
2. Local/regional/global social and environmental impacts of computing systems and their use (e.g., carbon footprints, resource usage, e-waste) in hardware (e.g., e-waste, data centers, rare element and resource utilization, recycling) and software (e.g., cloud-based services, blockchain, AI model training and use), not neglecting the impact of everyday use such as hardware (cheap hardware replaced frequently) and software (web-browsing, email, and other services with hidden/remote computational demands).
3. Guidelines for sustainable design standards

KA Core:

4. Systemic effects of complex computing technologies and phenomena (e.g., generative AI, data centers, social media, offshoring, remote work)
5. Pervasive computing: Information processing that has been integrated into everyday objects and activities, such as smart energy systems, social networking, and feedback systems to promote sustainable behavior, transportation, environmental monitoring, citizen science and activism
6. How the sustainability of software systems is interdependent with social systems, including the knowledge and skills of its users, organizational processes and policies, and its societal context (e.g., market forces, government policies)

Illustrative Learning Outcomes:

CS Core:

1. Identify ways to be a sustainable practitioner in a specific area or with a specific project.
2. Assess the environmental impacts of a given project's deployment (e.g., energy consumption, contribution to e-waste, impact of manufacturing).
3. Describe global social and environmental impacts of computer use and disposal.
4. List the sustainable effects of modern practices and activities (e.g., remote work, e-commerce, cryptocurrencies, AI models, data centers).

KA Core:

5. Describe the environmental impacts of design choices within the field of computing that relate to algorithm design, operating system design, networking design, database design, etc.
6. Analyze the social and environmental impacts of new system designs.
7. Design guidelines for sustainable IT design or deployment in areas such as smart energy systems, social networking, transportation, agriculture, supply-chain systems, environmental monitoring, and citizen activism.
8. Assess computing applications in respect to environmental issues (e.g., energy, pollution, resource usage, recycling and reuse, food management and production).

SEP-History: Computing History

History is important because it provides a mechanism for understanding why our computing systems operate the way they do, the societal contexts in which current approaches arose, and how those continue to echo through the discipline today. Not only does computing affect society but vice-versa, resulting in a complex socio-technical context that is constantly changing, requiring the perspective of history to put the present, as well as possible futures, into appropriate perspective. It also informs decisions based on successes and failures of the past including harms done and how to not repeat them. The history of computing is often taught in context with foundational concepts, such as system fundamentals and software development fundamentals. A focus should be put on those who, due to marginalization, have not historically featured as prominently as they should.

CS Core:

1. The history of computing: hardware, software, and human/organizational.
2. The role of history in the present including within different social contexts, and the relevance of this history on the future.

KA Core:

3. Age I (Pre-digital): Ancient analog computing (Stonehenge, Antikythera mechanism, Salisbury Cathedral clock, etc.), human-calculated number tables, Euclid, Lovelace, Babbage, Gödel, Church, Turing, pre-electronic (electro-mechanical and mechanical) hardware
4. Age II (Early modern computing): ENIAC, UNIVAC, Bombes (Bletchley Park and codebreakers), computer companies (e.g., IBM), mainframes, etc.

5. Age III (PC era): PCs, modern computer hardware and software, Moore's Law
6. Age IV (Internet): Networking, internet architecture, browsers and their evolution, standards, born-on-the-internet companies, and services (e.g., Google, Amazon, Microsoft, etc.), distributed computing
7. Age V (Mobile & Cloud): Mobile computing and smartphones, cloud computing and models thereof (e.g., SaaS), remote servers, security and privacy, social media
8. Age VI (AI): Decision making systems, recommender systems, generative AI and other machine learning driven tools and technologies

Illustrative Learning Outcomes:

CS Core:

1. Understand the relevance and impact of computing history on recent events, present context, and possible future outcomes, from more than one cultural perspective.
2. Discuss how perspectives held today have been shaped by history, and that alternative perspectives exist (e.g., fears of AI replacing human workers vs AI augmenting human work, various views on Moore's Law).

KA Core:

3. Identify formative and consequential trends in the history of the computing field.
4. Identify the contributions of several pioneering individuals or organizations (research labs, computer companies, government offices) in the computing field.
5. Discuss the historical context for important moments in history of computing, such as the move from vacuum tubes to transistors (TRADIC), early seminal operating systems (e.g., OS 360), Xerox PARC and the first Apple computer with a GUI, the creation of specific programming language paradigms, the first computer virus, the creation of the internet, the creation of the WWW, the dot com bubble, Y2K, the introduction of smartphones, etc.
6. Compare daily life before and after the advent of milestone developments (e.g., personal computers or the internet).

SEP-Economies: Economies of Computing

The economies of computing are important to those who develop and provide computing resources and services to others as well as society in general. They impact users of these resources and services, both professional and non-professional. Computing professionals have a duty to know the impact of these topics on their own roles and activities and how choices made will impact users and society.

KA Core:

1. Economic models: regulated and unregulated, monopolies, network effects, and open market; knowledge and attention economies
2. Pricing and deployment strategies: planned obsolescence, subscriptions, freemium, software licensing, open-source, free software, adware
3. Impacts of differences in access to computing resources, and the effect of skilled labor supply and demand on the quality of computing products

4. Automation, AI, and their effects on job markets, developers, and users
5. Ethical concerns surrounding the attention economy and other economies of computing (e.g. informed consent, data collection, use of verbose legalese in user agreements)

Illustrative Learning Outcomes:

KA Core:

1. Summarize the social effects of economic models (e.g., the knowledge and attention economies).
2. Describe the differences and similarities of competing strategies (e.g., subscription vs freemium vs free).
3. Discuss several examples of digital divides.
4. Understand the effects of automation and AI on society.
5. Understand the ethical implications of computing economies that rely on attention and data.

SEP-Security: Security Policies, Laws and Computer Crimes

While security policies, laws and computer crimes are important topics, it is essential they are viewed with the foundation of other social and professional knowledge units, such as [Intellectual Property](#), [Privacy and Civil Liberties](#), [Social Context](#), and [Professional Ethics](#). Computers, the internet, and artificial intelligence, perhaps more than any other technologies, have transformed society over the past 75 years. At the same time, they have contributed to unprecedented threats to privacy; new categories of computer crime and antisocial behavior; major disruptions to organizations; and the large-scale concentration of risk into information systems.

CS Core:

1. Computer crimes, legal redress for computer criminals and impact on victims and society
2. Social engineering, computing-enabled fraud, identity theft and recovery from these
3. Cyber terrorism, criminal hacking, and hacktivism
4. Malware, viruses, worms
5. Attacks on critical infrastructure such as electrical grids and pipelines
6. Non-technical fundamentals of security (e.g., human engineering, policy, confidentiality)

KA Core:

7. Benefits and challenges of existing and proposed computer crime laws
8. Security policies and the challenges of change and compliance
9. Responsibility for security throughout the computing life cycle
10. International and local laws and how they intersect

Illustrative Learning Outcomes:

CS Core:

1. List classic examples of computer crimes and social engineering incidents with societal impact.
2. Identify issues with laws that apply to computer crimes.

3. Describe the motivation and ramifications of cyber terrorism, data theft, hacktivism, ransomware, and other attacks.
4. Examine the ethical and legal issues surrounding the misuse of access and various breaches of security.
5. Discuss the professional's role in security and the trade-offs and challenges involved.

KA Core:

6. Investigate measures that can be taken by both individuals and organizations including governments to prevent or mitigate the undesirable effects of computer crimes.
7. Design a company-wide security policy, which includes procedures for managing passwords and employee monitoring.
8. Understand how legislation from one region may affect activities in another (e.g., how EU GDPR applies globally, when EU persons are involved).

SEP-DEIA: Diversity, Equity, Inclusion, and Accessibility

Despite being a creative, often highly compensated field with myriad job (and other) opportunities, racial, gender, and many other inequities in representation are pervasive in many regions. For too many students, their first computer science course is their last. There are many factors including the legacy of systemic racism, ableism, sexism, classism, and other injustices that contribute to the lack of diverse identities within computer science, and there is no single or quick fix.

CS2023's sponsoring organizations are ACM, IEEE, and AAAI. Each of these places a high value on diversity, equity, inclusion and accessibility; our computer science classrooms should promote and model those principles. We should welcome and seek diversity - the gamut of human differences including gender identity, ethnicity, race, politics, abilities and attributes, religion, nationality, etc. - in our classrooms, departments, and campuses. We should strive to make our classrooms, labs, and curricula accessible and to promote inclusion. We should aim for students and all community members to have a sense of belonging that comes with being respected, wanted, and valued. To achieve equity, we must allocate resources, promote fairness, and check our biases to ensure persons of all identities achieve success. Accessibility should be considered and implemented in all computing activities and products.

Explicitly infusing diversity, equity, inclusion and accessibility (DEIA) across the computer science curriculum demonstrates its importance for the department, institution, and our field - all of which likely have a DEIA statement and/or initiative(s). This emphasis on DEIA is important ethically and a bellwether issue of our time. Many professionals in computing already recognize attention to DEIA as integral to disciplinary practice. Regardless of the degree to which these values appear in any one computer science class, research suggests that a lack of attention to them results in inferior designs in addition to harm and its perpetuation. Not only does data support that diverse teams outperform homogeneous ones, but diverse teams may better prevent egregious technology failures in recent headlines such as facial recognition misuse, airbag injuries and deaths, and other well-known failures of science and computing.

CS Core:

1. How identity impacts and is impacted by computing technologies and environments (academic and professional)
2. The benefits of diverse development teams and the impacts of teams that are not diverse
3. Inclusive language and charged terminology, and why their use matters
4. Inclusive behaviors and why they matter
5. Designing and developing technology with accessibility in mind
6. How computing professionals can influence and impact diversity, equity, inclusion and accessibility, including but not only through the software they create

KA Core:

7. Experts and their practices that who reflect the identities of the classroom and the world through practical DEIA principles
8. Historic marginalization due to systemic social mechanisms, technological supremacy and global infrastructure challenges to diversity, equity, inclusion and accessibility
9. Cross-cultural differences in, and needs for, diversity, equity, inclusion and accessibility

Illustrative Learning Outcomes:

CS Core:

1. Define and distinguish equity, equality, diversity, inclusion, and accessibility.
2. Identify language, practices, and behaviors that may make someone feel included in a workplace and/or a team, and why is it relevant. Avoid charged terminology - see *Words Matter* (www.acm.org/diversity-inclusion/words-matter) - this includes identifying and accommodating users that are often excluded without thought and not considered at all.
3. Evaluate the accessibility of your classroom or lab. Evaluate the accessibility of your webpage. (See www.w3.org/WAI and www.w3.org/TR/WCAG21).
4. Demonstrate collegiality and respect when working with team members who do not share your identity. *It is not enough to merely assign team projects. Faculty should prepare students for teamwork and monitor, mentor, and assess the effectiveness of their student teams throughout a project.*
5. Compare the demographics of your institution's computer science and STEM majors to the overall institutional demographics. If they differ, identify factors that contribute to inequitable access, engagement, and achievement in computer science among marginalized groups. If they do not, assess why not.
6. Identify developers' assumptions and values embedded in hardware and software design, especially as they pertain to usability by diverse populations.
7. Identify examples of the benefits that diverse teams can bring to software products, and how a lack of diversity has costs.

KA Core:

9. Analyze the work of experts who reflect the identities of the classroom and the world.
10. Assess the impact of power and privilege in the computing profession as it relates to culture, industry, products, and society.

11. Develop examples of systemic changes that could positively address diversity, equity, inclusion, and accessibility in a familiar context (i.e., in an introductory computing course) and an unfamiliar context and when these might be different, or the same.
12. Compare the demographics of your institution to the overall community demographics. If they differ, identify factors that contribute to inequitable access, engagement, and achievement among marginalized groups. If they do not, assess why not.

Professional Dispositions

- **Critically Self-reflective:** Students should be able to inspect their own actions, thoughts, biases, privileges, and motives to discover places where professional activity is not up to current standards. They must strive to understand both conscious and unconscious biases and continuously work to counteract them.
- **Responsive:** Students must quickly and accurately respond to changes in the field and adapt in a professional manner, such as shifting from in-person office work to remote work at home. These shifts require rethinking one's entire approach to what is considered "professional".
- **Proactive:** Students must be able to identify areas of importance (e.g., in accessibility and inclusion) and understand how to address them for a more professional working environment.
- **Culturally Competent:** Students must prioritize cultural competence—the ability to work with people from cultures different from one's own—by using inclusive language, watching for and counteracting conscious and unconscious biases, and encouraging honest and open communication.
- **Advocative:** Students must think, speak, and act in ways that foster and promote diversity, equity, inclusion and accessibility in all ways including but not limited to teamwork, communication, and product development (hardware and software).
- **Responsible:** Students must act responsibly in all areas of their work towards: all users and stakeholders including society at large; colleagues and their profession in general.

Course Packaging Suggestions

In computing, societal, and ethical considerations arise in all other knowledge areas and therefore should arise in the context of other computing courses, not just siloed in an "SEP course". These topics should be covered in courses starting from the first year (the only likely exception is [SEP-Ethical-Analysis: Methods for Ethical Analysis](#)) although this could be delivered as part of a first-year course or via a seminar or an online independent study.

Presenting SEP topics as advanced topics only covered in later courses could create the incorrect perception that SEP topics are only important at a certain level or complexity. While it is true that the importance and consequence of SEP topics *increases* with level and complexity, introductory topics are not devoid of SEP topics. Further, many SEP topics are *best* presented early to lay a foundation for more intricate topics later in the curriculum.

Instructor choice for some of these topics is complex. When SEP topics arise in other courses these are naturally often taught by the instructor teaching that course, although at times bringing in expert educators from other disciplines (e.g., law, ethics, etc.) could be advantageous. Stand-alone courses in

SEP - should they be needed - are likely best delivered by an interdisciplinary team. However this brings additional complexity. Regardless, who teaches SEP topics and/or courses warrants careful consideration.

At a minimum the SEP CS Core learning outcomes are best covered in the context of courses covering other knowledge areas - ideally the SEP KA Core hours are also, with the likely exception of [SEP-Ethical-Analysis](#). This knowledge unit (KU) underlies every other KU in the SEP knowledge area (KA). However, this KU is the only one in the SEP KA that does not readily lend itself to being taught in the context of other KAs. Delivering these topics warrants even more careful consideration as to how/where they will be covered, and who will teach them. In conjunction with covering SEP topics as they occur naturally in other KAs, dedicated SEP courses can add value. However a sole, stand-alone course in a program where SEP topics are not covered in other courses should be a last resort.

At some institutions, an **in-depth dedicated course** at the mid- or advanced-level may be offered covering all recommended topics in both the CS Core and KA Core KUs **in close coordination with learning outcomes best covered in the context of courses covering other KAs.** Such a course could include:

- [SEP-Context](#) (5 hours)
- [SEP-Ethical-Analysis](#) (3 hours)
- [SEP-Professional-Ethics](#) (4 hours)
- [SEP-IP](#) (2 hours)
- [SEP-Privacy](#) (3 hours)
- [SEP-Communication](#) (3 hours)
- [SEP-Sustainability](#) (2 hours)
- [SEP-History](#) (2 hours)
- [SEP-Economies](#) (1 hour)
- [SEP-Security](#) (3 hours)
- [SEP-DEIA](#) (4 hours)

Skill Statement

A student who completes this course should be able to contribute to systemic change by applying societal and ethical knowledge using relevant underpinnings and frameworks to their work in the computing profession in a culturally competent manner including contributing to positive developments in inclusion, equity, diversity, and accessibility in computing.

At some institutions, a **dedicated minimal course** may be offered covering the CS Core knowledge units **in close coordination with learning outcomes best covered in the context of courses covering other knowledge areas.** Such a course could include:

- [SEP-Context](#) (3 hours)
- [SEP-Ethical-Analysis](#) (2 hours)
- [SEP-Professional-Ethics](#) (2 hours)
- [SEP-IP](#) (1 hour)
- [SEP-Privacy](#) (2 hours)
- [SEP-Communication](#) (2 hours)
- [SEP-Sustainability](#) (1 hour)

- [SEP-History](#) (1 hour)
- [SEP-Security](#) (2 hours)
- [SEP-DEIA](#) (2 hours)

Skill Statement

A student who completes this course should be able to apply societal and ethical knowledge to their work in the computing profession while fostering and contributing to inclusion, equity, diversity, and accessibility in computing.

Exemplary Materials

- Emanuelle Burton, Judy Goldsmith, Nicholas Mattei, Cory Siler, and Sara-Jo Swiatek. 2023. Teaching Computer Science Ethics Using Science Fiction. In Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. 2 (SIGCSE 2023). Association for Computing Machinery, New York, NY, USA, 1184. <https://doi.org/10.1145/3545947.3569618>
- Randy Connolly. 2020. Why computing belongs within the social sciences. Commun. ACM 63, 8 (August 2020), 54–59. <https://doi.org/10.1145/3383444>
- Casey Fiesler. Tech Ethics Curricula: A Collection of Syllabi Used to Teach Ethics in Technology Across Many Universities
 - a. <https://cfiesler.medium.com/tech-ethics-curricula-a-collection-of-syllabi-3eedfb76be18>
 - b. [Tech Ethics Curricula](#)
- Casey Fiesler. Tech Ethics Readings: A Spreadsheet of Readings Used to Teach Ethics in Technology [Tech Ethics Class Readings](#)
- Stanford Embedded EthiCS, Embedding Ethics in Computer Science. <https://embeddedethics.stanford.edu/>
- Jeremy, Weinstein, Rob Reich, and Mehran Sahami. *System Error: Where Big Tech Went Wrong and How We Can Reboot*. Hodder Paperbacks, 2023.
- Baecker, R. *Computers in Society: Modern Perspectives*, Oxford University Press. (2019).
- Embedded EthiCS @ Harvard: bringing ethical reasoning into the computer science curriculum. <https://embeddedethics.seas.harvard.edu/about>

Committee

Chair: Brett A. Becker, University College Dublin, Dublin, Ireland

Members:

- Richard L. Blumenthal, Regis University, Denver, CO, USA
- Mikey Goldweber, Denison University, Granville, OH, USA
- James Prather, Abilene Christian University, Abilene, TX, USA
- Susan Reiser, University of North Carolina Asheville, Asheville, NC, USA
- Michelle Trim, University of Massachusetts Amherst, Amherst, MA, USA
- Titus Winters, Google, Inc, New York, NY, USA

Contributors:

- Jake Baskin, Computer Science Teachers Association, Chicago, IL, USA
- Johanna Blumenthal, Regis University, Denver, CO, USA

- Chris Stephenson, Google, Portland, OR, USA
- MaryAnne Egan, Siena College, Loudonville, NY, USA
- Catherine Mooney, University College Dublin, Dublin, Ireland
- Fay Cobb Payton, North Carolina State University, Raleigh, NC, USA
- Keith Quille, Technological University of Dublin, Dublin, Ireland
- Mehran Sahami, Stanford University, Stanford, CA, USA
- Mark Scanlon, University College Dublin, Dublin, Ireland
- Karren Shorofsky, University of San Francisco School of Law, San Francisco, CA, USA
- Andreas Stefik, University of Nevada, Las Vegas, Las Vegas, NV, USA
- Ellen Walker, Hiram College, Cleveland, OH, USA

Systems Fundamentals (SF)

Preamble

A computer system is a set of hardware and software infrastructures upon which applications are constructed. Computer systems have become a pillar of people's daily life. As such, it is essential for students to learn knowledge about computer systems, grasp the skills to use and design these systems, and understand the fundamental rationale and principles in computer systems. It could equip students with the necessary competence for a career related to computer science.

In the curriculum of computer science, the study of computer systems typically spans multiple knowledge areas, including, but not limited to, operating systems, parallel and distributed systems, communications networks, computer architecture and organization, and software engineering. The System Fundamentals knowledge area, as suggested by its name, focuses on the fundamental concepts and design principles in computer systems that are shared by these courses within their respective cores. The goal of this knowledge area is to present an integrative view of these fundamental concepts and design principles in a unified albeit simplified fashion, providing a common foundation for the different specialized mechanisms and policies appropriate to the particular domain area. Specifically, the fundamental concepts in this knowledge area include an overview of computer systems, basic concepts such as state and state transition, resource allocation and scheduling, and so on. Moreover, this knowledge area introduces basic design principles to improve the reliability, availability, efficiency, and security of computer systems.

Changes since CS2013

Compared to CS2013, the SF knowledge area incorporated significant changes to the knowledge units:

1. Added two new knowledge units: System Security and System Design;
2. Added a new knowledge unit named System Performance, which includes the topics from the deprecated knowledge unit of Proximity and the deprecated knowledge unit of Virtualization and Isolation;
3. Added a new knowledge unit named Performance Evaluation, which includes the topics from the deprecated unit of Evaluation and the deprecated unit of Quantitative Evaluation;
4. Renamed the Computational Paradigms knowledge unit to Overview of Computer Systems, deprecated some topics in the unit, and added topics from the deprecated unit of Cross-Layer Communications;
5. Renamed the State and State Machines knowledge unit to Basic Concepts, and added topics such as finite state machines;
7. Deprecated the Cross-Layer Communications knowledge unit and moved parts of its topics to the unit of Overview of Computer Systems;
8. Deprecated the Evaluation and Quantitative Evaluation knowledge units and moved parts of their topics to the unit of Performance Evaluation;
9. Deprecated the Proximity and Virtualization and Isolation knowledge units and moved parts of their topics to the unit of System Performance;

10. Deprecated the Parallelism knowledge unit and moved parts of its topic to the Basic Concepts knowledge unit;
11. Renamed the Reliability through Redundancy knowledge unit to System Reliability;
12. Added the Society, Ethics, and the Profession knowledge unit.

Core Hours

Knowledge Unit	CS Core	KA Core
Overview of Computer Systems	3	0
Basic Concepts	4	0
Resource Management	1	1
System Performance	2	2
Performance Evaluation	2	2
System Reliability	2	1
System Security	2	1
System Design	2	1
Society, Ethics, and Profession	Included in SEP hours	
Total	18	8

Knowledge Units

SF-Overview: Overview of Computer Systems

CS Core:

1. Basic building blocks and components of a computer (gates, flip-flops, registers, interconnections; datapath + control + memory).
2. Hardware as a computational paradigm: Fundamental logic building blocks; Logic expressions, minimization, sum of product forms. (See also: [AR-Logic](#))
3. Programming abstractions, interfaces, use of libraries. (See also: [PDC-Programs](#))
4. Distinction and interaction between application and OS services, remote procedure call. (See also: [OS-Purpose](#))
5. Basic concept of pipelining, overlapped processing stages. (See also: [AR-Organization](#))
6. Basic concept of scaling: performance vs problem size.

Illustrative Learning Outcomes:

CS Core:

1. Describe the basic building blocks of computers and their role in the historical development of computer architecture.
2. Design a simple logic circuit using the fundamental building blocks of logic design to solve a simple problem (e.g., adder).
3. Describe how computing systems are constructed of layers upon layers, based on separation of concerns, with well-defined interfaces, hiding details of low layers from the higher layers.
4. Describe that hardware, OS, VM, and application are additional layers of interpretation/processing.
5. Describe the mechanisms of how errors are detected, signaled back, and handled through the layers.
6. Construct a simple program (e.g., a TCP client/server) using methods of layering, error detection and recovery, and reflection of error status across layers.
7. Identify bugs in a layered program by using tools for program tracing, single stepping, and debugging.
8. Understand the concept of strong vs weak scaling, i.e., how performance is affected by the scale of the problem vs the scale of resources to solve the problem. This can be motivated by simple, real-world examples.

SF-Foundations: Basic Concepts**CS Core:**

1. Digital vs Analog/Discrete vs Continuous Systems.
2. Simple logic gates, logical expressions, Boolean logic simplification.
3. Clocks, State, Sequencing.
4. State and state transition (e.g., starting state, final state, life cycle of states). (See also: [AL-Models](#))
5. Finite state machines (e.g., NFA, DFA). (See also: [AL-Models](#))
6. Combinational Logic, Sequential Logic, Registers, Memories. (See also: [AR-Logic](#))
7. Computers and Network Protocols as examples of State Machines. (See also: NC-Fundamentals)
8. Sequential vs parallel processing. (See also: [PDC-Programs](#), [OS-Concurrency](#))
9. Application-level sequential processing: single thread. (See also: [PDC-Programs](#), [OS-Concurrency](#))
10. Simple application-level parallel processing: request level (web services/client-server/distributed), single thread per server, multiple threads with multiple servers, pipelining. (See also: [PDC-Programs](#), [OS-Concurrency](#))

Illustrative Learning Outcomes:**CS Core:**

1. Describe the differences between digital and analog systems, and between discrete and continuous systems. Can give real-world examples of these systems.
2. Describe computations as a system characterized by a known set of configurations with transitions from one unique configuration (state) to another (state).
3. Describe the distinction between systems whose output is only a function of their input (stateless) and those with memory/history (stateful).
4. Develop state machine descriptions for simple problem statement solutions (e.g., traffic light sequencing, pattern recognizers).
5. Describe a computer as a state machine that interprets machine instructions.

