Artificial Intelligence Knowledge Area

Draft Version Alpha 2022-07-20

Artificial intelligence (AI) is the study of solutions for problems that are difficult or impractical to solve with traditional algorithmic approaches. AI solution strategies are often reminiscent of those considered to require human intelligence, and typically generalize over classes of problems. AI techniques are used pervasively in support of 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 the environment, such as self-driving vehicles and other robots.

Traditionally, AI has included a mix of symbolic and subsymbolic approaches. The solutions rely on a broad set of general and specialized knowledge representation schemes, problem solving mechanisms, and optimization techniques. They deal with perception (e.g., speech recognition, natural language understanding, computer vision), problem solving (e.g., search, planning, optimization), acting (e.g., robotics, task-automation, control), and the architectures needed to support them (e.g., single agents, multi-agent systems). The study of Artificial Intelligence prepares students to determine when an AI approach is appropriate for a given problem, identify appropriate representation 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 caused terms such as "neural networks" and "deep learning" to become part of everyday vernacular. Businesses now advertise AI-based solutions as value-additions to their services, using "artificial intelligence" as both a technical term and a marketing keyword. Other disciplines, such as biology, art, architecture, and finance, increasingly use AI techniques to solve problems within their disciplines. Additionally, AI technology is increasingly relied upon by the broader population and can have significant societal impacts, from its use in stock trading, to automated evaluation of job applicants, to detecting medical conditions, to influencing prison sentencing through recidivism prediction.

How this KA has changed since CS 2013

To reflect this recent growth and societal impact, the knowledge area has been revised from CS 2013 in the following ways:

- The name has changed from "Intelligent Systems" to "Artificial Intelligence," reflecting the most common term used for these topics within the field and the more widespread use of the term outside of the field.
- There is increased emphasis on neural networks and representation learning, reflecting the recent advances in the field. Search is still emphasized due to its key role throughout AI, but there is a slight reduction on symbolic methods in favor of understanding subsymbolic methods and learned representations. However, it is important to retain knowledge-based and symbolic approaches within the AI curriculum since these methods offer unique capabilities and are used in practice, in order to 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.).
- 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 cross-connections with the Data Management knowledge area.

- 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.

Allocation of Hours

AI. Artificial Intelligence (10 CS core hours, 10 KA core hours)

* 3 hrs overlaps with AL Algorithms (breadth and depth-first search)
$ Partial overlap with discrete mathematics / prob & stats (probability and Bayes' rule)

	CS Core	KA Core	CS2013 Core 2 (for comparison)	Electives
AI/Fundamental Issues	2		1	**Y**
AI/Basic Search Strategies	5 *	3	4	**N**
AI/Basic Symbolic Knowledge Representation and Reasoning	2 $	2	3 (includes probability)	**N**
AI/Basic Machine Learning	3	3	2	**N**
AI/Applications and Societal Impact	1	2		**Y**
AI/Advanced Search				**Y**
AI/Probabilistic Representation and Reasoning				**Y**
AI/Planning				**Y**
AI/Advanced Machine Learning				**Y**
AI/Advanced Representation and Reasoning				**Y**
AI/Agents				**Y**
AI/Natural Language Processing				**Y**

AI/Robotics				**Y**
AI/Perception and Computer Vision				**Y**

AI/Fundamental Issues

[2 CS Core hours]

Topics (CS Core):

- Overview of AI problems, Examples of successful recent AI applications
- What is intelligent behavior?
    - The Turing test
    - Multimodal input and output
    - Simulation of intelligent behavior
    - Rational versus non-rational reasoning
- Problem characteristics
    - Fully versus partially observable
    - Single versus multi-agent
    - Deterministic versus stochastic
    - Static versus dynamic
    - Discrete versus continuous
- Nature of agents
    - Autonomous, semi-autonomous, mixed-initiative autonomy
    - Reflexive, goal-based, and utility-based
    - Decision making under uncertainty and with incomplete information
    - The importance of perception and environmental interactions
    - Learning-based agents
    - Embodied agents
        - sensors, dynamics, effectors
- AI Applications, Growth, and Impact (Economic, Societal, Ethics)
- Philosophical issues. [elective]

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. Determine the characteristics of a specific problem.

AI/Basic Search Strategies

[5 CS Core hours, 3 of which overlap with AL Algorithms (Uninformed search); 3 additional KA hours]

(Cross-reference AL/Basic Analysis, AL/Algorithmic Strategies, AL/Fundamental Data Structures and Algorithms)

Topics (CS Core):

- State space representation of a problem
    - Specifying states, goals, and operators
    - Factoring states into representations
    - Problem solving by graph search
        - Dynamic construction of the graph (you're not given it upfront)
- Uninformed graph search for problem solving
    - Breadth-first search
    - Depth-first search
        - With iterative deepening
    - Uniform cost search
- Heuristic graph search for problem solving
    - Heuristic construction and admissibility
    - Hill-climbing
    - Local minima and the search landscape
        - Local vs global solutions
    - Greedy best-first search
    - A* search
- Space and time complexities of graph search algorithms

Topics (KA Core):

- *Bidirectional search*
- *Beam search*
- *Two-player adversarial games*
    - *Minimax search*
- *Implementation of A* search*

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 the necessary 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]

AI/Basic Knowledge Representation and Reasoning

[2 CS Core hours; 2 KA Core hours]

Topics (CS Core):

- Types of representations
    - Symbolic, logical
        - Creating a representation from a natural language problem statement
    - Learned subsymbolic representations
    - Graphical models
- Review of probabilistic reasoning, Bayes theorem (cross-reference with DS/Discrete Probability)
- Bayesian reasoning
    - Bayesian inference

Topics (KA Core):

- Random variables and probability distributions
    - Axioms of probability
    - Probabilistic inference
    - Bayes' Rule (derivation)
    - Bayesian inference (more complex examples)
- Independence
- Conditional Independence
- Utility and decision making

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. Make 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/Basic Machine Learning

[3 CS Core hours; 3 additional KA hours]

Topics:

- Definition and examples of a broad variety of machine learning tasks
    - Supervised learning
        - Classification
        - Regression
    - Reinforcement learning
    - Unsupervised learning
        - Clustering
- Simple statistical-based supervised learning such as Naive Bayes, Decision trees
- The overfitting problem and controlling solution complexity (regularization, pruning)
    - The bias (underfitting) - variance (overfitting) tradeoff
- Working with Data
    - Data preprocessing
        - Importance and pitfalls of
    - Handling missing values (imputing, flag-as-missing)
        - Implications of imputing vs flag-as-missing
    - Encoding categorical variables, encoding real-valued data
    - Normalization/standardization
    - Emphasis on real data, not textbook examples
- Representations
    - Simple basis feature expansion, such as squaring univariate features
    - Learned feature representations
- Machine learning evaluation
    - Measuring classifier accuracy
    - Separation of train, validation, and test sets
    - Estimation of test performance, using held-out data
        - Tuning the parameters of a machine learning model on held-out validation data

    - Importance of understanding what your model is actually doing, where its pitfalls/shortcomings are, and the implications of its decisions
- Basic neural networks
    - Fundamentals of understanding how neural networks work and their training process, without details of the calculations

Topics (KA core):

- Formulation of simple machine learning as an optimization problem, such as least squares linear regression or logistic regression
    - Objective function
    - Gradient descent
    - Regularization to avoid overfitting
- Ensembles of models
    - Simple weighted majority combination
- Deep learning
    - Deep feed-forward networks (intuition only, no math)
    - Convolutional neural networks (intuition only, no math)
    - Visualization of learned feature representations from