6. Explain how a program or network protocol can also be expressed as a state machine and that alternative representations for the same computation can exist.
7. Derive the time-series behavior of a state machine from its state machine representation (e.g., TCP connection management state machine).
8. Write a simple sequential problem and a simple parallel version of the same program.
9. Evaluate the performance of simple sequential and parallel versions of a program with different problem sizes, and be able to describe the speed-ups achieved.
10. Describe on an execution timeline that parallelism events and operations can take place simultaneously (i.e., at the same time). Explain how work can be performed in less elapsed time if this can be exploited.

SF-Resource: Resource Management

CS Core:

1. Different types of resources (e.g., processor share, memory, disk, net bandwidth). (See also: [OS-Scheduling](#), [OS-Memory](#), [OS-Files](#), [NC-Fundamentals](#))
2. Common resource allocation/scheduling algorithms (e.g., first-come-first-serve, priority-based scheduling, fair scheduling and preemptive scheduling). (See also: [OS-Scheduling](#))

KA Core:

3. Advantages and disadvantages of common scheduling algorithms. (See also: [OS-Scheduling](#))

Illustrative Learning Outcomes:

CS Core:

1. Define how finite computer resources (e.g., processor share, memory, storage and network bandwidth) are managed by their careful allocation to existing entities.
2. Describe how common resource allocation/scheduling algorithms work.
3. Develop common scheduling algorithms, and evaluate their performances.

KA Core:

4. Describe the pros and cons of common scheduling algorithms

SF-Performance: System Performance

CS Core:

1. Latencies in computer systems:
 - a. Speed of light and computers (one foot per nanosecond vs one GHz clocks). (See also: [AR-Organization](#))
 - b. Memory vs disk latencies vs across-the-network memory. (See also: [AR-Memory](#), [AR-Performance-Energy](#))
2. Caches and the effects of spatial and temporal locality on performance in processors and systems. (See also: [AR-Memory](#), [AR-Performance-Energy](#), [OS-Memory](#))
3. Caches and cache coherency in databases, operating systems, distributed systems, and computer architecture. (See also: [OS-Memory](#), [AR-Memory](#), [DM-Internals](#))
4. Introduction to the processor memory hierarchy. (See also: [AR-Memory](#), [AR-Performance-Energy](#))

KA Core:

5. The formula for average memory access time. (See also: [AR-Memory](#))
6. Rationale of virtualization and isolation: protection and predictable performance. (See also: [OS-Virtualization](#))
7. Levels of indirection, illustrated by virtual memory for managing physical memory resources. (See also: [OS-Virtualization](#))
8. Methods for implementing virtual memory and virtual machines. (See also: [OS-Virtualization](#))

Illustrative Learning Outcomes:**CS Core:**

1. Describe the breakdown of the latency of computer systems in terms of memory, disk and network.
2. Explain the importance of locality in determining system performance.
3. Calculate average memory access time and describe the tradeoffs in memory hierarchy performance in terms of capacity, miss/hit rate, and access time.

KA Core:

4. Explain why it is important to isolate and protect the execution of individual programs and environments that share common underlying resources.
5. Describe how the concept of indirection can create the illusion of a dedicated machine and its resources even when physically shared among multiple programs and environments.
6. Evaluate the performance of two application instances running on separate virtual machines and determine the effect of performance isolation.

SF-Evaluation: Performance Evaluation**CS Core:**

1. Performance figures of merit. (See also: [AR-Performance-Energy](#), [PDC-Evaluation](#))
2. Workloads and representative benchmarks, and methods of collecting and analyzing performance figures of merit. (See also: [AR-Performance-Energy](#), [PDC-Evaluation](#))
3. CPI (Cycles per Instruction) equation as a tool for understanding tradeoffs in the design of instruction sets, processor pipelines, and memory system organizations. (See also: [AR-Performance-Energy](#), [PDC-Evaluation](#))
4. Amdahl's Law: the part of the computation that cannot be sped up limits the effect of the parts that can. (See also: [AR-Performance-Energy](#), [PDC-Evaluation](#))
5. Order of magnitude analysis. (Big O notation) (See also: [AL-Complexity](#))
6. Analysis of slow and fast paths of a system. (See also: [AR-Organization](#))
7. Events on their effect on performance (e.g., instruction stalls, cache misses, page faults). (See also: [OS-Memory](#), [AR-Organization](#))

KA Core:

7. Analytical tools to guide quantitative evaluation
8. Understanding layered systems, workloads, and platforms, their implications for performance, and the challenges they represent for evaluation
9. Microbenchmark pitfalls

Illustrative Learning Outcomes:

CS Core:

1. Explain how the components of system architecture contribute to improving its performance.
2. Explain the circumstances in which a given figure of a system performance metric is useful.
3. Explain the usage and inadequacies of benchmarks as a measure of system performance.
4. Describe Amdahl's law and discuss its limitations.
5. Apply limit studies or simple calculations to produce order-of-magnitude estimates for a given performance metric in a given context.
6. Apply software tools to profile and measure program performance.

KA Core:

7. Design and conduct a performance-oriented experiment of a common system (e.g., an OS and Spark).
8. Design a performance experiment on a layered system to determine the effect of a system parameter on system performance.

SF-Reliability: System Reliability

CS Core:

1. Distinction between bugs, faults and failures. (See also: [PDC-Coordination](#), [SE-Reliability](#))
2. Reliability vs availability.
3. Reliability through redundancy.
 - a. check and retry. (See also: [OS-Faults](#), [NC-Reliability](#))
 - b. redundant encoding (error correction codes, CRC, FEC, RAID) (See also: [AR-Memory](#), [NC-Reliability](#), [DM-Distributed](#))
 - c. duplication/mirroring/replicas. (See also: [DM-Distributed](#))

KA Core:

4. Other approaches to reliability (e.g., journaling). (See also: [OS-Faults](#), [NC-Reliability](#), [SE-Reliability](#))

Illustrative Learning Outcomes:

CS Core:

1. Explain the distinction between program errors, system errors, and hardware faults (e.g., corrupted memory) and exceptions (e.g., attempt to divide by zero).
2. Articulate the distinction between detecting, handling, and recovering from faults and the methods for their implementation.
3. Describe the role of error correction codes in providing error checking and correction techniques in memories, storage, and networks.
4. Apply simple algorithms for exploiting redundant information for the purposes of data correction.

KA Core:

5. Compare different error detection and correction methods for their data overhead, implementation complexity, and relative execution time for encoding, detecting, and correcting errors.

SF-Security: System Security

CS Core:

1. Common system security issues (e.g., viruses, denial-of-service attacks and eavesdropping). (See also: [OS-Protection](#), [NC-Security](#), [SEC-Foundations](#), [SEC-Engineering](#))
2. Countermeasures (See also: [OS-Principles](#), [OS-Protection](#), [NC-Security](#))
 - a. Cryptography (See also: [SEC-Crypto](#))
 - b. Security architecture (See also: [SEC-Engineering](#))

KA Core:

3. Representative countermeasure systems
 - a. Intrusion detection systems, firewalls. (See also: [NC-Security](#))
 - b. Antivirus systems

Illustrative Learning Outcomes:

CS Core:

1. Describe some common system security issues and give examples.
2. Describe some countermeasures against system security issues.

KA Core:

3. Describe representative countermeasure systems.

SF-Design: System Design

CS Core:

1. Common criteria of system design (e.g., liveness, safety, robustness, scalability and security) .(See also: [PDC-Evaluation](#))

KA Core:

2. Designs of representative systems (e.g., Apache web server, Spark and Linux).

Illustrative Learning Outcomes:

CS Core:

1. Describe common criteria of system design.
2. Given the functionality requirements of a system and its key design criteria, provide a high-level design of this system.

KA Core:

3. Describe the design of some representative systems.

SF-SEP: Society, Ethics, and the Profession

KA Core:

1. Intellectual property rights of computer systems. (See also: [SEP-IP](#))
2. Common software licenses. (See also: [SEP-IP](#))
3. Computer crimes. (See also: [SEP-Security](#))

Illustrative Learning Outcomes:

KA Core:

1. Describe the intellectual property rights of computer systems.
2. List representative software licenses and compare their differences.
3. List representative computer crimes.

Professional Dispositions

- **Meticulous:** Students must pay attention to details of different perspectives when learning about and evaluating systems.
- **Adaptive:** Students must be flexible and adaptive when designing systems. Different systems have different requirements, constraints and working scenarios. As such, they require different designs. Students must be able to make appropriate design decisions correspondingly.

Mathematics Requirements

Required:

- Discrete Mathematics (See also: [MSF-Discrete](#)):
 - Sets and relations
 - Basic graph theory
 - Basic logic
- Linear Algebra (See also: [MSF-Linear](#)):
 - Basic matrix operations
- Probability and Statistics (See also: [MSF-Probability](#), [MSF-Statistics](#))
 - Random variable
 - Bayes theorem
 - Expectation and Variation
 - Cumulative distribution function and probability density function

Desirable:

- Basic queueing theory
- Basic stochastic process

Course Packaging Suggestions

Introductory Course to include the following:

- [SF-Overview](#) (2 hours)
- [SF-Foundations](#) (6 hours)
- [SF-Resource](#) (4 hours)
- [SF-Performance](#) (6 hours)
- [SF-Evaluation](#) (6 hours)

- [SF-Reliability](#) (4 hours)
- [SF-Security](#) (5 hours)
- [SF-SEP](#) (1 hour)
- [SF-Design](#) (6 hours)

Prerequisites:

- Sets and relations, basic graph theory and basic logic from Discrete Mathematics (See also: [MSF-Discrete](#))
- Basic matrix operations from Linear Algebra (See also: [MSF-Linear](#))
- Random variable, Bayes theorem, expectation and variation, cumulative distribution function and probability density function from Probability and Statistics (See also: [MSF-Probability](#), [MSF-Statistics](#))

Course objectives: Students should be able to (1) understand the fundamental concepts in computer systems; (2) understand the key design principles, in terms of performance, reliability and security, when designing computer systems; (3) deploy and evaluate representative complex systems (e.g., MySQL and Spark) based on their documentations, and (4) design and implement simple computer systems (e.g., an interactive program, a simple web server, and a simple data storage system).

Advanced Course to include the following:

- [SF-Overview](#): Overview of Computer Systems (2 hours)
- [SF-Design](#): System Design (8 hours)
- [OS-Purpose](#), [OS-Principles](#) (2 hours)
- [NC-Fundamentals](#), [NC-Networked-Applications](#) (2 hours)
- [PDC-Programs](#) (2 hours)
- [AR-IO](#), [AR-Performance-Energy](#) (2 hours)
- [SF-Reliability](#) (8 hours)
- [SF-Performance](#) (6 hours)
- [SF-Security](#) (6 hours)
- [SF-SEP](#): (2 hours)

Prerequisites:

- Basic queueing theory and stochastic process (See also: [MSF-Probability](#), [MSF-Statistics](#))
- Introductory Course of the [SF](#) KA

Course objectives: Students should be able to (1) have a deeper understanding in the key design principles of computer system design, (2) map such key principles to the designs of classic systems (e.g., Linux, SQL and TCP/IP network stack) as well as that of more recent systems (e.g., Hadoop, Spark and distributed storage systems), and (3) design and implement more complex computer systems (e.g., a file system and a high-performance web server).

Committee

Chair: Qiao Xiang, Xiamen University, Xiamen, China

Members:

- Doug Lea, State University of New York at Oswego, Oswego, NY, USA
- Monica D. Anderson, University of Alabama, Tuscaloosa, AL, USA
- Matthias Hauswirth, University of Lugano, Lugano, Switzerland
- Ennan Zhai, Alibaba Group, Hangzhou, China
- Yutong Liu, Shanghai JiaoTong University, Shanghai, China

Contributors:

- Michael S. Kirkpatrick, James Madison University, Harrisonburg, VA, USA
- Linghe Kong, Shanghai JiaoTong University, Shanghai, China

Specialized Platform Development (SPD)

Preamble

The Specialized Platform Development (SPD) Knowledge Area (KA) refers to attributes involving the creation of software targeting non-traditional hardware platforms. Developing for each specialized platform, for example, robots, mobile systems, web-based systems, and embedded systems typically involve unique considerations.

Societal and industry needs have created a high demand for developers on specialized platforms, such as mobile applications, web platforms, robotic platforms, and embedded systems. Some unique professional abilities relevant to this KA include:

- Creating applications that provide a consistent user experience across various devices, screen sizes, and operating systems.
- Developing application programming interfaces (APIs) to support the functionality of each specialized platform.
- Managing challenges related to resource constraints such as computation, memory, storage, and networking and communication.
- Applying cross-cutting concerns such as optimization, security, better development practices, and others.

Changes since CS2013

The assessment of these factors has led to significant modifications from the CS2013 version, including:

Knowledge Area Name Revision: The knowledge area name has been changed to reflect the specialized development platforms which serve as the target for software development.

Expansion of Computer Science Core Hours: Reflecting the increased deployment of specialized hardware platforms, the number of CS Core hours has been increased.

Revised and Introduced Knowledge Units: Reflecting modern computing systems, knowledge units in Robotics, Embedded Systems, and Society, Ethics, and the Profession (SEP) have been introduced. Other changes include: 1) renamed Introduction knowledge unit to Common Aspects/Shared Concerns and 2) renamed Industrial Platforms to Robot Platforms.

Core Hours

Knowledge Units	CS Core	KA Core
Common Aspects	3 + (1 SE)	2

Web Platforms		5 + (1 HCI)
Mobile Platforms		3 + (2 DM) + (1 GIT) + (1 HCI) + (1 NC)
Robot Platforms		4 + (3 GIT) + (3 AI)
Embedded Platforms		4 + (4 AR) + (1 FPL) + (1 GIT) + (3 OS) + (1 SF)
Game Platforms		4 + (1 AL) + (1 AI) + (4 AR) + (5 GIT) + (2 HCI) + (1 SDF) + (1 SE) + (1 MSF)
Interactive Computing Platforms		3 + (1 DM) + (2 GIT) + (1 AR) + (1 FPL)
SPD-SEP	Included in SEP Hours	
Total	4	68

Note: The CS Core hours total includes 1 hour shared with SE.

Knowledge Units

SPD-Common: Common Aspects/Shared Concerns

CS Core:

1. Overview of development platforms (i.e., web, mobile, game, robotics, embedded, and Interactive).
 - a. Input/sensors/control devices/haptic devices.
 - b. Resource constraints:
 - i. Computational
 - ii. Data storage
 - iii. Memory
 - iv. Communication
 - c. Requirements - security, uptime availability, fault tolerance. (See also: [SE-Reliability](#), [SEC-Engineering](#))
 - d. Output/actuators/haptic devices.
2. Programming via platform-specific Application Programming Interface (API) vs traditional application construction.
3. Overview of platform Languages (e.g., Python, Swift, Lua, Kotlin).
4. Programming under platform constraints and requirements (e.g., available development tools, development, security considerations). (See also: [SEC-Foundations](#))
5. Techniques for learning and mastering a platform-specific programming language.

Illustrative Learning Outcomes:

CS Core:

1. List the constraints of mobile programming.
2. List the characteristics of scripting languages.
3. Describe the three-tier model of web programming.
4. Describe how the state is maintained in web programming.

SPD-Web: Web Platforms

KA Core:

1. Web programming languages (e.g., HTML5, JavaScript, PHP, CSS)
2. Web platforms, frameworks, or meta-frameworks:
 - a. Cloud services
 - b. API, Web Components
3. Software as a Service (SaaS).
4. Web standards such as document object model, accessibility. (See also: [HCI-Accessibility](#))
5. Security and Privacy Considerations. (See also: [SEP-Security](#))

Non-core:

6. Analyzing requirements for web applications.
7. Computing services (See also: [DM-NoSQL](#))
 - a. Cloud Hosting.
 - b. Scalability (e.g., Autoscaling, Clusters).
 - c. Cost estimation for services.
8. Data management: (See also: [DM-Core](#))
 - a. Data residency: where the data is located and what paths can be taken to access it.
 - b. Data integrity: guaranteeing data is accessible and that data is deleted when required.
9. Architecture
 - a. Monoliths vs Microservices.
 - b. Micro-frontends.
 - c. Event-Driven vs RESTful architectures: advantages and disadvantages.
 - d. Serverless, cloud computing on demand.
10. Storage solutions: (See also: [DM-Relational](#), [DM-NoSQL](#))
 - a. Relational Databases.
 - b. NoSQL databases.

Illustrative Learning Outcomes:

KA Core:

1. Design and Implement a web-based application using a microservice architecture design.
2. Describe the constraints, such as hosting, services, and scalability, related to web platforms.
3. Compare and contrast web programming with general-purpose programming.
4. Describe the differences between Software-as-a-Service (SaaS) and traditional software products.
5. Discuss how web standards impact software development.
6. Evaluate an existing web application against current web standards.

SPD-Mobile: Mobile Platforms

KA Core:

1. Development with:
 - a. Mobile programming languages.
 - b. Mobile programming environments.
2. Mobile platform constraints:
 - a. User interface design. (See also: [HCI-User](#))
 - b. Security.
3. Access:
 - a. Accessing data through APIs. (See also: [DM-Querying](#))
 - b. Designing API endpoints for mobile apps: pitfalls and design considerations.
 - c. Network and the web interfaces. (See also: [NC-Fundamentals](#), [DM-Modeling](#))

Non-core:

4. Development:
 - a. Native versus cross-platform development.
 - b. Software design/architecture patterns for mobile applications. (See also: [SE-Design](#))
5. Mobile platform constraints:
 - a. Responsive user interface design. (See also: [HCI-Accessibility](#))
 - b. Heterogeneity and mobility of devices.
 - c. Differences in user experiences (e.g., between mobile and web-based applications).
 - d. Power and performance tradeoff.
6. Mobile computing affordances:
 - a. Location-aware applications.
 - b. Sensor-driven computing (e.g., gyroscope, accelerometer, health data from a watch).
 - c. Telephony and instant messaging.
 - d. Augmented reality. (See also: [GIT-Immersion](#))
7. Specification and testing. (See also: [SDF-Practices](#), [SE-Validation](#))
8. Asynchronous computing: (See also: [PDC-Algorithms](#))
 - a. Difference from traditional synchronous programming.
 - b. Handling success via callbacks.
 - c. Handling errors asynchronously.
 - d. Testing asynchronous code and typical problems in testing.

Illustrative Learning Outcomes:

KA Core:

1. Compare mobile programming with general-purpose programming.
2. Develop a location-aware mobile application with data API integration.
3. Build a sensor-driven mobile application capable of logging data on a remote server.
4. Create a communication app incorporating telephony and instant messaging.
5. Evaluate the pros and cons of native and cross-platform mobile application development.

SPD-Robot: Robot Platforms

KA Core:

1. Types of robotic platforms and devices. (See also: [AI-Robotics](#))
2. Sensors, embedded computation, and effectors (actuators). (See also: [GIT-Physical](#))
3. Robot-specific languages and libraries. (See also: [AI-Robotics](#))
4. Robotic software architecture (e.g., using the Robot Operating System (ROS)).
5. Robotic platform constraints and design considerations. (See also: [AI-Robotics](#))
6. Interconnections with physical or simulated systems. (See also: [GIT-Physical](#), [GIT-Simulation](#))
7. Robotic Algorithms (See also: [AI-Robotics](#), [GIT-Animation](#))
 - a. Forward kinematics.
 - b. Inverse kinematics.
 - c. Dynamics.
 - d. Navigation and path planning.
 - e. Grasping and manipulation.
8. Safety and interaction considerations. (See also: [SEP-Professional-Ethics](#), [SEP-Context](#))

Illustrative Learning Outcomes:

KA Core:

1. Design and implement an application on a given robotic platform.
2. Integrate an Arduino-based robot kit and program it to navigate a maze.
3. Compare robot-specific languages and techniques with those used for general-purpose software development.
4. Explain the rationale behind the design of the robotic platform and its interconnections with physical or simulated systems.
5. Given a high-level application, design a robot software architecture using ROS specifying all components and interconnections (ROS topics) to accomplish that application.
6. Discuss the constraints a given robotic platform imposes on developers.

SPD-Embedded: Embedded Platforms

KA Core:

1. Introduction to the unique characteristics of embedded systems:
 - a. Real-time vs soft real-time and non-real-time systems.
 - b. Resource constraints, such as memory profiles and deadlines. (See also: [AR-Memory](#))
2. API for custom architectures:
 - a. GPU technology. (See also: [AR-Heterogeneity](#), [GIT-Shading](#))
 - b. Field Programmable Gate Arrays (FPGA). (See also: [AR-Logic](#))
 - c. Cross-platform systems.
3. Embedded Systems:
 - a. Microcontrollers.
 - b. Interrupts and feedback.
 - c. Interrupt handlers in high-level languages. (See also: [SF-Overview](#))
 - d. Hard and soft interrupts and trap-exits. (See also: [OS-Principles](#))
 - e. Interacting with hardware, actuators, and sensors.

- f. Energy efficiency.
 - g. Loosely timed coding and synchronization.
 - h. Software adapters.
4. Embedded programming.
 5. Hard real-time systems vs soft real-time systems: (See also: [OS-Real-time](#))
 - a. Timeliness.
 - b. Time synchronization/scheduling.
 - c. Prioritization.
 - d. Latency.
 - e. Compute jitter.
 6. Real-time resource management.
 7. Memory management:
 - a. Mapping programming construct (variable) to a memory location. (See also: [AR-Memory](#))
 - b. Shared memory. (See also: [OS-Memory](#))
 - c. Manual memory management.
 - d. Garbage collection. (See also: [FPL-Translation](#))
 8. Safety considerations and safety analysis. (See also: [SEP-Context](#), [SEP-Professional-Ethics](#))
 9. Sensors and actuators.
 10. Analysis and verification.
 11. Application design.

Illustrative Learning Outcomes:

KA Core:

1. Design and implement a small embedded system for a given platform (e.g., a smart alarm clock or a drone).
2. Describe the unique characteristics of embedded systems versus other systems.
3. Interface with sensors/actuators.
4. Debug a problem with an existing embedded platform.
5. Identify different types of embedded architectures.
6. Evaluate which architecture is best for a given set of requirements.
7. Design and develop software to interact with and control hardware.
8. Design methods for real-time systems.
9. Evaluate real-time scheduling and schedulability analysis.
10. Evaluate formal specification and verification of timing constraints and properties.

SPD-Game: Game Platforms

KA Core:

1. Historical and contemporary platforms for games: (See also: [AR-Logic](#))
 - a. Evolution of Game Platforms (e.g., Brown Box to Metaverse and beyond; Improvement in Computing Architectures (CPU and GPU); Platform Convergence and Mobility).
 - b. Typical Game Platforms (e.g., Personal Computer; Home Console; Handheld Console; Arcade Machine; Interactive Television; Mobile Phone; Tablet; Integrated Head-Mounted Display; Immersive Installations and Simulators; Internet of Things enabled Devices; CAVE Systems; Web Browsers; Cloud-based Streaming Systems).

- c. Characteristics and Constraints of Different Game Platforms (e.g., Features (local storage, internetworking, peripherals); Run-time performance (GPU/CPU frequency, number of cores); Chipsets (physics processing units, vector co-processors); Expansion Bandwidth (PCIe); Network throughput (Ethernet); Memory types and capacities (DDR/GDDR); Maximum stack depth; Power consumption; Thermal design; Endian)
 - d. Typical Sensors, Controllers, and Actuators (e.g., distinctive control system designs—peripherals (mouse, keypad, joystick), game controllers, wearables, interactive surfaces; electronics and bespoke hardware; computer vision, inside-out tracking, and outside-in tracking; IoT-enabled electronics and I/O. (See also: [GIT-Interaction](#))
 - e. eSports Ecosystems (e.g., evolution of gameplay across platforms; games and eSports; game events such as LAN/arcade tournaments and international events such as the Olympic eSports Series; streamed media and spectatorship; multimedia technologies and broadcast management; professional play; data and machine learning for coaching and training)
2. Real-time Simulation and Rendering Systems
- a. CPU and GPU architectures: (e.g., Flynn’s taxonomy; parallelization; instruction sets; standard components—graphics compute array, graphics memory controller, video graphics array basic input/output system; bus interface; power management unit; video processing unit; display interface). (See also: [AR-Heterogeneity](#))
 - b. Pipelines for physical simulations and graphical rendering: (e.g., tile-based, immediate-mode). (See also: [GIT-Rendering](#))
 - c. Common Contexts for Algorithms, Data Structures, and Mathematical Functions (e.g., game loops; spatial partitioning, viewport culling, and level of detail; collision detection and resolution; physical simulation; behavior for intelligent agents; procedural content generation). (See also: [MSF-Discrete](#), [AL-Foundational](#))
 - d. Media representations (e.g., I/O, and computation techniques for virtual worlds: audio; music; sprites; models and textures; text; dialogue; multimedia (e.g., olfaction, tactile). (See also: [GIT-Fundamentals](#))
3. Game Development Tools and Techniques:
- a. Programming Languages (e.g., C++; C#; Lua; Python; JavaScript).
 - b. Shader Languages (e.g., HLSL, GLSL; ShaderGraph).
 - c. Graphics Libraries and APIs (e.g., DirectX; SDL; OpenGL; Metal; Vulkan; WebGL). (See also: [GIT-Rendering](#), [HCI-Design](#))
 - d. Common Development Tools and Environments (e.g., IDEs; Debuggers; Profilers; Version Control Systems including those handling binary assets; Development Kits and Production/Consumer Kits; Emulators). (See also: [SDF-Practices](#), [SE-Tools](#))
4. Game Engines
- a. Open Game Engines (e.g., Unreal; Unity; Godot; CryEngine; Phyre; Source 2; Pygame and Ren’Py; Phaser; Twine; SpringRTS)
 - b. Techniques (e.g., Ideation, Prototyping, Iterative Design and Implementation, Compiling Executable Builds, Development Operations and Quality Assurance—Play Testing and Technical Testing, Profiling; Optimization, Porting; Internationalization and Localization, Networking). (See also: [AR-Performance-Energy](#), [SE-Requirements](#))
5. Game Design

- a. Vocabulary (e.g., game definitions; mechanics-dynamics-aesthetics model; industry terminology; experience design; models of experience and emotion).
- b. Design Thinking and User-Centered Experience Design (e.g., methods of designing games; iteration, incrementing, and the double-diamond; phases of pre- and post-production; quality assurance, including alpha and beta testing; stakeholder and customer involvement; community management). (See also: [SE-Design](#))
- c. Genres (e.g., adventure; walking simulator; first-person shooter; real-time strategy; multiplayer online battle arena (MOBA); role-playing game (rpg)).
- d. Audiences and Player Taxonomies (e.g., people who play games; diversity and broadening participation; pleasures, player types, and preferences; Bartle, yee). (See also: [HCI-User](#))
- e. Proliferation of digital game technologies to domains beyond entertainment (e.g., Education and Training; Serious Games; Virtual Production; eSports; Gamification; Immersive Experience Design; Creative Industry Practice; Artistic Practice; Procedural Rhetoric). (See also: [AI-SEP](#))

Illustrative Learning Outcomes:

KA Core:

1. Recall the characteristics of common general-purpose graphics processing architectures.
2. Identify the key stages of the immediate-mode rendering pipeline.
3. Describe the key constraints a given game platform will likely impose on developers.
4. Explain how eSports are streamed to large audiences over the Internet.
5. Translate complex mathematical functions into performant source code.
6. Use an industry-standard graphics API to render a 3D model in a virtual scene.
7. Modify a shader to change a visual effect according to stated requirements.
8. Implement a game for a particular platform according to the specification.
9. Optimize a function for processing collision detection in a simulated environment.
10. Assess a game's run-time and memory performance using an industry-standard tool and development environment.
11. Compare the interfaces of different game platforms, highlighting their respective implications for human-computer interaction.
12. Recommend an appropriate set of development tools and techniques for implementing a game of a particular genre for a given platform.
13. Discuss the key challenges in making a digital game that is cross-platform compatible.
14. Express how game developers can enhance the accessibility of a game interface.
15. Create novel forms of gameplay using frontier game platforms.

SPD-Interactive: Interactive Computing Platforms

Non-core:

1. Data Analysis Platforms:
 - a. Jupyter notebooks; Google Colab; R; SPSS; Observable.
 - b. Cloud SQL/data analysis platforms (e.g., BigQuery). (See also: [DM-Querying](#))
 - i. Apache Spark.
 - ii. Data Visualizations. (See also: [GIT-Visualization](#))
 - c. Interactive presentations backed by data.

- d. Design tools requiring low-latency feedback loops:
 - i. Rendering tools.
 - ii. Graphic design tools.
- 2. Prompt programming:
 - a. Generative AI (e.g., OpenAI’s ChatGPT, OpenAI’s Codex, GitHub’s Copilot) and LLMs are accessed/interacted.
- 3. Quantum Platforms: (See also: [AR-Quantum](#))
 - a. Program quantum logic operators in quantum machines.
 - b. Use API for available quantum services.
 - c. Signal analysis/Fourier analysis/Signal processing (for music composition, audio/RF analysis). (See also: [GIT-Image](#))

Illustrative Learning Outcomes:

Non-core:

1. Analyze large datasets interactively.
2. Create a backing track for a musical performance, such as live coding.
3. Create compelling computational notebooks that construct a narrative for a given journalistic goal/story.
4. Implement interactive code that uses a dataset and generates exploratory graphics.
5. Create a program that performs a task using LLM systems.
6. Contrast a program developed by an AI platform and by a human.
7. Implement a system that interacts with a human without using a screen.
8. Contextualize the attributes of different data analysis styles, such as interactive vs engineered pipeline.
9. Write a program using a notebook computing platform (e.g., searching, sorting, or graph manipulation).
10. Demonstrate a quantum gate outcome using a quantum platform.

SPD-SEP/Mobile

Non-core:

1. Privacy and data protection.
2. Accessibility in mobile design.
3. Security and cybersecurity.
4. Social impacts of mobile technology,
5. Ethical use of AI and algorithms.

Illustrative Learning Outcomes:

Non-core:

1. Understand and uphold ethical responsibilities for safeguarding user privacy and data protection in mobile applications.
2. Design mobile applications with accessibility in mind, ensuring effective use by people with disabilities.
3. Demonstrate proficiency in secure coding practices to mitigate risks associated with various security threats in mobile development.

4. Analyze the broader social impacts of mobile technology, including its influence on communication patterns, relationships, and mental health.
5. Comprehend the ethical considerations of using AI in mobile applications, ensuring unbiased and fair algorithms.

SPD-SEP/Web

Non-core:

1. Privacy concerns in mobile apps.
2. Designing for inclusivity and accessibility.
3. Ethical use of AI in mobile apps.
4. Sustainable app development and server hosting.
5. Avoiding spam or intrusive notifications.
6. Addressing cyberbullying and harassment.
7. Promoting positive online communities.
8. Monetization and advertising.
9. Ethical use of gamification.

Illustrative Learning Outcomes:

Non-core:

1. Understand how mobile computing impacts communications and the flow of information within society.
2. Design mobile apps that have made daily tasks easier/faster.
3. Recognize how the ubiquity of mobile computing has affected work-life balance.
4. Understand how mobile computing impacts health monitoring and healthcare services.
5. Define how mobile apps are used to educate on and help achieve UN sustainability goals.

SPD-SEP/Game

Non-core:

1. Intellectual Property Rights in Creative Industries:
 - a. Intellectual Property Ownership: copyright, trademark; design right, patent, trade secret, civil versus criminal law; international agreements; procedural content generation and the implications of generative artificial intelligence.
 - b. Licensing: Usage and fair usage exceptions; open-source license agreements; proprietary and bespoke licensing; enforcement.
2. Fair Access to Play:
 - a. Game Interface Usability: user requirements, affordances, ergonomic design, user research, experience measurement, and heuristic evaluation methods for games.
 - b. Game Interface Accessibility: forms of impairment and disability; means to facilitate game access; universal design; legislated requirements for game platforms; compliance evaluation; challenging game mechanics and access.
3. Game-Related Health and Safety:
 - a. Injuries in Play: ways of mitigating common upper body injuries, such as repetitive strain injury; exercise psychology and physiotherapy in eSports.

- b. Risk Assessment for Events and Manufacturing: control of substances hazardous to health (COSHH); fire safety; electrical and electronics safety; risk assessment for games and game events; risk assessment for manufacturing.
 - c. Mental Health: motivation to play; gamification and gameful design; game psychology—internet gaming disorder.
- 4. Platform Hardware Supply Chain and Sustainability:
 - a. Platform Lifecycle: platform composition—materials, assembly; mineral excavation and processing; power usage; recycling; planned obsolescence.
 - b. Modern Slavery: supply chains; forced labor and civil rights; working conditions; detection and remission; certification bodies and charitable endeavors.
- 5. Representation in the Media and Industry:
 - a. Inclusion: identity and identification; Inclusion of a broad range of characters for diverse audiences; media representation and its effects; media literacy; content analysis; stereotyping; sexualization.
 - b. Equality: histories and controversies, such as gamergate, quality of life in the industry, professional discourse and conduct in business contexts, pathways to game development careers, social mobility, the experience of developers from different backgrounds and identities, gender, and technology.

Illustrative Learning Outcomes:

Non-core:

1. Recall how creators can protect their intellectual property.
2. Identify common pitfalls in game interfaces that exclude players with impaired or non-functional vision.
3. Describe how heuristic evaluation can be used to identify usability problems in game interfaces.
4. Explain why upper body injuries are common in eSports.
5. Discuss how to reform characters and dialogues in a scene to reduce stereotype threat.
6. Illustrate how the portrayal of race in a game can influence the risk of social exclusion in the associated online community around the game.
7. Modify a policy for a LAN party event to include mitigations that lower the risk of fire.
8. Design a gamification strategy to motivate serious play for an awareness-raising game.
9. Analyze the role of company hiring policies and advocacy on social mobility.
10. Assess the appropriateness of two manufacturers for producing a new game console.
11. Compare options for open-source licensing of a game development tool.
12. Recommend changes to a game interface to improve access to players who are deaf or whose hearing is otherwise impaired.
13. Discuss whether games are addictive.
14. Suggest how the portrayal of women in video games influences how players perceive members of those groups.
15. Create a video game that successfully advocates for climate science.

SPD-SEP/Robotics

Non-core:

1. Fairness, transparency, and accountability in robotic algorithms.

2. Mitigating biases in robot decision-making.
3. Public safety in shared spaces with robots.
4. Compliance with data protection laws.
5. Patient consent and trust in medical robots.

Illustrative Learning Outcomes:

Non-core:

1. Identify instances of bias in robotic algorithms and propose strategies to mitigate them.
2. Evaluate and critique robotic systems for ethical and fairness considerations, suggesting improvements where necessary.
3. Analyze real-world examples of biases in robot decision-making and develop strategies to reduce bias in robotic systems.
4. Assess the potential risks associated with robots in public spaces and propose safety measures to mitigate those risks.
5. Evaluate the impact of patient consent and trust on the effectiveness of medical robot deployments in healthcare contexts.

SPD-SEP/Interactive

Non-core:

1. Ethical guidelines when using AI models to assist in journalism and content creation.
2. Accountability for AI-generated outputs.
3. Behavior among prompt programmers and AI developers.
4. Trust with the public when using AI models.

Illustrative Learning Outcomes:

Non-core:

1. Indicate a framework for accountability in AI model deployment, including clear documentation and attribution.
2. Discuss ethical codes of conduct and professional standards relevant to prompt programming and AI development.
3. Create communication plans and materials to educate the public about AI capabilities, limitations, and ethical safeguards.

Professional Dispositions

- **Self-Directed:** Students should be able to learn new platforms and languages with a growth-oriented mindset and thrive in dynamic environments, while continually enhancing skills.
- **Inventive:** Students should demonstrate excellence in designing software architecture within unconventional constraints, emphasizing adaptability and creative problem-solving for innovative solutions.
- **Adaptable:** Students should adapt to diverse challenges, showing resilience, open-mindedness, and a proactive approach to changing requirements and constraints.

Mathematics Requirements

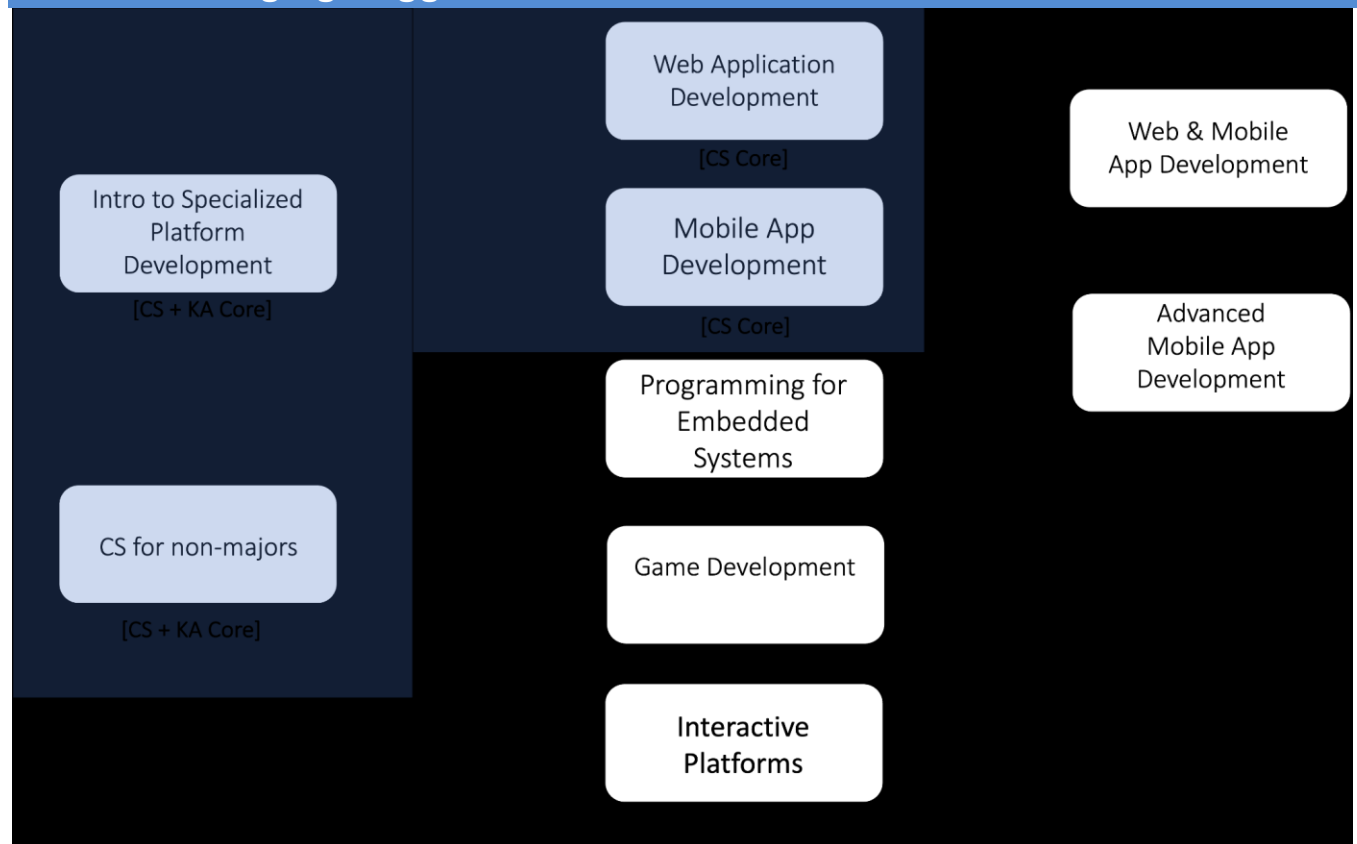
Required:

- MSF-Discrete

Desired:

- MSF-Calculus
- MSF-Linear
- MSF-Statistics

Course Packaging Suggestions



Introduction to Special Platform Development Course to include the following:

- [SPD-Common](#) (10 hours)
- [SPD-Web](#) (4 hours)
- [SPD-Mobile](#) (4 hours)
- [SPD-SEP/Robotics](#) (3 hours)
- [SPD-Embedded](#) (3 hours)
- [SPD-Game](#) (4 hours)
- [SPD-SEP/Interactive](#) (2 hours)
- [SEP-Context](#) (2 hours)
- [SDF-Practices](#) (4 hours)
- [SE-Design](#) (2 hours)

- [FPL-Scripting](#) (2 hours)

Course objectives: Students should be able to grasp common aspects of platform development, acquire foundational knowledge in web development, and attain proficiency in web techniques. They will apply comprehensive mobile development skills and explore challenges in robotics platforms. Expertise in developing platforms for embedded systems, along with skills in game development and creating interactive platforms, will be developed. Students will analyze societal, ethical, and professional implications of platform development, fostering a well-rounded understanding of this field within a concise curriculum.

Mobile Development Course to include the following:

- [SPD-Common](#): Common Aspects (3 hours)
- [SPD-Mobile](#) (25 hours)
- [SDF-Practices](#): Software Development Practices (2 hours)
- [SE-Design](#): Software Design (3 hours)
- [SE-Construction](#): Software Construction (2 hours)
- [SPD-SEP/Mobile](#) (3 hours)

Course objectives: Students should be able to design, develop, and deploy cross-platform mobile applications using languages like Java, Kotlin, Swift, or React Native. Proficiency in implementing user experience best practices, exploring cross-platform development tools, and utilizing platform-specific APIs for seamless integration is emphasized. The course covers security vulnerability identification, testing methodologies, and distribution/versioning of mobile applications. Students gain insights into user behavior and application performance through analytics tools. Additionally, they learn version control, release management, and ethical considerations relevant to mobile development, providing a well-rounded skill set for successful and responsible mobile application development across diverse platforms.

Web Development Course to include the following:

- [SPD-Web](#) (19 hours)
- [FPL-OOP](#) (3 hours)
- [SE-Construction](#) (2 hours)
- [DM-Querying](#) (2 hours)
- [SE-Tools](#) (4 hours)
- [SDF-Practices](#) (4 hours)
- [SE-Design](#) (2 hours)
- [FPL-Scripting](#) (2 hours)
- [SPD-SEP/Web](#) (2 hours)

Course objectives: Students should be able to gain expertise in designing, developing, and deploying modern web applications. The curriculum covers key concepts, ensuring proficiency in HTML, CSS, and JavaScript for responsive and visually appealing pages. Students explore and implement frontend frameworks (e.g., React, Angular) for efficient development, understand server-side languages (e.g., Node.js, Python) for dynamic applications, and design effective architectures prioritizing scalability and

security. They learn version control (e.g., [GIT](#)), integrate APIs for enhanced functionality, implement responsive design, optimize for performance, and ensure security through best practices. Testing, debugging, accessibility, deployment, and staying current with industry trends are also emphasized.

Game Development Course to include the following:

- [SPD-Game](#) (16 hours)
- [SPD-SEP/Game](#) (4 hours)
- [SDF-Practices](#) (4 hours)
- [GIT-Interaction](#) (1 hour)
- [HCI-Design](#) (3 hours)
- [HCI-User](#) (1 hour)
- [SE-Tools](#) (1 hour)
- [AL-Foundational](#) (2 hours)
- [GIT-Rendering](#) (4 hours)
- [SE-Design](#) (4 hours)

Course objectives: Students should be able to master designing, developing, and deploying interactive games. The curriculum covers fundamental game design principles, proficiency in languages like C++, C#, or Python, and utilization of popular engines such as Unity or Unreal. Students gain 3D modeling and animation skills, implement physics and simulations for realism, and create AI algorithms for intelligent non-player characters. They design multiplatform games, optimize UI/UX for engagement, apply game-specific testing and debugging techniques, integrate audio effectively, and explore industry monetization models. The course emphasizes ethical considerations, ensuring students analyze and address content, diversity, and inclusivity in game development.

Committee

Chair: Christian Servin, El Paso Community College, El Paso, TX, USA

Members:

- Sherif G. Aly, The American University in Cairo, Cairo, Egypt
- Yoonsik Cheon, The University of Texas at El Paso, El Paso, TX, USA
- Eric Eaton, University of Pennsylvania, Philadelphia, PA, USA
- Claudia L. Guevara, Jochen Schweizer mydays Holding GmbH, Munich, Germany
- Larry Heimann, Carnegie Mellon University, Pittsburgh, PA, USA
- Amruth N. Kumar, Ramapo College of New Jersey, Mahwah, NJ, USA
- R. Tyler Pirtle, Google, USA
- Michael James Scott, Falmouth University, Falmouth, Cornwall, UK

Contributors:

- Sean R. Piotrowski, Rider University, Lawrenceville, NJ, USA
- Mark O'Neil, Blackboard Inc., Newport, NH, USA
- John DiGennaro, Qwickly, Cleveland, OH, USA
- Rory K. Summerley, London South Bank University, London, England, UK

Core Topics Table

In the following 17 tables, CS and KA core topics have been listed, one table per knowledge area. For each topic, desired skill levels have been identified and used to estimate the time needed for the instruction of CS Core and KA Core topics. The skill levels should be treated as recommended, not prescriptive. The time needed to cover CS Core and KA Core topics is expressed in terms of **instructional hours**. Instructional hours are hours spent in the classroom imparting knowledge regardless of the pedagogy used. Students are expected to spend additional time after class practicing related skills and exercising professional dispositions.

For convenience, the tables have been listed under three competency areas: Software, Systems and Applications. The tables on Society, Ethics and the Profession (SEP) and Mathematical and Statistical Foundations (MSF) have been listed last as crosscutting topics that apply to all the competency areas.

Software Competency Area

The core topics in Software Development Fundamentals (SDF) and Algorithmic Foundations (AL) typically constitute the introductory course sequence in computer science and have been listed first.

	Knowledge Area	Knowledge Units	CS Core	KA Core
SDF	Software Development Fundamentals	5	43	
AL	Algorithmic Foundations	5	32	32
FPL	Foundations of Programming Languages	22	21	19
SE	Software Engineering	9	6	21
	Total		102	72

SDF: Software Development Fundamentals

KU	Topic	Skill Level	Core	Hours
SDF-Fundamentals	<ol style="list-style-type: none"> 1. Basic concepts such as variables, primitive data types, and expression evaluation. 2. How imperative programs work: state and state transitions on execution of statements, flow of control 3. Basic constructs such as assignment statements, conditional and iterative 	Develop	CS	18

	<p>statements, basic I/O</p> <ol style="list-style-type: none"> 4. Key modularity constructs such as functions and related concepts like parameter passing, scope, abstraction, data encapsulation, etc. 5. Input and output using files and APIs 6. Structured data types available in the chosen programming language like sequences, associative containers, others and when and how to use them 7. Libraries and frameworks provided by the language (when/where applicable) 8. Recursion 			
SDF-Fundamentals	<ol style="list-style-type: none"> 9. Dealing with runtime errors in programs 10. Basic concepts of programming errors, testing, and debugging 11. Documenting/commenting code at the program and module level 	Evaluate, Apply	CS	2
SDF-Data-Structures	<ol style="list-style-type: none"> 1. Standard abstract data types such as lists, stacks, queues, sets, and maps/dictionaries and operations on them 4. Strings and string processing 	Develop	CS	10
SDF-Data-Structures	<ol style="list-style-type: none"> 2. Selecting and using appropriate data structures 3. Performance implications of choice of data structure(s) 	Evaluate	CS	2
SDF-Algorithms	<ol style="list-style-type: none"> 1. Concept of algorithm and notion of algorithm efficiency 2. Some common algorithms (e.g. sorting, searching, tree traversal, graph traversal, etc. 3. Impact of algorithms on time/space efficiency of programs 	Explain	CS	6
SDF-Practices	<ol style="list-style-type: none"> 1. Basic testing including test case design 4. Specifying functionality of a module in a natural language 	Develop	CS	3
SDF-Practices	<ol style="list-style-type: none"> 3. Programming style that improves readability 	Evaluate	CS	1
SDF-Practices	<ol style="list-style-type: none"> 2. Use of a general-purpose IDE, including its debugger 	Apply	CS	1

AL: Algorithmic Foundations

KU	Topic	Skill Level	Core	Hours
AL-Foundational AL-Complexity	2. Arrays 1. Abstract Data Types and Operations 2b i. Foundational Complexity Classes: <i>Constant</i>	Explain Apply Evaluate	CS	1
AL-Foundational AL-Complexity AL-Strategies	11a. Search Algorithms $O(n)$ (e.g., linear array search) 2b iii. Foundational Complexity Classes: <i>Linear</i> 1a. Brute Force	Apply Evaluate Explain	CS	1
AL-Foundational AL-Complexity AL-Strategies	12a. Sorting $O(n^2)$, (e.g., selection sort of an array) 2b v. Foundational Complexity classes: <i>Quadratic</i> 1a. Brute Force	Apply Evaluate Explain	CS	1
AL-Foundational AL-Complexity AL-Strategies	11b. Search $O(\log_2 n)$, (e.g., Binary search of an array) 2b ii. Foundational Complexity Classes: <i>Logarithmic</i> 1b ii. Decrease-and-Conquer	Apply Evaluate Explain	CS	1
AL-Foundational AL-Complexity AL-Strategies	12b. Sorting $O(n \log n)$, (e.g., Quick, Merge, Tim) 2b iv. Foundational Complexity Classes: <i>Log Linear</i> 1c. Divide-and-Conquer	Apply Evaluate Explain	CS	1
AL-Foundational AL-Complexity AL-Strategies	4. Linked Lists 1. Abstract Data Types and Operations 11a. Search $O(n)$, (e.g., linear linked list search) 2b iii. Foundational Complexity Classes: <i>Linear</i> 1a. Brute Force	Explain Apply Apply Evaluate Explain	CS	1
AL-Foundational AL-Complexity AL-Foundational	5. Stacks 1. Abstract Data Types and Operations 2b i. Foundational complexity classes: <i>Constant</i> 6. Queues and Deques 1. Abstract Data Types and Operations	Explain, Apply Explain, Apply	CS	1

AL-Foundational AL-Complexity AL-Strategies	7. Hash Tables / Maps 7a. Collision resolution and complexity 1. Abstract Data Types and Operations 2b i. Foundational complexity classes: <i>Constant</i> 1f. Time vs Space tradeoff	Explain Explain Apply Explain Explain	CS	1 1
AL-Foundational AL-Strategies AL-Foundational AL-Strategies	9. Trees 1. Abstract Data Types and Operations 11c. Search Algorithms DFS/BFS 2b. Decrease-and-Conquer 9b. Balanced Trees (e.g., AVL, 2-3, Red-Black, Heap) 1e ii. Transform-and-Conquer: Representation Change (e.g., heapsort)	Explain Apply Apply Explain Apply Explain	CS	1 3
AL-Foundational AL-Foundational AL-Strategies	8. Graphs (e.g., [un]directed, [a]cyclic, [un]connected, [un]weighted) 8a. Representation: Adjacency List vs Matrix 13. Graph Algorithms 13a. Shortest Path (e.g., Dijkstra's, Floyd's) 13b. Minimal, spanning tree (e.g. Prim's, Kruskal's) 1d. Greedy 1e. iv. Dynamic Programming	Explain Apply Apply Explain	CS	1 4
AL-Foundational	1. Abstract Data Types and Operations 3. Records/Structures/Tuples and Objects 10. Sets	Explain	CS	1
AL-Strategies AL-Strategies AL-Strategies AL-Strategies	1. Paradigms (demonstrated in AL-Foundational) 1a. Brute-Force, 1b. Decrease-and-Conquer, 1c. Divide-and-Conquer, 1f. Time-Space Tradeoff, 3. Iteration vs Recursion, 1e. Transform-and-Conquer 1e i. Instance Simplification (e.g., pre-sorting) 1e iii. Problem Reduction (e.g., least-common-multiple) 2. Handling Exponential Growth (e.g., heuristic A*, backtracking, ranch-and-bound) 1e iv. Dynamic Programming (e.g., Bellman-Ford, Knapsack, Floyd, Warshall)	Explain Explain Explain Explain	CS	3 1 1 1

AL-Complexity	1. Complexity Analysis Framework 2. Asymptotic Complexity Analysis 2a. Big O, Big Omega, and Big Theta 2b. Foundational complexity classes demonstrated by AL-Foundational algorithms (with complexity): <i>Constant, Logarithmic, Linear, Log Linear, Quadratic, and Cubic</i>	Explain Explain	CS	1
AL-Complexity	4. Tractability and Intractability 4a. P, NP and NP-C complexity classes 4b. NP-Complete Problems (e.g., SAT, Knapsack, TSP) 4c. Reductions	Explain		1
AL-Strategies	1a. Paradigms: Exhaustive brute force, 1e. iv. Dynamic Programming	Explain		1
AL-Complexity	2 vii. Foundational Complexity Classes: <i>Exponential</i> 2b viii. Factorial complexity classes: Factorial $O(n!)$ (e.g. All Permutations, Hamiltonian Circuit)	Explain		1
AL-Models	1. Formal Automata 1a. Finite State Automata 2a. Regular language, grammar, and expressions 1b. Pushdown Automata 2b. Context-Free language and grammar 1d. Turing Machine 2d. Recursively Enumerable language and grammar 1c. Linear-Bounded 2c. Context-Sensitive language and grammar	Explain, Apply	CS	1
	2. Formal Languages and Grammars	Explain		1
	4. Decidability, Computability, Halting problem	Explain		2
	5. The Church-Turing Thesis	Explain		1
	6. Algorithmic Correctness Invariants (e.g., in iteration, recursion, sorting, heaps)	Explain		
AL-SEP	1. Social, Ethical, and Secure Algorithms 2. Algorithmic Fairness (e.g. differential privacy) 3. Accountability/Transparency 4. Responsible algorithms 5. Economic and other impacts of algorithms 6. Sustainability	Explain	CS	In SEP Hours

FPL: Foundations of Programming Languages

KU	Topic	Skill Level	Core	Hours
FPL-OOP	<ol style="list-style-type: none"> 1. Imperative programming as a subset of object-oriented programming. 2. Object-oriented design: <ol style="list-style-type: none"> a. Decomposition into objects carrying state and having behavior. b. Class-hierarchy design for modeling. 3. Definition of classes: fields, methods, and constructors. 4. Subclasses, inheritance (including multiple inheritance), and method overriding. 5. Dynamic dispatch: definition of method-call. 6. Exception handling. 7. Object-oriented idioms for encapsulation: <ol style="list-style-type: none"> a. Privacy, data hiding, and visibility of class members. b. Interfaces revealing only method signatures. c. Abstract base classes, traits and mixins. 8. Dynamic vs static properties. 9. Composition vs inheritance. 10. Subtyping: <ol style="list-style-type: none"> a. Subtype polymorphism; implicit upcasts in typed languages. b. Notion of behavioral replacement: subtypes acting like supertypes. c. Relationship between subtyping and inheritance. 	Develop	CS	1. 5
	<ol style="list-style-type: none"> 11. Collection classes, iterators, and other common library components. 12. Metaprogramming and reflection. 	Develop	KA	1
FPL-Functional	<ol style="list-style-type: none"> 1. Lambda expressions and evaluation: <ol style="list-style-type: none"> a. Variable binding and scope rules. b. Parameter-passing. c. Nested lambda expressions and reduction order. 2. Effect-free programming: <ol style="list-style-type: none"> a. Function calls have no side effects, facilitating compositional reasoning. b. Immutable variables and data copying vs reduction. 	Develop	CS	4

	<ul style="list-style-type: none"> c. Use of recursion vs loops vs pipelining (map/reduce). <ol style="list-style-type: none"> 3. Processing structured data (e.g., trees) via functions with cases for each data variant: <ul style="list-style-type: none"> a. Functions defined over compound data in terms of functions applied to the constituent pieces. b. Persistent data structures. 4. Using higher-order functions (taking, returning, and storing functions). 			
	<ol style="list-style-type: none"> 5. Metaprogramming and reflection. 6. Function closures (functions using variables in the enclosing lexical environment): <ul style="list-style-type: none"> a. Basic meaning and definition - creating closures at run-time by capturing the environment. b. Canonical idioms: call-backs, arguments to iterators, reusable code via function arguments. c. Using a closure to encapsulate data in its environment. d. Lazy versus eager evaluation. 	Explain	KA	3
FPL-Logic	<ol style="list-style-type: none"> 1. Universal vs existential quantifiers. 2. First order predicate logic vs higher order logic. 3. Expressing complex relations using logical connectives and simpler relations. 4. Definitions of Horn clause, facts, goals and subgoals. 5. Unification and unification algorithm; unification vs assertion vs expression evaluation. 6. Mixing relations with functions. 7. Cuts, backtracking, and non-determinism. 8. Closed-world vs open-world assumptions. 	Explain	KA	3
FPL-Scripting	<ol style="list-style-type: none"> 1. Error/exception handling. 2. Piping. 3. System commands: <ul style="list-style-type: none"> a. Interface with operating systems. 4. Environment variables. 5. File abstraction and operators. 6. Data structures, such as arrays and lists. 7. Regular expressions. 8. Programs and processes. 9. Workflow. 	Develop	CS	2
FPL-Event-Driven	<ol style="list-style-type: none"> 1. Procedural programming vs reactive programming: advantages of reactive programming in capturing 	Develop	CS	2

	<p>events.</p> <ol style="list-style-type: none"> 2. Components of reactive programming: event-source, event signals, listeners and dispatchers, event objects, adapters, event-handlers. 3. Stateless and state-transition models of event-based programming. 4. Canonical uses such as GUIs, mobile devices, robots, servers. 			
	<ol style="list-style-type: none"> 5. Using a reactive framework: <ol style="list-style-type: none"> a. Defining event handlers/listeners. b. Parameterization of event senders and event arguments. c. Externally generated events and program-generated events. 6. Separation of model, view, and controller. 7. Event-driven and reactive programs as state-transition systems. 	Develop	KA	2
FPL-Parallel	<ol style="list-style-type: none"> 1. Safety and liveness: <ol style="list-style-type: none"> a. Race conditions. b. Dependencies/preconditions. c. Fault models. d. Termination. 2. Programming models. One or more of the following: <ol style="list-style-type: none"> a. Actor models. b. Procedural and reactive models. c. Synchronous/asynchronous programming models. d. Data parallelism. 3. Properties <ol style="list-style-type: none"> a. Order-based properties: <ol style="list-style-type: none"> i. Commutativity. ii. Independence. b. Consistency-based properties: <ol style="list-style-type: none"> i. Atomicity. ii. Consensus. 4. Execution control: <ol style="list-style-type: none"> a. Async await. b. Promises. c. Threads. 5. Communication and coordination: <ol style="list-style-type: none"> a. Mutexes. b. Message-passing. 	Develop	CS	3

	<ul style="list-style-type: none"> c. Shared memory. d. Cobegin-coend. e. Monitors. f. Channels. g. Threads. h. Guards. 			
	<ul style="list-style-type: none"> 6. Futures. 7. Language support for data parallelism such as forall, loop unrolling, map/reduce. 8. Effect of memory-consistency models on language semantics and correct code generation. 9. Representational State Transfer Application Programming Interfaces (REST APIs). 10. Technologies and approaches: cloud computing, high performance computing, quantum computing, ubiquitous computing 11. Overheads of message-passing 12. Granularity of program for efficient exploitation of concurrency. 13. Concurrency and other programming paradigms (e.g., functional). 	Explain	KA	2
FPL-Types	<ul style="list-style-type: none"> 1. A type as a set of values together with a set of operations: <ul style="list-style-type: none"> a. Primitive types (e.g., numbers, Booleans). b. Compound types built from other types (e.g., records/structs, unions, arrays, lists, functions, references using set operations). 2. Association of types to variables, arguments, results, and fields. 3. Type safety as an aspect of program correctness. 4. Type safety and errors caused by using values inconsistently given their intended types. 5. Goals and limitations of static and dynamic typing: <ul style="list-style-type: none"> a. Detecting and eliminating errors as early as possible. 6. Generic types (parametric polymorphism): <ul style="list-style-type: none"> a. Definition and advantages of polymorphism: parametric, subtyping, overloading and coercion. b. Comparison of monomorphic and polymorphic types. c. Comparison with ad-hoc polymorphism (overloading) and subtype polymorphism. 	Develop	CS	3

	<ul style="list-style-type: none"> d. Generic parameters and typing. e. Use of generic libraries such as collections. f. Comparison with ad hoc polymorphism (overloading) and subtype polymorphism. g. Prescriptive vs descriptive polymorphism. h. Implementation models of polymorphic types. i. Subtyping. 			
	<ul style="list-style-type: none"> 7. Type equivalence: structural vs name equivalence. 8. Complementary benefits of static and dynamic typing: <ul style="list-style-type: none"> a. Errors early vs errors late/avoided. b. Enforce invariants during code development and code maintenance vs postpone typing decisions while prototyping and conveniently allow flexible coding patterns such as heterogeneous collections. c. Typing rules: <ul style="list-style-type: none"> i. Rules for function, product, and sum types. d. Avoid misuse of code vs allow more code reuse. e. Detect incomplete programs vs allow incomplete programs to run. f. Relationship to static analysis. g. Decidability. 	Develop	KA	4
FPL-Systems	<ul style="list-style-type: none"> 1. Data structures for translation, execution, translation and code mobility such as stack, heap, aliasing (sharing using pointers), indexed sequence and string. 2. Direct, indirect, and indexed access to memory location. 3. Run-time representation of data abstractions such as variables, arrays, vectors, records, pointer-based data elements such as linked-lists and trees, and objects. 4. Abstract low-level machine with simple instruction, stack and heap to explain translation and execution. 5. Run-time layout of memory: activation record (with various pointers), static data, call-stack, heap: <ul style="list-style-type: none"> a. Translating selection and iterative constructs to control-flow diagrams. b. Translating control-flow diagrams to low level abstract code. c. Implementing loops, recursion, and tail calls. 	Develop	CS	3

	<ul style="list-style-type: none"> d. Translating function/procedure calls and return from calls, including different parameter-passing mechanisms using an abstract machine. <p>6. Memory management:</p> <ul style="list-style-type: none"> a. Low level allocation and accessing of high-level data structures such as basic data types, n-dimensional array, vector, record, and objects. b. Return from procedure as automatic deallocation mechanism for local data elements in the stack. c. Manual memory management: allocating, deallocating, and reusing heap memory. d. Automated memory management: garbage collection as an automated technique using the notion of reachability. <p>7. Green computing.</p>			
FPL-Translation	<ol style="list-style-type: none"> 1. Execution models for JIT (Just-In-Time), compiler, interpreter. 2. Use of intermediate code, e.g., bytecode. 3. Limitations and benefits of JIT, compiler, and interpreter. 4. Cross compilers/transpilers. 5. BNF and extended BNF representation of context-free grammar. 6. Parse tree using a simple sentence such as arithmetic expression or if-then-else statement. 7. Execution as native code or within a virtual machine. 8. Language translation pipeline: syntax analysis, parsing, optional type-checking, translation/code generation and optimization, linking, loading, execution. 	Explain	CS	2
	<ol style="list-style-type: none"> 9. Run-time representation of core language constructs such as objects (method tables) and first-class functions (closures). 10. Secure compiler development. 	Explain	KA	2
FPL-Abstraction	<ol style="list-style-type: none"> 1. BNF and regular expressions. 2. Programs that take (other) programs as input such as interpreters, compilers, type-checkers, documentation generators. 3. Components of a language: <ul style="list-style-type: none"> a. Definitions of alphabets, delimiters, sentences, 	Explain	KA	3

	<p>syntax and semantics.</p> <p>b. Syntax vs semantics.</p> <p>4. Program as a set of non-ambiguous meaningful sentences.</p> <p>5. Basic programming abstractions: constants, variables, declarations (including nested declarations), command, expression, assignment, selection, definite and indefinite iteration, iterators, function, procedure, modules, exception handling.</p> <p>6. Mutable vs immutable variables: advantages and disadvantages of reusing existing memory location vs advantages of copying and keeping old values; storing partial computation vs recomputation.</p> <p>7. Types of variables: static, local, nonlocal, global; need and issues with nonlocal and global variables.</p> <p>8. Scope rules: static vs dynamic; visibility of variables; side-effects.</p> <p>9. Side-effects induced by nonlocal variables, global variables and aliased variables.</p>			
--	---	--	--	--

SE: Software Engineering

KU	Topic	Skill Level	Core	Hours
SE-Teamwork	<ol style="list-style-type: none"> 1. Effective communication 2. Common causes of team conflict, and approaches for conflict resolution 3. Cooperative programming 4. Roles and responsibilities in a software team 5. Team processes 6. Importance of team diversity and inclusion 	Evaluate	CS	2
SE-Teamwork	<ol style="list-style-type: none"> 7. Interfacing with stakeholders, as a team 8. Risks associated with physical, distributed, hybrid and virtual teams 	Explain	KA	2
SE-Tools	<ol style="list-style-type: none"> 1. Software configuration management and version control 	Evaluate	CS	1
SE-Tools	<ol style="list-style-type: none"> 2. Release management 3. Testing tools including static and dynamic analysis tools 			

	<ol style="list-style-type: none"> 4. Software process automation 5. Design and communication tools (docs, diagrams, common forms of design diagrams like UML) 6. Tool integration concepts and mechanisms 7. Use of modern IDE facilities - debugging, refactoring, searching/indexing, etc. 	Explain	KA	3
SE-Requirements	<ol style="list-style-type: none"> 1. Describe functional requirements using, for example, use cases or user stories 2. Properties of requirements including consistency, validity, completeness, and feasibility 3. Requirements elicitation 4. Non-functional requirements, for example, security, usability, or performance (aka Quality Attributes) 5. Risk identification and management 6. Communicating and/or formalizing requirement specifications 	Apply	KA	2
SE-Design	<ol style="list-style-type: none"> 1. System design principles 2. Software architecture 3. Programming in the large vs programming in the small 4. Code smells and other indications of code quality, distinct from correctness. 	Explain	CS	1
SE-Design	<ol style="list-style-type: none"> 5. API design principles 6. Identifying and codifying data invariants and time invariants 7. Structural and behavioral models of software designs 8. Data design 9. Requirement traceability 	Apply	KA	4
SE-Construction	<ol style="list-style-type: none"> 1. Practical small-scale testing 2. Documentation 	Apply	CS	1
SE-Construction	<ol style="list-style-type: none"> 3. Coding Style 4. "Best Practices" for coding 5. Debugging 6. Logging 7. Use of libraries and frameworks developed by others 	Apply	KA	3
SE-Validation	<ol style="list-style-type: none"> 1. Verification and validation concepts 2. Why testing matters 	Explain	CS	1

	<ol style="list-style-type: none"> 3. Testing objectives 4. Test kinds 5. Stylistic differences between tests and production code 			
SE-Validation	<ol style="list-style-type: none"> 6. Test planning and generation 7. Test development (see SDF) 8. Verification and validation in the development cycle 9. Domain specific verification and validation challenges 	Explain	KA	4
SE-Refactoring	<ol style="list-style-type: none"> 1. Hyrum's Law 2. Backward Compatibility 3. Refactoring 4. Versioning 	Explain	KA	1
SE-Reliability	<ol style="list-style-type: none"> 1. Concept of reliability 2. Identifying reliability requirements (see SEP) 3. Software failures vs defect injection/detection 4. Software reliability, system reliability and failure behavior (cross-reference SF/Reliability Through Redundancy) 5. Defect injection and removal cycle , and different approaches for defect removal 6. Compare the "error budget" approach to reliability with the "error-free" approach, and identify domains where each is relevant 	Explain	KA	4

Systems Competency Area

The core topics in Architecture and Organization (AR) and Operating Systems (OS) are typically covered early in the curriculum and have been listed first. Data Management (DM) and Security (SEC) topics listed in this section can be applied to all three competency areas.

	Knowledge Area	Knowledge Units	CS Core	KA Core
AR	Architecture and Organization	9	9	16
OS	Operating Systems	14	8	14
NC	Networking and Communication	8	7	24

PDC	Parallel and Distributed Computing	5	9	26
SF	Systems Fundamentals	8	18	8
DM	Data Management	12	10	26
SEC	Security	6	6	35
	Total		67	149

AR: Architecture and Organization

KU	Topic	Skill Level	Core	Hours
AR-Logic	1. Combinational vs sequential logic/field programmable gate arrays (FPGAs) <ul style="list-style-type: none"> a. Fundamental combinational b. Sequential logic building block 	Explain	KA	3
	2. Computer-aided design tools that process hardware and architectural representations 3. High-level synthesis <ul style="list-style-type: none"> a. Register transfer notation b. Hardware description language (e.g., Verilog/VHDL/Chisel) 4. System-on-chip (SoC) design flow 5. Physical constraints <ul style="list-style-type: none"> a. Gate delays b. Fan-in and fan-out c. Energy/power d. Speed of light 	Evaluate		
AR-Representation	1. Overview and history of computer architecture 2. Bits, bytes, and words 3. Unsigned, signed and two's complement representations 4. Numeric data representation and number bases <ul style="list-style-type: none"> a. Fixed-point b. Floating-point 5. Representation of non-numeric data 6. Representation of records, arrays and UTF data types	Apply	CS	1

AR-Assembly	<ol style="list-style-type: none"> 1. von Neumann machine architecture 2. Control unit: instruction fetch, decode, and execution 3. Introduction to SIMD vs MIMD and the Flynn taxonomy 4. Shared memory multiprocessors/multicore organization 	Explain	CS	1
AR-Assembly	<ol style="list-style-type: none"> 5. Instruction set architecture (ISA) (e.g., x86, ARM, and RISC-V) <ol style="list-style-type: none"> a. Instruction formats b. Data manipulation, control, I/O c. Addressing modes d. Machine language programming e. Assembly language programming 6. Subroutine call and return mechanisms 7. I/O and interrupts 8. Heap, static, stack, and code segments 	Develop	KA	2
AR-Memory	<ol style="list-style-type: none"> 1. Memory hierarchy: the importance of temporal and spatial locality 2. Main memory organization and operations 3. Persistent memory (e.g., SSD, standard disks) 4. Latency, cycle time, bandwidth, and interleaving 7. Virtual memory (hardware support) 8. Fault handling and reliability 9. Reliability <ol style="list-style-type: none"> a. Error coding b. Data compression c. Data integrity 10. In-Memory Processing (PIM) 	Explain	CS	6
	<ol style="list-style-type: none"> 5. Cache memories <ol style="list-style-type: none"> a. Address mapping b. Block size c. Replacement and store policy 6. Multiprocessor cache coherence 	Evaluate		
	<ol style="list-style-type: none"> 1. I/O fundamentals <ol style="list-style-type: none"> a. Handshaking and buffering b. Programmed I/O c. Interrupt-driven I/O 2. Interrupt structures: vectored and prioritized, interrupt acknowledgment 	Explain	CS	1

AR-IO	<ol style="list-style-type: none"> 3. I/O devices (e.g., mouse, keyboard, display, camera, sensors, actuators) 4. External storage, physical organization, and drives 5. Bus fundamentals <ol style="list-style-type: none"> a. Bus protocols b. Arbitration c. Direct-memory access (DMA) 			
AR-Organization	<ol style="list-style-type: none"> 1. Implementation of simple datapaths, including instruction pipelining, hazard detection, and resolution 2. Control unit <ol style="list-style-type: none"> a. Hardwired implementation b. Microprogrammed realization 	Develop	KA	2
	<ol style="list-style-type: none"> 3. Instruction pipelining 4. Introduction to instruction-level parallelism (ILP) 	Explain		
AR-Performance-Energy	<ol style="list-style-type: none"> 1. Performance-energy evaluation (introduction): performance, power consumption, memory, and communication costs 2. Branch prediction, speculative execution, out-of-order execution, Tomasulo's algorithm 	Evaluate	KA	2
AR-Performance-Energy	<ol style="list-style-type: none"> 3. Enhancements for vector processors and GPUs 4. Hardware support for Multithreading <ol style="list-style-type: none"> a. Race conditions b. Lock implementations c. Point-to-point synchronization d. Barrier implementation 5. Scalability 6. Alternative architectures, such as VLIW/EPIC, accelerators, and other special-purpose processors 7. Dynamic voltage and frequency scaling (DVFS) 8. Dark Silicon 	Explain	KA	1
AR-Heterogeneity	<ol style="list-style-type: none"> 1. SIMD and MIMD architectures (e.g., General-Purpose GPUs, TPUs, and NPUs) 2. Heterogeneous memory system <ol style="list-style-type: none"> a. Shared memory versus distributed memory b. Volatile vs non-volatile memory c. Coherence protocols 3. Domain-Specific Architectures (DSAs) <ol style="list-style-type: none"> a. Machine Learning Accelerator 	Explain	KA	2

	<ul style="list-style-type: none"> b. In-networking computing c. Embedded systems for emerging applications d. Neuromorphic computing e. Edge computing devices <ul style="list-style-type: none"> 4. Packaging and integration solutions such as 3DIC and Chiplets 5. Machine learning in architecture design <ul style="list-style-type: none"> a. AI algorithms for workload analysis b. Optimization of architecture configurations for performance and power efficiency 			
AR-Security	<ul style="list-style-type: none"> 1. Principles of Secure Hardware <ul style="list-style-type: none"> a. Security Risk Analysis, Asset Protection, and Threat Model b. Cryptographic Acceleration with Hardware c. Support for virtualization (e.g., OS isolation) 2. Roots of trust in hardware, Physically Unclonable Functions (PUF) 3. Hardware Random Number Generators 4. Memory protection extensions <ul style="list-style-type: none"> a. Runtime pointer bounds checking (e.g., buffer overflow) b. Protection at the microarchitectural level c. Protection at the ISA level 5. Trusted Execution Environment (TEE) <ul style="list-style-type: none"> a. Trusted Computer Base Protections b. Protecting virtual machines c. Protecting containers d. Trusted software modules (Enclaves) 	Explain	KA	2
AR-Quantum	<ul style="list-style-type: none"> 1. Principles <ul style="list-style-type: none"> a. The wave-particle duality principle b. The uncertainty principle in the double-slit experiment c. What is a Qubit? Superposition and measurement. Photons as qubits. d. Systems of two qubits. Entanglement. Bell states. The No-Signaling theorem. 2. Axioms of QM: superposition principle, measurement axiom, unitary evolution 3. Single qubit gates for the circuit model of quantum computation: X, Z, H. 	Explain	KA	2

	<ol style="list-style-type: none"> 4. Two qubit gates and tensor products. Working with matrices. 5. The No-Cloning Theorem. The Quantum Teleportation protocol. 6. Algorithms <ol style="list-style-type: none"> a. Simple quantum algorithms: Bernstein-Vazirani, Simon's algorithm. b. Implementing Deutsch-Josza with Mach-Zehnder Interferometers. c. Quantum factoring (Shor's Algorithm) d. Quantum search (Grover's Algorithm) 7. Implementation aspects <ol style="list-style-type: none"> a. The physical implementation of qubits b. Classical control of a Quantum Processing Unit (QPU) c. Error mitigation and control. NISQ and beyond. 8. Emerging Applications <ol style="list-style-type: none"> a. Post-quantum encryption b. The Quantum Internet c. Adiabatic quantum computation (AQC) and quantum annealing 			
--	--	--	--	--

OS: Operating Systems

KU	Topic	Skill Level	Core	Hours
OS-Purpose	<ol style="list-style-type: none"> 1. Operating system as mediator between general purpose hardware and application-specific software 2. Universal operating system functions 3. Extended and/or specialized operating system functions 4. Design issues 5. Influences of security, networking, multimedia, parallel and distributed computing 6. Overarching concern of security/protection: Neglecting to consider security at every layer creates an opportunity to inappropriately access resources. 	Explain	CS	2
OS-Principles	<ol style="list-style-type: none"> 1. Operating system software design and approaches 2. Abstractions, processes, and resources 	Explain	CS	2

	<ol style="list-style-type: none"> 3. Concept of system calls and links to application program interfaces 4. The evolution of the link between hardware architecture and the operating system functions 5. Protection of resources means protecting some machine instructions/functions 6. Leveraging interrupts from hardware level: service routines and implementations 7. Concept of user/system state and protection, transition to kernel mode using system calls 8. Mechanism for invoking of system calls, the corresponding mode and context switch and return from interrupt 9. Performance costs of context switches and associated cache flushes when performing process switches in Spectre-mitigated environments 			
OS-Concurrency	<ol style="list-style-type: none"> 1. Thread abstraction relative to concurrency 2. Race conditions, critical regions (role of interrupts if needed) 3. Deadlocks and starvation 4. Multiprocessor issues (spin-locks, reentrancy) 5. Multiprocess concurrency vs multithreading 	Explain	CS	2
	<ol style="list-style-type: none"> 6. Thread creation, states, structures 7. Thread APIs 8. Deadlocks and starvation (necessary conditions/mitigations) 9. Implementing thread safe code (semaphores, mutex locks, cond vars) 10. Race conditions in shared memory 	Apply	KA	1
OS-Protection	<ol style="list-style-type: none"> 1. Overview of operating system security mechanisms 2. Attacks and antagonism (scheduling, etc.) 3. Review of major vulnerabilities in real operating systems 4. Operating systems mitigation strategies such as backups 	Apply	CS	2
	<ol style="list-style-type: none"> 5. Policy/mechanism separation 6. Security methods and devices 7. Protection, access control, and authentication 	Apply	KA	1
OS-Scheduling	<ol style="list-style-type: none"> 1. Preemptive and non-preemptive scheduling 2. Schedulers and policies. 	Explain	KA	2

	<ol style="list-style-type: none"> 3. Concepts of Symmetric Multi-Processor (SMP) multiprocessor scheduling and cache coherence 4. Timers (e.g., building many timers out of finite hardware timers) 5. Fairness and starvation 			
OS-Process	<ol style="list-style-type: none"> 1. Processes and threads relative to virtualization - Protected memory, process state, memory isolation, etc. 2. Memory footprint/segmentation (stack, heap, etc.) 3. Creating and loading executables, shared libraries, and dynamic linking 4. Dispatching and context switching 5. Interprocess communication 	Explain	KA	2
OS-Memory	<ol style="list-style-type: none"> 1. Review of physical memory, address translation and memory management hardware 2. Impact of memory hierarchy including cache concept, cache lookup, etc. on operating system mechanisms and policy 3. Logical and physical addressing, address space virtualization 4. Concepts of paging, page replacement, thrashing and allocation of pages and frames 5. Allocation/deallocation/storage techniques (algorithms and data structure) performance and flexibility 6. Memory caching and cache coherence and the effect of flushing the cache to avoid speculative execution vulnerabilities 7. Security mechanisms and concepts in memory management including sandboxing, protection, isolation, and relevant vectors of attack 	Explain	KA	2
OS-Devices	<ol style="list-style-type: none"> 1. Buffering strategies 2. Direct Memory Access and Polled I/O, Memory-mapped I/O 3. Historical and contextual - Persistent storage device management (magnetic, SSD, etc.) 	Explain	KA	1
OS-Files	<ol style="list-style-type: none"> 1. Concept of a file including Data, Metadata, Operations and Access-mode 2. File system mounting 3. File access control 4. File sharing 	Explain	KA	2

	<ol style="list-style-type: none"> 5. Basic file allocation methods including linked, allocation table, etc. 6. File system structures comprising file allocation including various directory structures and methods for uniquely identifying files (name, identified or metadata storage location) 7. Allocation/deallocation/storage techniques (algorithms and data structure) impact on performance and flexibility (i.e. Internal and external fragmentation and compaction) 8. Free space management such as using bit tables vs linking 9. Implementation of directories to segment and track file location 			
OS-Advanced-Files	<ol style="list-style-type: none"> 1. File systems: partitioning, mount/unmount, virtual file systems 2. In-depth implementation techniques 3. Memory-mapped files 4. Special-purpose file systems 5. Naming, searching, access, backups 6. Journaling and log-structured file systems 	Explain	KA	1
OS-Virtualization	<ol style="list-style-type: none"> 1. Using virtualization and isolation to achieve protection and predictable performance 2. Advanced paging and virtual memory 3. Virtual file systems and virtual devices 4. Containers and their comparison to virtual machine 5. Thrashing 	Explain	KA	1
OS-Real-time	<ol style="list-style-type: none"> 1. Process and task scheduling 2. Deadlines and real-time issues (See also: SPD-Embedded) 3. Low-latency vs "soft real-time" vs "hard real time" 	Explain	KA	1
OS-Faults	<ol style="list-style-type: none"> 1. Reliable and available systems 2. Software and hardware approaches to address tolerance (RAID) 	Explain	KA	1
OS-SEP	<ol style="list-style-type: none"> 1. Open source in operating systems 2. End-of-life issues with sunsetting operating systems 	Explain	KA	-

KU	Topic	Skill Level	Core	Hours
NC-Fundamentals	1. Importance of networking in contemporary computing, and associated challenges.	Explain	CS	3
	2. Organization of the Internet <ul style="list-style-type: none"> a. Users, b. Internet Service Providers c. Autonomous systems d. Content providers e. Content delivery networks 	Explain		
	3. Switching techniques <ul style="list-style-type: none"> a. Circuit Switching b. Packet Switching 	Evaluate		
	4. Layers and their roles. <ul style="list-style-type: none"> a. Application b. Transport c. Network d. Datalink e. Physical 	Explain		
	5. Layering principles <ul style="list-style-type: none"> a. Encapsulation b. Hourglass model 	Explain		
	6. Network elements <ul style="list-style-type: none"> a. Routers b. Switches c. Hubs d. Access points e. Hosts 	Explain		
	7. Basic queueing concepts <ul style="list-style-type: none"> a. Relationship with latency b. Relationship with Congestion c. Relationship with Service levels 	Explain		
NC-Applications	1. Naming and address schemes. <ul style="list-style-type: none"> a. DNS b. IP addresses c. Uniform Resource Identifiers 	Explain	CS	4

	<p>2. Distributed application paradigms</p> <ul style="list-style-type: none"> a. Client/server b. Peer-to-peer c. Cloud d. Edge e. Fog 	Evaluate		
	<p>3. Diversity of networked application demands</p> <ul style="list-style-type: none"> a. Latency b. Bandwidth c. Loss tolerance 	Explain		
	<p>4. Application-layer development using one or more protocols:</p> <ul style="list-style-type: none"> a. HTTP b. SMTP c. POP3 	Develop		
	<p>5. Interactions with TCP, UDP, and Socket APIs.</p>	Explain		
NC-Reliability	<p>1. Unreliable delivery</p> <ul style="list-style-type: none"> a. UDP b. Other 	Explain	KA	6
	<p>2. Principles of reliability</p> <ul style="list-style-type: none"> a. Delivery without loss b. Duplication c. Out of order 	Develop		
	<p>3. Error control</p> <ul style="list-style-type: none"> a. Retransmission b. Error correction 	Evaluate		
	<p>4. Flow control</p> <ul style="list-style-type: none"> a. Stop and wait b. Window based 	Develop		
	<p>5. Congestion control</p> <ul style="list-style-type: none"> a. Implicit congestion notification b. Explicit congestion notification 	Explain		
	<p>6. TCP and performance issues</p> <ul style="list-style-type: none"> a. Tahoe b. Reno c. Vegas 	Evaluate		

	<ul style="list-style-type: none"> d. Cubic e. QUIC 			
NC-Routing	<ul style="list-style-type: none"> 1. Routing paradigms and hierarchy <ul style="list-style-type: none"> a. Intra/inter domain b. Centralized and decentralized c. Source routing d. Virtual circuits e. QoS 	Evaluate	KA	4
	<ul style="list-style-type: none"> 2. Forwarding methods <ul style="list-style-type: none"> a. Forwarding tables b. Matching algorithms 	Apply		
	<ul style="list-style-type: none"> 3. IP and Scalability issues <ul style="list-style-type: none"> a. NAT b. CIDR c. BGP d. Different versions of IP 	Explain		
NC-SingleHop	<ul style="list-style-type: none"> 1. Introduction to modulation, bandwidth, and communication media. 	Explain	KA	3
	<ul style="list-style-type: none"> 2. Encoding and Framing. 	Evaluate		
	<ul style="list-style-type: none"> 3. Medium Access Control (MAC) <ul style="list-style-type: none"> a. Random access b. Scheduled access 	Evaluate		
	<ul style="list-style-type: none"> 4. Ethernet 	Explain		
	<ul style="list-style-type: none"> 5. Switching 	Apply		
	<ul style="list-style-type: none"> 6. Local Area Network Topologies (e.g. data center networks) 	Explain		
NC-Security	<ul style="list-style-type: none"> 1. General intro about security <ul style="list-style-type: none"> a. Threats b. Vulnerabilities c. Countermeasures 	Explain	KA	4
	<ul style="list-style-type: none"> 2. Network specific threats and attack types <ul style="list-style-type: none"> a. Denial of service b. Spoofing c. Sniffing d. Traffic redirection 	Explain		

	<ul style="list-style-type: none"> e. Man-in-the-middle f. Message integrity attacks g. Routing attacks h. Traffic analysis 			
	<ul style="list-style-type: none"> 3. Countermeasures [Shared with Security] <ul style="list-style-type: none"> a. Cryptography (e.g. SSL, symmetric/asymmetric). b. Architectures for secure networks (e.g., secure channels, secure routing protocols, secure DNS, VPNs, DMZ, Zero Trust Network Access, hyper network security, anonymous communication protocols, isolation) c. Network monitoring, intrusion detection, firewalls, spoofing and DoS protection, honeypots, tracebacks, BGP Sec. 	Explain		
NC-Mobility	1. Principles of cellular communication (e.g. 4G, 5G)	Explain	KA	3
	2. Principles of Wireless LANs (mainly 802.11)	Explain		
	3. Device to device communication	Explain		
	4. Multihop wireless networks	Explain		
	5. Examples (e.g ad hoc networks, opportunistic, delay tolerant)	Explain		
NC-Emerging	1. Middleboxes (e.g. filtering, deep packet inspection, load balancing, NAT, CDN)	Explain	KA	4
	2. Virtualization (e.g. SDN, Data Center Networks)	Explain		
	3. Quantum Networking (e.g. Intro to the domain, teleportation, security, Quantum Internet)	Explain		

PDC: Parallel and Distributed Computing

KU	Topic	Skill Level	Core	Hours
PDC-	1. Parallelism	Explain	CS	2

<p>Programs</p>	<ul style="list-style-type: none"> a. Declarative parallelism: Determining which actions may, or must not, be performed in parallel, at the level of instructions, functions, closures, composite actions, sessions, tasks, and services is the main idea underlying PDC algorithms; failing to do so is the main source of errors. b. Defining order: for example, using happens-before relations or series/parallel directed acyclic graphs representing programs. c. Independence: determining when ordering doesn't matter, in terms of commutativity, dependencies, preconditions. d. Ensuring ordering among otherwise parallel actions when necessary, including locking, safe publication; and imposing communication: <ul style="list-style-type: none"> – sending a message happens before receiving it; conversely relaxing when unnecessary. <p>2. Distribution</p> <ul style="list-style-type: none"> a. Defining places as devices executing actions, including hardware components, remote hosts, may also include external, uncontrolled devices, hosts, and users. b. One device may time-slice or otherwise emulate multiple parallel actions by fewer processors by scheduling and virtualization. c. Naming or identifying places (e.g., device IDs) and actions as parties (e.g., thread IDs). d. Activities across places may communicate across media. <p>3. Starting activities</p> <ul style="list-style-type: none"> a. Options that enable actions to be performed (eventually) at places range from hardwiring to configuration scripts; also establishing communication and resource management; these are expressed differently across languages and contexts, usually relying on automated provisioning and management by platforms. b. Procedural: Enabling multiple actions to start at a given program point; for example, starting 			
---------------------------------	---	--	--	--

	<p>new threads, possibly scoping or otherwise organizing them in hierarchical groups.</p> <ul style="list-style-type: none"> c. Reactive: Enabling upon an event by installing an event handler, with less control of when actions begin or end. d. Dependent: Enabling upon completion of others; for example, sequencing sets of parallel actions. e. Granularity: Execution cost of action bodies should outweigh the overhead of arranging. <p>4. Execution Properties</p> <ul style="list-style-type: none"> a. Nondeterministic execution of unordered actions. b. Consistency: Ensuring agreement among parties about values and predicates when necessary to avoid races, maintain safety and atomicity, or arrive at consensus. c. Fault tolerance: Handling failures in parties or communication, including (Byzantine) misbehavior due to untrusted parties and protocols, when necessary to maintain progress or availability. d. Tradeoffs are one focus of evaluation. 			
	<p>5. One or more of the following mappings and mechanisms across layered systems:</p> <ul style="list-style-type: none"> a. CPU data- and instruction-level- parallelism. b. SIMD and heterogeneous data parallelism. c. Multicore scheduled concurrency, tasks, actors. d. Clusters, clouds; elastic provisioning. e. Networked distributed systems. f. Emerging technologies such as quantum computing and molecular computing. 	Explain, Develop	KA	2
PDC-Communication	<p>1. Media</p> <ul style="list-style-type: none"> a. Varieties: channels (message passing or IO), shared memory, heterogeneous, data stores. b. Reliance on the availability and nature of underlying hardware, connectivity, and protocols; language support, emulation. <p>2. Channels</p> <ul style="list-style-type: none"> a. Explicit (usually named) party-to-party communication media. 	Explain	CS	2

	<ul style="list-style-type: none"> b. APIs: sockets, architectural and language-based constructs, and layered constructs such as RPC (remote procedure call). c. IO channel APIs. <p>3. Memory</p> <ul style="list-style-type: none"> a. Shared memory architectures in which parties directly communicate only with memory at given addresses, with extensions to heterogeneous memory supporting multiple memory stores with explicit data transfer across them; for example, GPU local and shared memory, Direct Memory Access (DMA). b. Memory hierarchies: Multiple layers of sharing domains, scopes and caches; locality: latency, false-sharing. c. Consistency properties: Bitwise atomicity limits, coherence, local ordering. <p>4. Data Stores</p> <ul style="list-style-type: none"> a. Cooperatively maintained data structures implementing maps and related ADTs. b. Varieties: Owned, shared, sharded, replicated, immutable, versioned. 			
	<p>5. One or more of the following properties and extensions</p> <ul style="list-style-type: none"> a. Topologies: Unicast, Multicast, Mailboxes, Switches; Routing via hardware and software interconnection networks b. Media concurrency properties: Ordering, consistency, idempotency, overlapping communication with computation c. Media performance: Latency, bandwidth (throughput) contention (congestion), responsiveness (liveness), reliability (error and drop rates), protocol-based progress (acks, timeouts, mediation) d. Media security properties: integrity, privacy, authentication, authorization e. Data formats: marshaling, validation, encryption, compression f. Channel policies: Endpoints, Sessions, Channel policies: Endpoints, sessions, 	Explain, Develop	KA	6

	<p>buffering, saturation response (waiting vs dropping), Rate control.</p> <ul style="list-style-type: none"> g. Multiplexing and demultiplexing many relatively slow I/O devices or parties; completion-based and scheduler-based techniques; async-await, select and polling APIs. h. Formalization and analysis of channel communication; for example, CSP. i. Applications of queuing theory to model and predict performance. j. Memory models: sequential and release/acquire consistency. k. Memory management; including reclamation of shared data; reference counts and alternatives. l. Bulk data placement and transfer; reducing message traffic and improving locality; overlapping data transfer and computation; impact of data layout such as array-of-structs vs struct-of-arrays. m. Emulating shared memory: distributed shared memory, Remote Direct Memory Access (RDMA). n. Data store consistency: Atomicity, linearizability, transactionality, coherence, causal ordering, conflict resolution, eventual consistency, blockchains. o. Faults, partitioning, and partial failures; voting; protocols such as Paxos and Raft. p. Design tradeoffs among consistency, availability, partition (fault) tolerance; impossibility of meeting all at once. q. Security and trust: Byzantine failures, proof of work and alternatives. 			
PDC-Coordination	<ol style="list-style-type: none"> 1. Dependencies <ul style="list-style-type: none"> a. Initiation or progress of one activity may be dependent on other activities, so as to avoid race conditions, ensure termination, or meet other requirements. b. Ensuring progress by avoiding dependency cycles, using monotonic conditions, removing inessential dependencies. 2. Control constructs and design patterns 	Explain	CS	2

	<ul style="list-style-type: none"> a. Completion-based: Barriers, joins, including termination control. b. Data-enabled: Queues, producer-consumer designs. c. Condition-based: Polling, retrying, backoffs, helping, suspension, signaling, timeouts d. Reactive: enabling and triggering continuations. <p>3. Atomicity</p> <ul style="list-style-type: none"> a. Atomic instructions, enforced local access orderings. b. Locks and mutual exclusion; lock granularity. c. Deadlock avoidance: Ordering, coarsening, randomized retries; encapsulation via lock managers. d. Common errors: Failing to lock or unlock when necessary, holding locks while invoking unknown operations. e. Avoiding locks: replication, read-only, ownership, and non-blocking constructions. 			
	<p>4. One or more of the following properties and extensions:</p> <ul style="list-style-type: none"> a. Progress properties including lock-free, wait-free, fairness, priority scheduling; interactions with consistency, reliability. b. Performance with respect to contention, granularity, convoying, scaling. c. Non-blocking data structures and algorithms. d. Ownership and resource control, e. Lock variants and alternatives: sequence locks, read-write locks; Read-Copy-Update (RCU), reentrancy; tickets; controlling spinning versus blocking. f. Transaction-based control: Optimistic and conservative. g. Distributed locking: reliability. h. Alternatives to barriers: Clocks; Counters, Virtual clocks; Dataflow and continuations; Futures and RPC; Consensus-based, Gathering results with reducers and collectors i. Speculation, selection, cancellation; observability and security consequences 	Explain, Develop	KA	6

	<ul style="list-style-type: none"> j. Resource control using Semaphores and condition variables k. Control flow: Scheduling computations, Series-parallel loops with (possibly elected) leaders, Pipelines and Streams, nested parallelism. l. Exceptions and failures. Handlers, detection, timeouts, fault tolerance, voting. 			
PDC-Evaluation:	<ol style="list-style-type: none"> 1. Safety and liveness requirements in terms of temporal logic constructs to express “always” and “eventually”. 2. Identifying, testing for, and repairing violations, including common forms of errors such as failure to ensure necessary ordering (race errors), atomicity (including check-then-act errors), and termination (livelock). 3. Performance requirements metrics for throughput, responsiveness, latency, availability, energy consumption, scalability, resource usage, communication costs, waiting and rate control, fairness; service level agreements. 4. Performance impact of design and implementation choices, including granularity, overhead, and energy consumption. 5. Estimating scalability limitations, for example using Amdahl’ Law or Universal Scalability Law. 	Explain, Evaluate	CS	1
	<ol style="list-style-type: none"> 6. One or more of the following methods and tools: <ul style="list-style-type: none"> a. Extensions to formal sequential requirements such as linearizability. b. Protocol, session, and transactional specifications. c. Use of tools such as Unified Modelling Language (UML), Temporal Logic of Actions (TLA), program logics. d. Security analysis: safety and liveness in the presence of hostile or buggy behaviors by other parties; required properties of communication mechanisms (for example lack of cross-layer leakage), input screening, rate limiting. e. Static analysis applied to correctness, throughput, latency, resources, energy 	Explain, Evaluate	KA	3

	<ul style="list-style-type: none"> f. Directed Acyclic Graph (DAG) model analysis of algorithmic efficiency (work, span, critical paths). g. Testing and debugging; tools such as race detectors, fuzzers, lock dependency checkers, unit/stress/torture tests, visualizations, continuous integration, continuous deployment, and test generators, h. Measuring and comparing throughput, overhead, waiting, contention, communication, data movement, locality, resource usage, behavior in the presence of excessive numbers of events, clients, or threads. i. Application domain specific analyses and evaluation techniques. 			
PDC- Algorithms	<ol style="list-style-type: none"> 1. Expressing and implementing algorithms in given languages and frameworks, to initiate activities (for example threads), use shared memory constructs, and channel, socket, and/or remote procedure call APIs. <ul style="list-style-type: none"> a. Data parallel examples including map/reduce b. Using channel, socket, and/or RPC APIs in a given language, with program control for sending (usually procedural) vs receiving (usually reactive or RPC-based) c. Using locks, barriers, and/or synchronizers to; maintain liveness without introducing races 2. Survey of common application domains across multicore, reactive, data parallel, cluster, cloud, open distributed systems and frameworks. 	Explain, Develop	CS	2
	<ol style="list-style-type: none"> 3. One or more of the following algorithmic domains: <ul style="list-style-type: none"> a. Linear algebra: Vector and matrix operations, numerical precision/stability, applications in data analytics and machine learning. b. Data processing: sorting, searching and retrieval, concurrent data structures. c. Graphs, search, and combinatorics: Marking, edge-parallelization, bounding, speculation, network-based analytics. d. Modeling and simulation: differential equations; randomization, N-body problems, genetic algorithms. 	Explain, Develop, Evaluate	KA	9

	<ul style="list-style-type: none"> e. Computational Logic: Satisfiability (SAT), concurrent logic programming. f. Graphics and computational geometry: Transforms, rendering, ray-tracing. g. Resource management: Allocating, placing, recycling and scheduling processors, memory, channels, and hosts; exclusive vs shared resources; static, dynamic and elastic algorithms; Real-time constraints; Batching, prioritization, partitioning; decentralization via work-stealing and related techniques. h. Services: Implementing web APIs, electronic currency, transaction systems, multiplayer games. 			
--	---	--	--	--

SF: Systems Fundamentals

KU	Topic	Skill Level	Core	Hours
SF-Overview	<ol style="list-style-type: none"> 1. Basic building blocks and components of a computer (gates, flip-flops, registers, interconnections; datapath + control + memory) 2. Hardware as a computational paradigm: Fundamental logic building blocks; Logic expressions, minimization, sum of product forms 3. Programming abstractions, interfaces, and use of libraries 4. Distinction and interaction between application and OS services, remote procedure call 5. Basic concept of pipelining, overlapped processing stages 6. Basic concept of scaling: going faster vs handling larger problems 	Explain	CS	3
SF-Foundations	<ol style="list-style-type: none"> 1. Digital vs Analog/Discrete vs Continuous Systems 2. Simple logic gates, logical expressions, Boolean logic simplification 3. Clocks, State, Sequencing 4. State and state transition (e.g., starting state, final state, life cycle of states) 5. Finite state machines (e.g., NFA, DFA) 	Apply	CS	4

	<ul style="list-style-type: none"> 6. Combinational Logic, Sequential Logic, Registers, Memories 7. Computers and Network Protocols as examples of State Machines 8. Sequential vs parallel processing 9. Application-level sequential processing: single thread 10. Simple application-level parallel processing: request level (web services/client-server/distributed), single thread per server, multiple threads with multiple servers, pipelining 			
SF-Resource	<ul style="list-style-type: none"> 1. Different types of resources (e.g., processor share, memory, disk, net bandwidth) 2. Common resource allocation/scheduling algorithms (e.g., first-come-first-serve, priority-based scheduling, fair scheduling and preemptive scheduling) 	Explain	CS	1
	<ul style="list-style-type: none"> 3. Advantages and disadvantages of common scheduling algorithms 	Explain	KA	1
SF-Performance	<ul style="list-style-type: none"> 1. Latencies in computer systems <ul style="list-style-type: none"> a. Speed of light and computers (one foot per nanosecond vs one GHz clocks) b. Memory vs disk latencies vs across-the-network memory 2. Caches and the effects of spatial and temporal locality on performance in processors and systems 3. Caches and cache coherency in databases, operating systems, distributed systems, and computer architecture 4. Introduction to the processor memory hierarchy 	Apply	CS	2
	<ul style="list-style-type: none"> 5. The formula for average memory access time 6. Rationale of virtualization and isolation: protection and predictable performance 7. Levels of indirection, illustrated by virtual memory for managing physical memory resources 8. Methods for implementing virtual memory and virtual machines 	Apply	KA	2
	<ul style="list-style-type: none"> 1. Performance figures of merit 	Evaluate	CS	2

SF-Evaluation	2. Workloads and representative benchmarks, and methods of collecting and analyzing performance figures of merit 3. CPI (Cycles per Instruction) equation as a tool for understanding tradeoffs in the design of instruction sets, processor pipelines, and memory system organizations. 4. Amdahl's Law: the part of the computation that cannot be sped up limits the effect of the parts that can 5. Order of magnitude analysis (Big O notation) 6. Analysis of slow and fast paths of a system 7. Events on their effect on performance (e.g., instruction stalls, cache misses, page faults)			
	8. Analytical tools to guide quantitative evaluation 9. Understanding layered systems, workloads, and platforms, their implications for performance, and the challenges they represent for evaluation 10. Microbenchmarking pitfalls	Evaluate	KA	2
SF-Reliability	1. Distinction between bugs and faults 2. Reliability vs availability 3. Reliability through redundancy a. check and retry b. redundant encoding (error correction codes, CRC, FEC) c. duplication/mirroring/replicas	Evaluate	CS	2
	4. Other approaches to reliability (e.g., journaling)	Evaluate	KA	1
SF-Security	1. Common system security issues (e.g., viruses, denial-of-service attacks and eavesdropping) 2. Countermeasures a. Cryptography b. Security architecture	Evaluate	CS	2
	3. Representative countermeasure systems a. Intrusion detection systems, firewalls	Evaluate	KA	1
SF-Design	1. Common criteria of system design (e.g., liveness, safety, robustness, scalability and security)	Design	CS	2
	2. Designs of representative systems (e.g., Apache web server, Spark and Linux)	Design	KA	1

DM: Data Management

KU	Topic	Skill Level	Core	Hours
DM-Data	1. The Data Life Cycle	Evaluate	CS	2
DM-Core	1. Purpose and advantages of database systems 2. Components of database systems 3. Design of core DBMS functions (e.g., query mechanisms, transaction management, buffer management, access methods) 4. Database architecture, data independence, and data abstraction 5. Transaction management 6. Normalization 7. Approaches for managing large volumes of data (e.g., NoSQL database systems, use of MapReduce) 8. How to support CRUD-only applications 9. Distributed databases/cloud-based systems 10. Structured, semi-structured, and unstructured databases	Explain	CS	2
	11. Use of a declarative query language	Develop		
DM-Core	12. Systems supporting structured and/or stream content	Explain	KA	1
DM-Modeling	1. Data modeling 2. Relational data models	Develop	CS	2
DM-Modeling	3. Conceptual models (e.g., entity-relationship, UML diagrams) 4. Semi-structured data model (expressed using DTD, XML, or JSON Schema, for example)	Explain	KA	3
DM-Relational	1. Entity and referential integrity a. Candidate key, superkeys 2. Relational database design	Explain	CS	1
DM-Relational	3. Mapping conceptual schema to a relational schema 4. Physical database design: file and storage structures 5. Introduction to Functional dependency theory 6. Normalization theory	Develop	KA	3

	<ul style="list-style-type: none"> a. Decomposition of a schema; lossless-join and dependency-preservation properties of a decomposition b. Normal forms (BCNF) c. Denormalization (for efficiency) 			
DM-Querying	<ul style="list-style-type: none"> 1. SQL Query Formation <ul style="list-style-type: none"> a. Interactive SQL execution b. Programmatic execution of an SQL query 	Develop	CS	2
DM-Querying	<ul style="list-style-type: none"> 2. Relational Algebra 3. SQL <ul style="list-style-type: none"> a. Data definition including integrity and other constraints specification b. Update sublanguage 	Develop	KA	4
DM-Processing	<ul style="list-style-type: none"> 1. Page structures 2. Index structures <ul style="list-style-type: none"> a. B+ trees b. Hash indices: static and dynamic c. Index creation in SQL 3. File Structures <ul style="list-style-type: none"> a. Heap files b. Hash files 4. Algorithms for query operators <ul style="list-style-type: none"> a. External Sorting b. Selection c. Projection; with and without duplicate elimination d. Natural Joins: Nested loop, Sort-merge, Hash join e. Analysis of algorithm efficiency 5. Query transformations 6. Query optimization <ul style="list-style-type: none"> a. Access paths b. Query plan construction c. Selectivity estimation d. Index-only plans 7. Parallel Query Processing (e.g., parallel scan, parallel join, parallel aggregation) 	Explain	KA	4
	<ul style="list-style-type: none"> 8. Database tuning/performance <ul style="list-style-type: none"> a. Index selection b. Impact of indices on query performance 	Develop		

	c. Denormalization			
DM-Internals	<ol style="list-style-type: none"> 1. DB Buffer Management 2. Transaction Processing <ol style="list-style-type: none"> a. Isolation Levels b. ACID c. Serializability d. Distributed transactions 3. Concurrency Control: <ol style="list-style-type: none"> a. 2-Phase Locking b. Deadlocks handling strategies c. Quorum-based consistency models 4. Recovery Manager <ol style="list-style-type: none"> a. Relation with Buffer Manager 	Explain	KA	4
DM-NoSQL	<ol style="list-style-type: none"> 1. Why NoSQL? (e.g., Impedance mismatch between Application [CRUD] and RDBMS) 2. Key-Value and Document data model 	Explain	KA	2
DM-Analytics	<ol style="list-style-type: none"> 1. Exploratory data techniques (motivation, representation, descriptive statistics, visualizations) 2. Data science lifecycle: business understanding, data understanding, data preparation, modeling, evaluation, deployment, and user acceptance 3. Data mining and machine learning algorithms: e.g., classification, clustering, association, regression 4. Data acquisition and governance 5. Data security and privacy considerations 6. Data fairness and bias 7. Data visualization techniques and their use in data analytics 8. Entity Resolution 	Explain	KA	3
DM-Security	<ol style="list-style-type: none"> 1. Differences between data security and data privacy 2. Protecting data and database systems from attacks, including injection attacks such as SQL injection 3. Personally identifying information (PII) and its protection 4. Ethical considerations in ensuring the security and privacy of data 	Explain	CS	1

DM-Security	<ol style="list-style-type: none"> 5. Need for, and different approaches to securing data at rest, in transit, and during processing 6. Database auditing and its role in digital forensics 7. Data inferencing and preventing attacks 8. Laws and regulations governing data security and data privacy 	Explain	KA	2
DM-SEP	<ol style="list-style-type: none"> 1. Issues related to scale 2. Data privacy overall <ol style="list-style-type: none"> a. Privacy compliance by design 3. Data anonymity 4. Data ownership/custodianship 5. Intended and unintended applications of stored data 	Explain	CS	
DM-SEP	<ol style="list-style-type: none"> 6. Reliability of data 7. Provenance, data lineage, and metadata management 8. Data security 	Explain	KA	

SEC: Security

KU	Topic	Skill Level	Core	Hours
SEC-Foundations	<ol style="list-style-type: none"> 1. Developing a security mindset incorporating crosscutting concepts: confidentiality, integrity, availability, risk assessment, systems thinking, adversarial thinking, human-centered thinking 2. Basic concepts of authentication and authorization/access control 3. Vulnerabilities, threats, attack surfaces, and attack vectors 4. Denial of Service (DoS) and Distributed Denial of Service (DDoS) 5. Principles and practices of protection, e.g., least privilege, open design, fail-safe defaults, defense in depth, and zero trust; and how they can be implemented 6. Optimization considerations between security, privacy, performance, and other design goals 	Develop	CS	1

	7. Impact of AI on security and privacy: using AI to bolster defenses as well as address increased adversarial capabilities due to AI			
SEC-Foundations	8. Access control models (e.g., discretionary, mandatory, role-based, and attribute-based) 9. Security controls 10. Concepts of trust and trustworthiness 11. Applications of a security mindset: web, cloud, and mobile devices 12. Protecting embedded and cyber-physical systems 13. Principles of usable security and human-centered computing 14. Security and trust in AI/machine learning systems, e.g., fit for purpose, ethical operating boundaries, authoritative knowledge sources, verified training data, repeatable system evaluation tests, system attestation, independent validation/certification; unintended consequences from: adverse effect 15. Security risks in building and operating AI/machine learning systems, e.g., algorithm bias, knowledge corpus bias, training corpus bias, copyright violation 16. Hardware considerations in security, e.g., principles of secure hardware, secure processor architectures, cryptographic acceleration, compartmentalization, software-hardware interaction	Develop	KA	5
SEC-SEP	1. Principles and practices of privacy 2. Societal impacts on breakdowns in security and privacy. 3. Applicability of laws and regulations on security and privacy 4. Professional ethical considerations when designing secure systems and maintaining privacy; ethical hacking	Develop	CS	1
SEC-SEP	5. Security by design 6. Privacy by design and privacy engineering 7. Security and privacy implications of malicious AI/machine learning actors, e.g., identifying deep fakes 8. Societal impacts of Internet of Things (IoT) devices and other emerging technologies on security and privacy	Develop	KA	2

SEC-Coding	<ol style="list-style-type: none"> 1. Common vulnerabilities and weaknesses 2. SQL injection and other injection attacks 3. Cross-site scripting techniques and mitigations 4. Input validation and data sanitization 5. Type safety and type-safe languages 6. Buffer overflows, stack smashing, and integer overflows 7. Security issues due to race conditions 	Develop	CS	2
SEC-Coding	<ol style="list-style-type: none"> 8. Principles of noninterference and non-deducibility 9. Preventing information flow attacks 10. Offensive security techniques as a defense 11. AI-assisted malware detection techniques 12. Ransomware: creation, prevention, and mitigation. 13. Secure use of third-party components 14. Malware: varieties, creation, reverse engineering, and defense against them 15. Assurance: testing (including fuzzing and penetration testing), verification and validation 16. Static and dynamic analyses 17. Secure compilers and secure code generation 	Develop	KA	5
SEC-Crypto	<ol style="list-style-type: none"> 1. Differences between algorithmic, applied, and mathematical views of cryptography 2. Mathematical preliminaries: modular arithmetic, Euclidean algorithm, probabilistic independence, linear algebra basics, number theory, finite fields, complexity, asymptotic analysis. 3. Basic cryptography: symmetric key and public key cryptography 4. Basic cryptographic building blocks, including symmetric encryption, asymmetric encryption, hashing, and message authentication 5. Classical cryptosystems, such as shift, substitution, transposition ciphers, code books, machines 6. Kerckhoff's principle and use of vetted libraries 7. History and real-world applications, e.g., electronic cash, secure channels between clients and servers, secure electronic mail, entity authentication, device pairing, voting systems 	Evaluate	CS	1
SEC-Crypto	<ol style="list-style-type: none"> 8. Additional mathematics: primality and factoring; elliptic curve cryptography. 	Develop	KA	4

	<ul style="list-style-type: none"> 9. Private-key cryptosystems: substitution-permutation networks, linear cryptanalysis, differential cryptanalysis, DES, AES 10. Public-key cryptosystems: Diffie-Hellman, RSA 11. Data integrity and authentication: hashing, digital signatures 12. Cryptographic protocols: challenge-response authentication, zero-knowledge protocols, commitment, oblivious transfer, secure two- or multi-party computation, hash functions, secret sharing, and applications 13. Attacker capabilities: chosen-message attack (for signatures), birthday attacks, side channel attacks, fault injection attacks 14. Quantum cryptography; Post Quantum/Quantum resistant cryptography 15. Blockchain and cryptocurrencies 			
SEC-Engineering	<ul style="list-style-type: none"> 1. Security engineering goals: building systems that remain dependable despite errors, accidents, or malicious adversaries 2. Privacy engineering goals: building systems that design, implement, and deploy privacy features and controls 3. Problem analysis and situational analysis to address system security 4. Engineering tradeoff analysis based on time, cost, risk tolerance, risk acceptance, return on investment, and so on 	Develop	CS	1
SEC-Engineering	<ul style="list-style-type: none"> 5. Security design and engineering, including functional requirements, secure subsystems, information protection, security testing, security assessment, and evaluation 6. Security analysis, covering security requirements analysis; security controls analysis; threat analysis; and vulnerability analysis 7. Security attack domains and attack surfaces, e.g., communications and networking, hardware, physical, social engineering, software, and supply chain 8. Security attack modes, techniques, and tactics, e.g., authentication abuse; brute force; buffer manipulation; code injection; content insertion; denial of service; eavesdropping; function bypass; impersonation; 	Develop	KA	8

	<p>integrity attack; interception; phishing; protocol analysis; privilege abuse; spoofing; and traffic injection</p> <p>9. Attestation of software products with respect to their specification and adaptiveness</p> <p>10. Design and development of cyber-physical systems</p> <p>11. Considerations for trustworthy computing, e.g., tamper resistant packaging, trusted boot, trusted kernel, hardware root of trust, software signing and verification, hardware-based cryptography, virtualization, and containers</p>			
SEC-Forensics	<ol style="list-style-type: none"> 1. Basic principles and methodologies for digital forensics 2. System design for forensics 3. Forensics in different situations: operating systems, file systems, application forensics, web forensics, network forensics, mobile device forensics, use of database auditing 4. Attacks on forensics and preventing such attacks 5. Incident handling processes 6. Rules of evidence – general concepts and differences between jurisdictions 7. Legal issues: digital evidence protection and management, chains of custody, reporting, serving as an expert witness. 	Develop	KA	6
SEC-Governance	<ol style="list-style-type: none"> 1. Protecting critical assets from threats 2. Security governance: organizational objectives and general risk assessment 3. Security management: achieve and maintain appropriate levels of confidentiality, integrity, availability, accountability, authenticity, and reliability 4. Security policy: Organizational policies, issue-specific policies, system-specific policies 5. Approaches to identifying and mitigating risks to computing infrastructure 6. Data lifecycle management policies: data collection, backups, and retention; cloud storage and services; breach disclosure 	Develop	KA	3

The core topics listed in this section are typically covered in advanced, often elective courses. Effort should be made to include the CS Core topics in earlier required courses in the curriculum.

	Knowledge Area	Knowledge Units	CS Core	KA Core
AI	Artificial Intelligence	12	12	18
GIT	Graphics and Interactive Techniques	11	4	70
HCI	Human-Computer Interaction	6	8	16
SPD	Specialized Platform Development	8	4	
	Total		28	N/A

AI: Artificial Intelligence

KU	Topic	Skill Level	Core	Hours
AI- Introduction	<ol style="list-style-type: none"> 1. Overview of AI problems, Examples of successful recent AI applications 2. Definitions of agents with examples (e.g., reactive, deliberative) 3. What is intelligent behavior? <ol style="list-style-type: none"> a. The Turing test and its flaws b. Multimodal input and output c. Simulation of intelligent behavior d. Rational versus non-rational reasoning 	Explain	CS	2
	<ol style="list-style-type: none"> 4. Problem characteristics <ol style="list-style-type: none"> a. Fully versus partially observable b. Single versus multi-agent c. Deterministic versus stochastic d. Static versus dynamic e. Discrete versus continuous 5. Nature of agents <ol style="list-style-type: none"> a. Autonomous, semi-autonomous, mixed-initiative autonomy b. Reflexive, goal-based, and utility-based c. Decision making under uncertainty and with incomplete information 	Evaluate		

	<ul style="list-style-type: none"> d. The importance of perception and environmental interactions e. Learning-based agents f. Embodied agents <ul style="list-style-type: none"> i. sensors, dynamics, effectors 			
	6. AI Applications, growth, and Impact (economic, societal, ethics)	Explain		
AI-Introduction	<ul style="list-style-type: none"> 7. Practice identifying problem characteristics in example environments 8. Additional depth on nature of agents with examples 9. Additional depth on AI Applications, growth, and Impact (economic, societal, ethics) 	Evaluate	KA	1
AI-Search	<ul style="list-style-type: none"> 1. State space representation of a problem <ul style="list-style-type: none"> a. Specifying states, goals, and operators b. Factoring states into representations (hypothesis spaces) c. Problem solving by graph search <ul style="list-style-type: none"> i. e.g., Graphs as a space, and tree traversals as exploration of that space ii. Dynamic construction of the graph (you're not given it upfront) 	Explain	CS	5
AL-Foundationa	<ul style="list-style-type: none"> 2. Uninformed graph search for problem solving <ul style="list-style-type: none"> a. Breadth-first search b. Depth-first search <ul style="list-style-type: none"> i. With iterative deepening c. Uniform cost search 	Develop, Apply		
!	<ul style="list-style-type: none"> 3. Heuristic graph search for problem solving <ul style="list-style-type: none"> a. Heuristic construction and admissibility b. Hill-climbing c. Local minima and the search landscape <ul style="list-style-type: none"> i. Local vs global solutions d. Greedy best-first search e. A* search 	Develop, Apply		
	<ul style="list-style-type: none"> 4. Space and time complexities of graph search algorithms 	Evaluate		
AI-Search	<ul style="list-style-type: none"> 5. Bidirectional search 6. Beam search 7. Two-player adversarial games 	Develop,	KA	6

	<ul style="list-style-type: none"> a. Minimax search b. Alpha-beta pruning <ul style="list-style-type: none"> i. Ply cutoff <p>8. Implementation of A* search</p> <p>9. Constraint Satisfaction</p>	Apply		
AI-KRR	<ul style="list-style-type: none"> 1. Types of representations <ul style="list-style-type: none"> a. Symbolic, logical <ul style="list-style-type: none"> i. Creating a representation from a natural language problem statement b. Learned subsymbolic representations c. Graphical models (e.g., naive Bayes, Bayes net) 2. Review of probabilistic reasoning, Bayes theorem 	Explain	CS	2
	<ul style="list-style-type: none"> 3. Bayesian reasoning <ul style="list-style-type: none"> a. Bayesian inference 	Apply		
AI-KRR	<ul style="list-style-type: none"> 4. Random variables and probability distributions <ul style="list-style-type: none"> a. Axioms of probability b. Probabilistic inference c. Bayes' Rule (derivation) d. Bayesian inference (more complex examples) 5. Independence 6. Conditional Independence 7. Markov chains and Markov models 8. Utility and decision making 	Apply	KA	2
AI-ML	<ul style="list-style-type: none"> 1. Definition and examples of a broad variety of machine learning tasks <ul style="list-style-type: none"> a. Supervised learning <ul style="list-style-type: none"> i. Classification ii. Regression b. Reinforcement learning c. Unsupervised learning <ul style="list-style-type: none"> i. Clustering 2. Fundamental ideas: <ul style="list-style-type: none"> a. No free lunch theorem: no one learner can solve all problems; representational design decisions have consequences b. sources of error and undecidability in machine learning 		CS	4

	<ul style="list-style-type: none"> 3. A simple statistical-based supervised learning such as linear regression or decision trees <ul style="list-style-type: none"> a. Focus on how they work without going into mathematical or optimization details; enough to understand and use existing implementations correctly 4. The overfitting problem / controlling solution complexity (regularization, pruning – intuition only) <ul style="list-style-type: none"> a. The bias (underfitting) - variance (overfitting) tradeoff 5. Working with Data <ul style="list-style-type: none"> a. Data preprocessing <ul style="list-style-type: none"> i. Importance and pitfalls of preprocessing choices b. Handling missing values (imputing, flag-as-missing) <ul style="list-style-type: none"> i. Implications of imputing vs flag-as-missing c. Encoding categorical variables, encoding real-valued data d. Normalization/standardization e. Emphasis on real data, not textbook examples 6. Representations <ul style="list-style-type: none"> a. Hypothesis spaces and complexity b. Simple basis feature expansion, such as squaring univariate features c. Learned feature representations 7. Machine learning evaluation <ul style="list-style-type: none"> a. Separation of train, validation, and test sets b. Performance metrics for classifiers c. Estimation of test performance on held-out data d. Tuning the parameters of a machine learning model with a validation set e. Importance of understanding what your model is actually doing, where its pitfalls/shortcomings are, and the implications of its decisions 8. Basic neural networks <ul style="list-style-type: none"> a. Fundamentals of understanding how neural networks work and their training process, without details of the calculations 	<p>Apply, Develop, Evaluate</p>		
--	---	---	--	--

	<ul style="list-style-type: none"> b. Basic introduction to generative neural networks (large language models, etc.) 			
	<ul style="list-style-type: none"> 9. Ethics for Machine Learning <ul style="list-style-type: none"> a. Focus on real data, real scenarios, and case studies b. Dataset/algorithmic/evaluation bias and unintended consequences 	Explain, Evaluate		
AI-ML	<ul style="list-style-type: none"> 10. Formulation of simple machine learning as an optimization problem, such as least squares linear regression or logistic regression <ul style="list-style-type: none"> a. Objective function b. Gradient descent c. Regularization to avoid overfitting (mathematical formulation) 11. Ensembles of models <ul style="list-style-type: none"> a. Simple weighted majority combination 12. Deep learning <ul style="list-style-type: none"> a. Deep feed-forward networks (intuition only, no mathematics) b. Convolutional neural networks (intuition only, no mathematics) c. Visualization of learned feature representations from deep nets d. Other architectures (generative NN, recurrent NN, transformers, etc.) 13. Performance evaluation <ul style="list-style-type: none"> a. Other metrics for classification (e.g., error, precision, recall) b. Performance metrics for regressors c. Confusion matrix d. Cross-validation <ul style="list-style-type: none"> i. Parameter tuning (grid/random search, via cross-validation) 14. Overview of reinforcement learning methods 15. Two or more applications of machine learning algorithms <ul style="list-style-type: none"> a. E.g., medicine and health, economics, vision, natural language, robotics, game play 	Apply, Develop, Evaluate	KA	6
	<ul style="list-style-type: none"> 16. Ethics for Machine Learning <ul style="list-style-type: none"> a. Continued focus on real data, real scenarios, and case studies 	Explain, Evaluate		

	<ul style="list-style-type: none"> b. Privacy c. Fairness d. Intellectual property e. Explainability 			
AI-SEP	<ol style="list-style-type: none"> 1. At least one application of AI to a specific problem and field, such as medicine, health, sustainability, social media, economics, education, robotics, etc. (choose at least one for the CS Core). <ul style="list-style-type: none"> a. Formulating and evaluating a specific application as an AI problem <ul style="list-style-type: none"> i. How to deal with underspecified or ill-posed problems b. Data availability/scarcity and cleanliness <ul style="list-style-type: none"> i. Basic data cleaning and preprocessing ii. Data set bias c. Algorithmic bias d. Evaluation bias e. Assessment of societal implications of the application 2. Deployed deep generative models <ul style="list-style-type: none"> a. High-level overview of deep image generative models (e.g. as of 2023, DALL-E, Midjourney, Stable Diffusion, etc.), their uses, and their shortcomings/pitfalls. b. High-level overview of large language models (e.g. as of 2023, ChatGPT, Bard, etc.), their uses, and their shortcomings/pitfalls. 3. Overview of societal impact of AI <ul style="list-style-type: none"> a. Ethics b. Fairness c. Trust / explainability d. Privacy and usage of training data e. Human autonomy and oversight/regulations/legal requirements f. Sustainability 	Explain, Evaluate	CS	3
AI-SEP	<ol style="list-style-type: none"> 4. One or more additional applications of AI to a broad set of problems and diverse fields, such as medicine, health, sustainability, social media, economics, education, robotics, etc. (choose a different area from that chosen for the CS Core). <ul style="list-style-type: none"> a. Formulating and evaluating a specific application as an AI problem 	Explain, Evaluate	KA	3

	<ul style="list-style-type: none"> i. How to deal with underspecified or ill-posed problems b. Data availability/scarcity and cleanliness <ul style="list-style-type: none"> i. Basic data cleaning and preprocessing ii. Data set bias c. Algorithmic bias d. Evaluation bias e. Assessment of societal implications of the application <p>5. Additional depth on deployed deep generative models</p> <ul style="list-style-type: none"> a. Introduction to how deep image generative models (e.g. as of 2023, DALL-E, Midjourney, Stable Diffusion, etc.) work, including discussion of attention b. Introduction to how large language models (e.g. as of 2023, ChatGPT, Bard, etc.) work, including discussion of attention c. Idea of foundational models, how to use them, and the benefits / issues with training them from big data <p>6. Analysis and discussion of the societal impact of AI</p> <ul style="list-style-type: none"> a. Ethics b. Fairness c. Trust / explainability d. Privacy and usage of training data e. Human autonomy and oversight/regulations/legal requirements f. Sustainability 			
--	---	--	--	--

GIT: Graphics and Interactive Techniques

KU	Topic	Skill Level	Core	Hours
GIT-Fundamentals	<ul style="list-style-type: none"> 1. Uses and potential risks 2. Graphic output 3. Human vision system 4. Standard image formats 5. Digitization of analog data 6. Color Models 7. Tradeoffs between storing image data and recomputing image data 	Explain	CS	4

	8. Spatialization 9. Animation			
	10. Applied computer graphics 11. Display characteristics	Develop	KA	3
GIT-Visualization	<ol style="list-style-type: none"> 1. Data Visualization and Information Visualization 2. Visualization of: <ol style="list-style-type: none"> a. 2D/3D scalar fields b. Vector fields and flow data c. Time-varying data d. High-dimensional data e. Non-spatial data 3. Visualization techniques (e.g., color mapping, isosurfaces, dimension reduction, parallel coordinates, multi-variate, tree/graph-structured, text) 4. Direct volume data rendering (e.g., ray-casting, transfer functions, segmentation) 5. Common data formats (e.g., HDF, netCDF, geotiff, GeoJSON, shape files, raw binary, CSV, ASCII to parse) 6. Common visualization software and libraries (e.g., R, Processing, D3.js, GIS, Matlab, IDL, Python) 7. Perceptual and cognitive foundations that drive visual abstractions 8. Visualization design 9. Evaluation of visualization methods and applications 10. Visualization bias 11. Applications of visualization 	Explain	KA	6
GIT-Rendering	<ol style="list-style-type: none"> 1. Object and scene modeling 2. Camera and projection modeling 3. Light models and radiometry 4. Rendering 	Explain	KA	6
GIT-Modeling	<ol style="list-style-type: none"> 1. Basic geometric operations 2. Surface representation/model 3. Volumetric representation/model 4. Procedural representation/model 5. Multi-resolution modeling 6. Reconstruction 	Explain	KA	6
GIT-Shading	<ol style="list-style-type: none"> 1. Solutions and approximations to the rendering equation 	Explain	KA	6

	<ol style="list-style-type: none"> 2. Time (motion blur), lens position (focus), and continuous frequency (color) and their impact on rendering 3. Shadow mapping 4. Occlusion culling 5. Bidirectional Scattering Distribution function (BSDF) theory and microfacets 6. Subsurface scattering 7. Area light sources 8. Hierarchical depth buffering 9. Image-based rendering 10. Non-photorealistic rendering 11. GPU architecture 12. Human visual systems including adaptation to light, sensitivity to noise, and flicker fusion 			
GIT-Animation	<ol style="list-style-type: none"> 1. Principles of Animation 2. Types of animation 3. Key-frame animation 4. Forward and inverse kinematics 5. Skinning algorithms 6. Motion capture 	Explain	KA	6
GIT-Simulation	<ol style="list-style-type: none"> 1. Collision detection and response 2. Procedural animation using noise 3. Particle systems 4. Grid-based fluids 5. Heightfields 6. Rule-based systems (e.g., L-systems, space-colonizing systems, Game of Life) 	Explain	KA	6
GIT-Immersion	<ol style="list-style-type: none"> 1. Immersion levels (i.e., Virtual Reality (VR), Augmented Reality (AR), and Mixed Reality (MR)) 2. The definition and distinction of immersion and presence 3. 360 Video 4. Stereoscopic display 5. Viewer tracking 6. Time-critical rendering to achieve optimal motion to photon (MTP) latency <ol style="list-style-type: none"> a. Branching movies 7. Distributed VR, collaboration over computer network 8. Presence and factors that impact level of immersion 9. 3D interaction 	Explain	KA	6

	10. Applications in medicine, simulation, training, and visualization 11. Safety in immersive applications			
GIT-Interaction	1. Event Driven Programming 2. Graphical User Interface (Single Channel) 3. Accessibility	Apply	KA	4
GIT-Image	1. Morphological operations 2. Color histograms 3. Image enhancement 4. Image restoration 5. Image coding 6. Connections to deep learning	Explain	KA	6
GIT-Physical	1. Interaction with the physical world 2. Connection to physical artifacts 3. Internet of Things	Explain	KA	6
GIT-SEP	1. Accessibility 2. Ethics/privacy 3. Intellectual Property law as it relates to computer graphics and interactive techniques 4. DEIA, current and past contributors to the field	Evaluate	CS KA	1 3

HCI: Human-Computer Interaction

KU	Topic	Skill Level	Core	Hours
HCI-User	1. User-centered design and evaluation methods a. "You are not the users", b. user needs-finding c. formative studies d. interviews e. surveys f. usability tests	Explain, Evaluate , Develop	CS	2
HCI-User	2. User-centered design and evaluation methods a. personas/persona spectrum b. user stories/storytelling and techniques for gathering stories c. empathy maps	Explain, Evaluate , Develop	KA	5

	<ul style="list-style-type: none"> d. needs assessment (techniques for uncovering needs and gathering requirements - e.g., interviews, surveys, ethnographic and contextual enquiry) e. journey maps f. evaluating the design g. interfacing with stakeholders, as a team h. risks associated with physical, distributed, hybrid and virtual teams <p>3. Physical and cognitive characteristics of the user</p> <ul style="list-style-type: none"> a. physical capabilities that inform interaction design (e.g., color perception, ergonomics) b. cognitive models that inform interaction design (e.g., attention, perception and recognition, movement, memory) c. topics in social/behavioral psychology (e.g., cognitive biases, change blindness) <p>4. Designing for diverse user populations</p> <ul style="list-style-type: none"> a. how differences (e.g., in race, ability, age, gender, culture, experience, and education) impact user experiences and needs b. internationalization, other cultures, and cross-cultural design c. designing for users from other cultures d. cross-cultural design e. challenges to effective design evaluation (e.g., sampling, generalization; disability and disabled experiences) f. universal design <p>5. Collaboration and communication</p> <ul style="list-style-type: none"> a. understanding the user in a multi-user context b. synchronous group communication (e.g., chat rooms, conferencing, online games) c. asynchronous group communication (e.g., email, forums, social networks) d. social media, social computing, and social network analysis e. online collaboration f. social coordination and online communities g. avatars, characters, and virtual worlds 			
HCL- Accountability	<p>1. Design impact</p> <ul style="list-style-type: none"> a. Sustainability b. Inclusivity 	<p>Explain, Apply,</p>	<p>CS</p>	<p>2</p>

	<ul style="list-style-type: none"> c. Safety, security and privacy d. Harm and disparate impact <p>2. Ethics in design methods and solutions</p> <ul style="list-style-type: none"> a. the role of artificial intelligence b. responsibilities for considering stakeholder impact and human factors c. the role of design to meet user needs. <p>3. Requirements in design</p> <ul style="list-style-type: none"> a. ownership responsibility b. legal frameworks and compliance requirements c. consideration beyond immediate user needs including via iterative reconstruction of problem analysis and “digital well-being” features 	Evaluate , Develop		
HCI- Accountability	<p>4. Value-sensitive design</p> <ul style="list-style-type: none"> a. identify stakeholders b. determine and include diverse stakeholder values and value systems <p>5. Persuasion through design</p> <ul style="list-style-type: none"> a. assessing persuasive content of a design b. persuasion as a design goal 	Explain, Apply, Evaluate , Develop	KA	2
HCI- Accessibility	<p>1. Background</p> <ul style="list-style-type: none"> a. societal and legal support for and obligations to people with disabilities b. accessible design benefits everyone <p>2. Techniques</p> <ul style="list-style-type: none"> a. accessibility standards (e.g., Web Content Accessibility Guidelines) <p>3. Technologies</p> <ul style="list-style-type: none"> a. features and products that enable accessibility and support inclusive development by designers and engineers <p>4. IDFs (Inclusive Design Frameworks)</p> <ul style="list-style-type: none"> a. recognizing differences <p>5. Universal design</p>	Explain, Apply, Evaluate , Develop	CS	2
HCI- Accessibility	<p>6. Background</p> <ul style="list-style-type: none"> a. Demographics and populations b. International perspectives c. Attitudes towards people with disabilities <p>7. Techniques</p> <ul style="list-style-type: none"> a. UX (user experience) design and research 	Explain, Apply, Evaluate , Develop	KA	2

	<ul style="list-style-type: none"> b. software engineering practices that enable inclusion and accessibility. <p>8. Technologies</p> <ul style="list-style-type: none"> a. examples of accessibility-enabling features, such as conformance to screen readers <p>9. Inclusive Design Frameworks</p> <ul style="list-style-type: none"> a. creating inclusive processes, such as participatory design; designing for larger impact b. designing for larger impact 			
HCI-Evaluation	<p>1. Methods for evaluation with users</p> <ul style="list-style-type: none"> a. formative (e.g. needs-finding and exploratory analysis) and summative assessment (e.g. functionality and usability testing) b. elements to evaluate (e.g., utility, efficiency, learnability, user satisfaction, affective elements such as pleasure and engagement) c. understanding ethical approval requirements before engaging in user research 	Explain, Apply, Evaluate, Develop	CS	1
HCI-Evaluation	<p>2. Methods for evaluation with users</p> <ul style="list-style-type: none"> a. qualitative methods (qualitative coding and thematic analysis) b. quantitative methods (statistical tests) c. mixed methods (e.g., observation, think-aloud, interview, survey, experiment) d. presentation requirements (e.g., reports, personas) e. user-centered testing f. heuristic evaluation g. challenges and shortcomings to effective evaluation (e.g., sampling, generalization) <p>3. Study planning</p> <ul style="list-style-type: none"> a. how to set study goals b. hypothesis design c. approvals from Institutional Research Boards and ethics committees d. how to pre-register a study e. within-subjects vs between-subjects design <p>4. Implications and impacts of design with respect to</p> <ul style="list-style-type: none"> a. the environment b. material c. society d. security 	Explain, Apply, Evaluate, Develop	KA	2

	<ul style="list-style-type: none"> e. privacy f. ethics g. broader impacts 			
HCI-Design	<ol style="list-style-type: none"> 1. Prototyping techniques and tools <ul style="list-style-type: none"> a. low-fidelity prototyping b. rapid prototyping c. throw-away prototyping d. granularity of prototyping 2. Design patterns <ul style="list-style-type: none"> a. iterative design b. universal design c. interaction design (e.g., data-driven design, event-driven design) 3. Design constraints <ul style="list-style-type: none"> a. platforms b. devices c. resources d. balance among usability, security and privacy 	Explain, Apply, Evaluate, Develop	CS	1
HCI-Design	<ol style="list-style-type: none"> 4. Design patterns and guidelines <ul style="list-style-type: none"> a. software architecture patterns b. cross-platform design c. synchronization 5. Design processes <ul style="list-style-type: none"> a. participatory design b. co-design c. double-diamond d. convergence and divergence 6. Interaction techniques <ul style="list-style-type: none"> a. input and output vectors (e.g., gesture, pose, touch, voice, force) b. graphical user interfaces c. controllers d. haptics e. hardware design f. error handling 7. Visual UI design <ul style="list-style-type: none"> a. color b. layout c. gestalt principles 	Explain, Apply, Evaluate, Develop	KA	5

HCI-SEP	<ol style="list-style-type: none"> 1. Universal and user-centered design 2. Accountability 3. Accessibility and inclusive design 4. Evaluating the design 5. System design 	Explain, Apply, Evaluate, Develop	CS	Share d with SEP
HCI-SEP	<ol style="list-style-type: none"> 6. Participatory and inclusive design processes 	Explain, Apply, Evaluate, Develop	KA	Share d with SEP

SPD: Specialized Platform Development

KU	Topic	Skill Level	Core	Hours
SPD-Common	<ol style="list-style-type: none"> 1. Overview of development platforms (i.e., web, mobile, game, robotics, embedded, and Interactive). <ol style="list-style-type: none"> a. Input/sensors/control devices/haptic devices. b. Resource constraints: <ol style="list-style-type: none"> i. Computational ii. Data storage iii. Memory iv. Communication c. Requirements - security, uptime availability, fault tolerance. d. Output/actuators/haptic devices. 2. Programming via platform-specific Application Programming Interface (API) vs traditional application construction. 3. Overview of platform Languages (e.g., Python, Swift, Lua, Kotlin). 4. Programming under platform constraints and requirements (e.g., available development tools, development, security considerations). 5. Techniques for learning and mastering a platform-specific programming language. 	Apply	CS	3
SPD-Web		Apply	KA	5

	<ol style="list-style-type: none"> 1. Web programming languages (e.g., HTML5, JavaScript, PHP, CSS) 2. Web platforms, frameworks, or meta-frameworks: <ol style="list-style-type: none"> a. Cloud services b. API, Web Components 3. Software as a Service (SaaS). 4. Web standards such as document object model, accessibility. 5. Security and Privacy Considerations. 			
SPD-Mobile	<ol style="list-style-type: none"> 1. Development with: <ol style="list-style-type: none"> a. Mobile programming languages. b. Mobile programming environments. 2. Mobile platform constraints: <ol style="list-style-type: none"> a. User interface design. b. Security. 3. Access: <ol style="list-style-type: none"> a. Accessing data through APIs. b. Designing API endpoints for mobile apps: pitfalls and design considerations. c. Network and the web interfaces. 	Apply	KA	3
SPD-Robot	<ol style="list-style-type: none"> 1. Types of robotic platforms and devices. 2. Sensors, embedded computation, and effectors (actuators). 3. Robot-specific languages and libraries. 4. Robotic software architecture (e.g., using the Robot Operating System (ROS)). 5. Robotic platform constraints and design considerations. 6. Interconnections with physical or simulated systems. 7. Robotic Algorithms 8. Forward kinematics. <ol style="list-style-type: none"> a. Inverse kinematics. b. Dynamics. c. Navigation and path planning. d. Grasping and manipulation. 9. Safety and interaction considerations. 	Apply	KA	4
SPD-Embedded	<ol style="list-style-type: none"> 1. Introduction to the unique characteristics of embedded systems: <ol style="list-style-type: none"> a. Real-time vs soft real-time and non-real-time systems. 	Apply	KA	4

	<ul style="list-style-type: none"> b. Resource constraints, such as memory profiles and deadlines. c. API for custom architectures: d. GPU technology. (See also: AR-Heterogeneity, GIT-Shading) e. Field Programmable Gate Arrays (FPGA). f. Cross-platform systems. <p>2. Embedded Systems:</p> <ul style="list-style-type: none"> a. Microcontrollers. b. Interrupts and feedback. c. Interrupt handlers in high-level languages. d. Hard and soft interrupts and trap-exits. e. Interacting with hardware, actuators, and sensors. f. Energy efficiency. g. Loosely timed coding and synchronization. h. Software adapters. <p>3. Embedded programming.</p> <p>4. Hard real-time systems vs soft real-time systems:</p> <ul style="list-style-type: none"> a. Timeliness. b. Time synchronization/scheduling. c. Prioritization. d. Latency. e. Compute jitter. <p>5. Real-time resource management.</p> <p>6. Memory management:</p> <ul style="list-style-type: none"> a. Mapping programming construct (variable) to a memory location. b. Shared memory. c. Manual memory management. d. Garbage collection. e. Safety considerations and safety analysis. (See also: SEP-Context, SEP-Professional) <p>7. Sensors and actuators.</p> <p>8. Analysis and verification.</p> <p>9. Application design.</p>			
<p>SPD-Game</p>	<p>1. Historical and contemporary platforms for games:</p> <ul style="list-style-type: none"> a. Evolution of Game Platforms (e.g., Brown Box to Metaverse and beyond; Improvement in Computing Architectures (CPU and GPU); Platform Convergence and Mobility). 	<p>Apply</p>	<p>KA</p>	<p>4</p>

	<ul style="list-style-type: none"> b. Typical Game Platforms (e.g., Personal Computer; Home Console; Handheld Console; Arcade Machine; Interactive Television; Mobile Phone; Tablet; Integrated Head-Mounted Display; Immersive Installations and Simulators; Internet of Things enabled Devices; CAVE Systems; Web Browsers; Cloud-based Streaming Systems). c. Characteristics and Constraints of Different Game Platforms (e.g., Features (local storage, internetworking, peripherals); Runtime performance (GPU/CPU frequency, number of cores); Chipsets (physics processing units, vector co-processors); Expansion Bandwidth (PCIe); Network throughput (Ethernet); Memory types and capacities (DDR/GDDR); Maximum stack depth; Power consumption; Thermal design; Endian) d. Typical Sensors, Controllers, and Actuators (e.g., distinctive control system designs—peripherals (mouse, keypad, joystick), game controllers, wearables, interactive surfaces; electronics and bespoke hardware; computer vision, inside-out tracking, and outside-in tracking; IoT-enabled electronics and I/O. e. eSports Ecosystems (e.g., evolution of gameplay across platforms; games and eSports; game events such as LAN/arcade tournaments and international events such as the Olympic eSports Series; streamed media and spectatorship; multimedia technologies and broadcast management; professional play; data and machine learning for coaching and training) <p>2. Real-time Simulation and Rendering Systems</p> <ul style="list-style-type: none"> a. CPU and GPU architectures: (e.g., Flynn’s taxonomy; parallelization; instruction sets; standard components—graphics compute array, graphics memory controller, video graphics array basic input/output system; 			
--	---	--	--	--

	<p>bus interface; power management unit; video processing unit; display interface).</p> <ul style="list-style-type: none"> b. Pipelines for physical simulations and graphical rendering: (e.g., tile-based, immediate-mode). c. Common Contexts for Algorithms, Data Structures, and Mathematical Functions (e.g., game loops; spatial partitioning, viewport culling, and level of detail; collision detection and resolution; physical simulation; behavior for intelligent agents; procedural content generation). d. Media representations (e.g., I/O, and computation techniques for virtual worlds: audio; music; sprites; models and textures; text; dialogue; multimedia (e.g., olfaction, tactile). <p>3. Game Development Tools and Techniques:</p> <ul style="list-style-type: none"> a. Programming Languages (e.g., C++; C#; Lua; Python; JavaScript). b. Shader Languages (e.g., HLSL, GLSL; ShaderGraph). c. Graphics Libraries and APIs (e.g., DirectX; SDL; OpenGL; Metal; Vulkan; WebGL). d. Common Development Tools and Environments (e.g., IDEs; Debuggers; Profilers; Version Control Systems including those handling binary assets; Development Kits and Production/Consumer Kits; Emulators). <p>4. Game Engines</p> <ul style="list-style-type: none"> a. Open Game Engines (e.g., Unreal; Unity; Godot; CryEngine; Phyre; Source 2; Pygame and Ren'Py; Phaser; Twine; SpringRTS) b. Techniques (e.g., Ideation, Prototyping, Iterative Design and Implementation, Compiling Executable Builds, Development Operations and Quality Assurance—Play Testing and Technical Testing, Profiling; Optimization, Porting; Internationalization and Localization, Networking). <p>5. Game Design</p>			
--	--	--	--	--

	<ul style="list-style-type: none"> a. Vocabulary (e.g., game definitions; mechanics-dynamics-aesthetics model; industry terminology; experience design; models of experience and emotion). b. Design Thinking and User-Centered Experience Design (e.g., methods of designing games; iteration, incrementing, and the double-diamond; phases of pre- and post-production; quality assurance, including alpha and beta testing; stakeholder and customer involvement; community management). c. Genres (e.g., adventure; walking simulator; first-person shooter; real-time strategy; multiplayer online battle arena (MOBA); role-playing game (rpg)). d. Audiences and Player Taxonomies (e.g., people who play games; diversity and broadening participation; pleasures, player types, and preferences; Bartle, yee). e. Proliferation of digital game technologies to domains beyond entertainment (e.g., Education and Training; Serious Games; Virtual Production; eSports; Gamification; Immersive Experience Design; Creative Industry Practice; Artistic Practice; Procedural Rhetoric). 			
--	---	--	--	--

Crosscutting Core Topics

The core topics in Society, Ethics, and the Profession (SEP) and Mathematical and Statistical Foundations (MSF) may be covered across the curriculum or in dedicated courses and benefit all the competency areas.

	Knowledge Area	Knowledge Units	CS Core	KA Core
SEP	Society, Ethics, and the Profession	11	18	14
MSF	Mathematical and Statistical Foundations	5	55	145
	Total		73	159

SEP: Society, Ethics, and the Profession

KU	Topic	Skill Level	Core	Hours
SEP-Context	<ol style="list-style-type: none"> 1. Social implications (e.g., political, and cultural ideologies) in a hyper-networked world where the capabilities and impact of social media, artificial intelligence and computing in general are rapidly evolving 2. Impact of computing applications (e.g., social media, artificial intelligence applications) on individual well-being, and safety of all kinds (e.g., physical, emotional, economic) 3. Consequences of involving computing technologies, particularly artificial intelligence, biometric technologies, and algorithmic decision-making systems, in civic life (e.g., facial recognition technology, biometric tags, resource distribution algorithms, policing software) and how human agency and oversight is crucial 4. How deficits in diversity and accessibility in computing affect society and what steps can be taken to improve equity in computing 	<p>Evaluate</p> <p>Evaluate</p> <p>Evaluate</p> <p>Evaluate</p>	CS	3
SEP-Context	<ol style="list-style-type: none"> 5. Growth and control of the internet, data, computing, and artificial intelligence 6. Often referred to as the digital divide, differences in access to digital technology resources and its resulting ramifications for gender, class, ethnicity, geography, and/or developing countries, including consideration of responsibility to those who might be less wealthy, under threat, or who would struggle to have their voices heard. 7. Accessibility issues, including legal requirements such as Web Content Accessibility Guidelines (www.w3.org/TR/WCAG21) 8. Context-aware computing 	<p>Explain</p> <p>Evaluate</p> <p>Explain</p> <p>Explain</p>	KA	2
SEP-Ethical-Analysis	<ol style="list-style-type: none"> 1. Avoiding fallacies and misrepresentation in argumentation 2. Ethical theories and decision-making (philosophical and social frameworks, e.g. [3]) 	<p>Apply</p> <p>Apply</p> <p>Evaluate</p>	CS	2

	<ol style="list-style-type: none"> 3. Recognition of the role culture plays in our understanding, adoption, design, and use of computing technology 4. Why ethics is important in computing, and how ethics is similar to, and different from, laws and social norms 	Explain		
SEP-Ethical-Analysis	<ol style="list-style-type: none"> 5. Professional checklists 6. Evaluation rubrics 7. Stakeholder analysis 8. Standpoint theory 9. Introduction to ethical frameworks (e.g., consequentialism such as utilitarianism, non-consequentialism such as duty, rights, or justice, agent-centered such as virtue or feminism, contractarianism, ethics of care) and their use for analyzing an ethical dilemma 	<p>Develop Develop Develop Apply Explain</p>	KA	1
SEP-Professional-Ethics	<ol style="list-style-type: none"> 1. Community values and the laws by which we live 2. The nature of being a professional including care, attention, discipline, fiduciary responsibility, and mentoring 3. Keeping up to date as a computing professional in terms of familiarity, tools, skills, legal and professional frameworks as well as the ability and responsibility to self-assess and progress in the computing field 4. Professional certification, codes of ethics, conduct, and practice, such as the ACM, IEEE, AAAI, and other international societies 5. Accountability, responsibility, and liability (e.g., software correctness, reliability and safety, warranty, negligence, strict liability, ethical approaches to security vulnerability disclosures) including whether a product/service should be built, not just doing so because it is technically possible. 6. Introduction to theories describing the human creation and use of technology including instrumentalism, sociology of technological systems, disability justice, neutrality thesis, pragmatism, and decolonial models, including developing and using technology to right wrongs and do good 	<p>Evaluate Apply</p> <p>Evaluate</p> <p>Evaluate</p> <p>Apply</p> <p>Explain</p>	CS	2

	7. Strategies for recognizing and reporting designs, systems, software, and professional conduct (or their outcomes) that may violate law or professional codes of ethics	Apply		
SEP-Professional-Ethics	8. The role of the computing professional and professional societies in public policy 9. Maintaining awareness of consequences 10. Ethical dissent and whistleblowing 11. The relationship between regional culture and ethical dilemmas 12. Dealing with harassment and discrimination 13. Forms of professional credentialing 14. Ergonomics and healthy computing environments 15. Time-to-market and cost considerations versus quality professional standards	Explain Apply Explain Evaluate Explain Explain Explain Explain	KA	2
SEP-IP	1. Intellectual property rights 2. Intangible digital intellectual property (IDIP) 3. Legal foundations for intellectual property protection 4. Common software licenses (e.g., MIT, GPL and its variants, Apache, Mozilla, Creative Commons) 5. Plagiarism and authorship	Explain Explain Evaluate Evaluate Explain	CS	1
SEP-IP	6. Philosophical foundations of intellectual property 7. Forms of intellectual property (e.g., copyrights, patents, trade secrets, trademarks) and the rights they protect 8. Limitations on copyright protections, including fair use and the first sale doctrine 9. Intellectual property laws and treaties that impact the enforcement of copyrights 10. Software piracy and technical methods for enforcing intellectual property rights, such as digital rights management and closed source software as a trade secret 11. Moral and legal foundations of the open source movement 12. Systems that use others' data (e.g., large language models)	Explain Explain Explain Explain Explain Explain	KA	1
SEP-Privacy	1. Privacy implications of widespread data collection including but not limited to transactional	Explain	CS	2

	<p>databases, data warehouses, surveillance systems, cloud computing, and artificial intelligence</p> <p>2. Conceptions of anonymity, pseudonymity, and identity</p> <p>3. Technology-based solutions for privacy protection (e.g., end-to-end encryption and differential privacy)</p> <p>4. Civil liberties, privacy rights, and cultural differences</p>	Evaluate		
SEP-Privacy	<p>5. Philosophical and legal conceptions of the nature of privacy including the right to privacy</p> <p>6. Legal foundations of privacy protection in relevant jurisdictions (e.g., GDPR in the EU)</p> <p>7. Privacy legislation in areas of practice (e.g., HIPAA in the US, AI Act in the EU)</p> <p>8. Basic Principles of human-subjects research and principles beyond what the law requires (e.g., Belmont Report, UN Universal Declaration on Human Rights and how this relates to technology)</p> <p>9. Freedom of expression and its limitations</p> <p>10. User-generated content, content moderation, and liability</p>	Evaluate	KA	1
SEP-Communication	<p>1. Oral, written, and electronic team and group communication</p> <p>2. Technical communication materials (e.g., source code, and documentation, tutorials, reference materials, API documentation)</p> <p>3. Communicating with different stakeholders such as customers, leadership, or the general public</p> <p>4. Team collaboration (including tools) and conflict resolution</p> <p>5. Accessibility and inclusivity requirements for addressing professional audiences</p> <p>6. Cultural competence in communication including considering the impact of difference in natural language</p>	Apply	CS	2
SEP-Communication	<p>7. Tradeoffs in competing factors that affect communication channels and choices</p> <p>8. Communicating to solve problems or make recommendations in the workplace, such as</p>	Evaluate	KA	1

	raising ethical concerns or addressing accessibility issues			
SEP-Sustainability	<ol style="list-style-type: none"> 1. Environmental, social, and cultural impacts of implementation decisions (e.g., sustainability goals, algorithmic bias/outcomes, economic viability, and resource consumption) 2. Local/regional/global social and environmental impacts of computing systems and their use (e.g., carbon footprints, resource usage, e-waste) in hardware (e.g., e-waste, data centers, rare element and resource utilization, recycling) and software (e.g., cloud-based services, blockchain, AI model training and use), not neglecting the impact of everyday use such as hardware (cheap hardware replaced frequently) and software (web-browsing, email, and other services with hidden/remote computational demands). 3. Guidelines for sustainable design standards 	Evaluate Evaluate Develop	CS	1
SEP-Sustainability	<ol style="list-style-type: none"> 4. Systemic effects of complex computing technologies and phenomena (e.g., generative AI, data centers, social media, offshoring, remote work) 5. Pervasive computing: Information processing that has been integrated into everyday objects and activities, such as smart energy systems, social networking, and feedback systems to promote sustainable behavior, transportation, environmental monitoring, citizen science and activism 6. How the sustainability of software systems is interdependent with social systems, including the knowledge and skills of its users, organizational processes and policies, and its societal context (e.g., market forces, government policies) 	Evaluate Evaluate Explain	KA	1
SEP-History	<ol style="list-style-type: none"> 1. The history of computing: hardware, software, and human/organizational 2. The role of history in the present including within different social contexts, and the relevance of this history on the future 	Explain Evaluate	CS	1

SEP-History	<ol style="list-style-type: none"> 3. Age I (Pre-digital): Ancient analog computing (Stonehenge, Antikythera mechanism, Salisbury Cathedral clock, etc.), human-calculated number tables, Euclid, Lovelace, Babbage, Gödel, Church, Turing, pre-electronic (electro-mechanical and mechanical) hardware 4. Age II (Early modern computing): ENIAC, UNIVAC, Bombes (Bletchley Park and codebreakers), computer companies (e.g., IBM), mainframes, etc. 5. Age III (PC era): PCs, modern computer hardware and software, Moore's Law 6. Age IV (Internet): Networking, internet architecture, browsers and their evolution, standards, born-on-the-internet companies, and services (e.g., Google, Amazon, Microsoft, etc.), distributed computing 7. Age V (Mobile & Cloud): Mobile computing and smartphones, cloud computing and models thereof (e.g., SaaS), remote servers, security and privacy, social media 8. Age VI (AI): Decision making systems, recommender systems, generative AI and other machine learning driven tools and technologies 	<p>Explain</p> <p>Explain</p> <p>Explain</p> <p>Explain</p> <p>Explain</p> <p>Explain</p>	<p>KA</p>	<p>1</p>
SEP-Economies	<ol style="list-style-type: none"> 1. Economic models: regulated and unregulated, monopolies, network effects, and open market; knowledge and attention economies 2. Pricing and deployment strategies: planned obsolescence, subscriptions, freemium, software licensing, open-source, free software, adware 3. Impacts of differences in access to computing resources, and the effect of skilled labor supply and demand on the quality of computing products 4. Automation, AI, and their effects on job markets, developers, and users 5. Ethical concerns surrounding the attention economy and other economies of computing (e.g. informed consent, data collection, use of verbose legalese in user agreements) 	<p>Explain</p> <p>Explain</p> <p>Evaluate</p> <p>Explain</p> <p>Evaluate</p>	<p>KA</p>	<p>1</p>
SEP-Security	<ol style="list-style-type: none"> 1. Computer crimes, legal redress for computer criminals and impact on victims and society 	<p>Explain</p>	<p>CS</p>	<p>2</p>

	<ol style="list-style-type: none"> 2. Social engineering, computing-enabled fraud, identity theft and recovery from these 3. Cyber terrorism, criminal hacking, and hacktivism 4. Malware, viruses, worms 5. Attacks on critical infrastructure such as electrical grids and pipelines 6. Non-technical fundamentals of security (e.g., human engineering, policy, confidentiality) 	<p>Explain</p> <p>Explain</p> <p>Explain</p> <p>Explain</p>		
SEP-Security	<ol style="list-style-type: none"> 7. Benefits and challenges of existing and proposed computer crime laws 8. Security policies and the challenges of change and compliance 9. Responsibility for security throughout the computing life cycle 10. International and local laws and how they intersect 	<p>Evaluate</p> <p>Explain</p> <p>Explain</p> <p>Explain</p>	KA	1
SEP-DEIA	<ol style="list-style-type: none"> 1. How identity impacts and is impacted by computing technologies and environments (academic and professional) 2. The benefits of diverse development teams and the impacts of teams that are not diverse 3. Inclusive language and charged terminology, and why their use matters 4. Inclusive behaviors and why they matter 5. Designing and developing technology with accessibility in mind 6. How computing professionals can influence and impact diversity, equity, inclusion and accessibility, including but not only through the software they create 	<p>Explain</p> <p>Explain</p> <p>Apply</p> <p>Explain</p> <p>Explain</p> <p>Explain</p>	CS	2
SEP-DEIA	<ol style="list-style-type: none"> 7. Experts and their practices that who reflect the identities of the classroom and the world through practical DEIA principles 8. Historic marginalization due to systemic social mechanisms, technological supremacy and global infrastructure challenges to diversity, equity, inclusion and accessibility 9. Cross-cultural differences in, and needs for, diversity, equity, inclusion and accessibility 	<p>Evaluate</p> <p>Explain</p> <p>Explain</p>	KA	2

MSF: Mathematical and Statistical Foundations

KU	Topic	Skill Level	Core	Hours
MSF-Discrete	<ol style="list-style-type: none"> 1. Sets, relations, functions, cardinality 2. Recursive mathematical definitions 3. Proof techniques (induction, proof by contradiction) 4. Permutations, combinations, counting, pigeonhole principle 5. Modular arithmetic 6. Logic: truth tables, connectives (operators), inference rules, formulas, normal forms, simple predicate logic 7. Graphs: basic definitions 8. Order notation 	Apply, Develop, Explain	CS/KA	29-40
MSF-Probability	<ol style="list-style-type: none"> 1. Basic notions: sample spaces, events, probability, conditional probability, Bayes' rule 2. Discrete random variables and distributions 3. Continuous random variables and distributions 4. Expectation, variance, law of large numbers, central limit theorem 5. Conditional distributions and expectation 6. Applications to computing, the difference between probability and statistics (as subjects) 	CS-Core: Apply KA-Core: Apply, Develop, Explain	CS/KA	11-40
MSF-Statistics	<ol style="list-style-type: none"> 1. Basic definitions and concepts: populations, samples, measures of central tendency, variance 2. Univariate data: point estimation, confidence intervals 	Develop	CS	10
MSF-Statistics	<ol style="list-style-type: none"> 3. Multivariate data: estimation, correlation, regression 4. Data transformation: dimension reduction, smoothing 5. Statistical models and algorithms 6. Hypothesis testing 	Apply, Explain	KA	30
MSF-Linear	<ol style="list-style-type: none"> 1. Vectors: definitions, vector operations, geometric interpretation, angles: Matrices: definition, matrix operations, meaning of $Ax=b$. 	Develop	CS	5
MSF-Linear	<ol style="list-style-type: none"> 2. Matrices, matrix-vector equation, geometric interpretation, geometric transformations with matrices 3. Solving equations, row-reduction 4. Linear independence, span, basis 5. Orthogonality, projection, least-squares, orthogonal bases 	Apply, Explain	KA	35

	<ul style="list-style-type: none"> 6. Linear combinations of polynomials, Bezier curves 7. Eigenvectors and eigenvalues 8. Applications to computer science: PCA, SVD, page-rank, graphics 			
MSF-Calculus	<ul style="list-style-type: none"> 1. Sequences, series, limits 2. Single-variable derivatives: definition, computation rules (chain rule etc.), derivatives of important functions, applications 3. Single-variable integration: definition, computation rules, integrals of important functions, fundamental theorem of calculus, definite vs indefinite, applications (including in probability) 4. Parametric and polar representations 5. Taylor series 6. Multivariate calculus: partial derivatives, gradient, chain-rule, vector valued functions, 7. Optimization: convexity, global vs local minima, gradient descent, constrained optimization and Lagrange multipliers. 8. ODEs: definition, Euler method, applications to simulation, Monte Carlo integration 9. CS applications: gradient descent for machine learning, forward and inverse kinematics, applications of calculus to probability 	Apply, Develop	KA	40

Curricular Packaging

A few curricular packaging options of various sizes are presented here. These can be adapted to local strengths and needs to create a customized computer science curriculum. In each case, effort should be made to include all the CS Core topics in required courses in the curriculum. The more KA Core topics covered, the greater the breadth of the curriculum. The more hours dedicated to KA Core topics, the greater the depth of the curriculum. Non-core topics add to the richness of the curriculum. In each curricular model, a capstone course is included to emphasize the importance of an integrative hands-on experience. It may also serve as the course where CS Core topics not covered elsewhere in the curriculum can be incorporated.

8 Course Model

This is a minimal course configuration that covers all the CS core topics. However, it does not leave much room for exploration:

1. CS I ([AL-2](#), [FPL-1](#), [SDF-34](#), [SEP-3](#))
2. CS II ([AL-18](#), [FPL-4](#), [MSF-4](#), [SDF-9](#), [SEC-1](#), [SEP-3](#))
3. Mathematical and Statistical Foundations ([MSF-40](#))
4. Algorithms ([AL-12](#), [MSF-11](#), [PDC-2](#), [SEP-3](#))
5. Introduction to Computing Systems ([SF-18](#), [OS-8](#), [AR-9](#), [NC-7](#), [SEP-2](#))
6. Programming Language Concepts ([FPL-16](#), [PDC-7](#), [SEP-2](#))
7. Introduction to Computing Applications ([SEC-5](#), [AI-12](#), [GIT-4](#), [DM-10](#), [SEP-3](#))
8. Capstone ([SE-6](#), [HCI-8](#), [SPD-3](#), [SEP-2](#))

10 Course Model

1. CS I ([SDF](#), [SEP](#))
2. CS II ([SDF](#), [FPL-4](#), [AL-12](#), [SEP](#))
3. Mathematical and Statistical Foundations ([MSF](#))
4. Data Structures and Algorithms ([AL-20](#), [AI](#), [MSF](#), [SEP](#))
5. Introduction to Computing Systems ([SF](#), [OS](#), [AR](#), [NC](#))
6. Programming Languages ([FPL-17](#), [AL](#), [PDC](#), [SEP](#))
7. Software Engineering ([SE](#), [HCI](#), [GIT](#), [PDC](#), [SPD](#), [DM](#), [SEP](#))
8. One Systems elective:
 - a. Operating Systems ([OS](#), [PDC](#))
 - b. Computer Architecture ([AR](#))
 - c. Parallel and Distributed Computing ([PDC](#))
 - d. Networking ([NC](#), [SEC](#), [SEP](#))
 - e. Databases ([DM](#), [SEP](#))
9. One elective from Applications:
 - a. Artificial Intelligence ([AI](#), [MSF](#), [SPD](#), [SEP](#))
 - b. Graphics ([GIT](#), [HCI](#), [MSF](#), [SEP](#))

- c. Application Security ([SEC](#), [SEP](#))
 - d. Human-Centered Design ([HCI](#), [GIT](#), [SEP](#))
10. Capstone ([SE](#), [SEP](#))

12 Course Model

1. CS I ([SDF](#), [SEP](#))
2. CS II ([SDF](#), [AL-12](#), [DM](#), [SEP](#))
3. Mathematical and Statistical Foundations ([MSF](#))
4. Algorithms ([AL-20](#), [AI](#), [MSF](#), [SEC](#), [SEP](#))
5. Introduction to Computing Systems ([SF](#), [OS](#), [AR](#), [NC](#))
6. Programming Languages ([FPL](#), [AL](#), [PDC](#), [SEP](#))
7. Software Engineering ([SE](#), [HCI](#), [GIT](#), [PDC](#), [SPD](#), [DM](#), [SEP](#))
8. Two from Systems electives:
 - a. Operating Systems ([OS](#), [PDC](#))
 - b. Computer Architecture ([AR](#))
 - c. Parallel and Distributed Computing ([PDC](#))
 - d. Networking ([NC](#), [SEC](#), [SEP](#))
 - e. Databases ([DM](#), [SEP](#))
9. Two electives from Applications:
 - a. Artificial Intelligence ([AI](#), [MSF](#), [SPD](#), [SEP](#))
 - b. Graphics ([GIT](#), [HCI](#), [MSF](#), [SEP](#))
 - c. Application Security ([SEC](#), [SEP](#))
 - d. Human-Centered Design ([HCI](#), [GIT](#), [SEP](#))
10. Capstone ([SE](#), [SEP](#))

16 Course Model

Three different models are presented here, each with its own benefits.

Model 1:

1. CS I ([SDF](#), [SEP](#))
2. CS II ([SDF](#), [AL-12](#), [DM](#), [SEP](#))
3. Mathematical and Statistical Foundations ([MSF](#))
4. Algorithms ([AL-20](#), [SEP](#))
5. Introduction to Computing Systems ([SF](#), [SEP](#))
6. Programming Languages ([FPL](#), [AL](#), [PDC](#), [SEP](#))
7. Theory of Computation ([AL-32](#), [SEP](#))
8. Software Engineering ([SE](#), [HCI](#), [GIT](#), [PDC](#), [SPD](#), [DM](#), [SEP](#))
9. Operating Systems ([OS](#), [PDC](#), [SEP](#))
10. Computer Architecture ([AR](#), [SEP](#))
11. Parallel and Distributed Computing ([PDC](#), [SEP](#))

12. Networking ([NC](#), [SEP](#))
13. Pick one of:
 - a. Introduction to Artificial Intelligence ([AI](#), [MSF](#), [SEP](#))
 - b. Machine Learning ([AI](#), [MSF](#), [SEP](#))
 - c. Robotics ([AI](#), [SPD](#), [SEP](#))
14. Pick one of:
 - a. Graphics ([GIT](#), [MSF](#), [SEP](#))
 - b. Human-Centered Design ([GIT](#), [SEP](#))
 - c. Animation ([GIT](#), [SEP](#))
 - d. Virtual Reality ([GIT](#), [SEP](#))
15. Security ([SEC](#), [SEP](#))
16. Capstone ([SE](#), [SEP](#))

Model 2:

1. CS I ([SDF](#), [SEP](#))
2. CS II ([SDF](#), [AL](#), [DM](#), [SEP](#))
3. Mathematical and Statistical Foundations ([MSF](#), [AI](#), [DM](#))
4. Algorithms ([AL](#), [MSF](#), [SEP](#))
5. Introduction to Computing Systems ([SF](#), [SEP](#))
6. Programming Languages ([FPL](#), [AL](#), [PDC](#), [SEP](#))
7. Theory of Computation ([AL](#), [SEP](#))
8. Software Engineering ([SE](#), [HCI](#), [GIT](#), [PDC](#), [SPD](#), [DM](#), [SEP](#))
9. Operating Systems ([OS](#), [PDC](#), [SEP](#))
10. Two electives from:
 - a. Computer Architecture ([AR](#), [SEP](#))
 - b. Parallel and Distributed Computing ([PDC](#), [SEP](#))
 - c. Networking ([NC](#), [SEP](#))
 - d. Network Security ([NC](#), [SEC](#), [SEP](#))
 - e. Security ([SEC](#), [SEP](#))
11. Pick three of:
 - a. Introduction to Artificial Intelligence ([AI](#), [MSF](#), [SEP](#))
 - b. Machine Learning ([AI](#), [MSF](#), [SEP](#))
 - c. Deep Learning ([AI](#), [MSF](#), [SEP](#))
 - d. Robotics ([AI](#), [SPD](#), [SEP](#))
 - e. Data Science ([AI](#), [DM](#), [GIT](#), [MSF](#))
 - f. Graphics ([GIT](#), [MSF](#), [SEP](#))
 - g. Human-Computer interaction ([HCI](#), [SEP](#))
 - h. Human-Centered Design ([GIT](#), [HCI](#), [SEP](#))
 - i. Animation ([GIT](#), [SEP](#))
 - j. Virtual Reality ([GIT](#), [SEP](#))
 - k. Physical Computing ([GIT](#), [SPD](#), [SEP](#))
12. Society, Ethics and Professionalism ([SEP](#))
13. Capstone ([SE](#), [SEP](#))

Model 3:

1. CS I ([SDF](#), [SEP](#))
2. CS II ([SDF](#), [AL](#), [DM](#), [SEP](#))
3. Mathematical and Statistical Foundations ([MSF](#))
4. Algorithms ([AL](#), [AI](#), [MSF](#), [SEC](#), [SEP](#))
5. Introduction to Computing Systems ([SF](#), [OS](#), [AR](#), [NC](#))
6. Programming Languages ([FPL](#), [AL](#), [PDC](#), [SEP](#))
7. Software Engineering ([SE](#), [HCI](#), [GIT](#), [PDC](#), [SPD](#), [DM](#), [SEP](#))
8. Two from Systems electives:
 - a. Operating Systems ([OS](#), [PDC](#))
 - b. Computer Architecture ([AR](#))
 - c. Parallel and Distributed Computing ([PDC](#))
 - d. Networking ([NC](#), [SEC](#), [SEP](#))
 - e. Databases ([DM](#), [SEP](#))
9. Two electives from Applications:
 - a. Artificial Intelligence ([AI](#), [MSF](#), [SPD](#), [SEP](#))
 - b. Graphics ([GIT](#), [HCI](#), [MSF](#), [SEP](#))
 - c. Application Security ([SEC](#), [SEP](#))
 - d. Human-Centered Design ([HCI](#), [GIT](#), [SEP](#))
10. Three open CS electives
11. Society, Ethics and Professionalism ([SEP](#)) course
12. Capstone ([SE](#), [SEP](#))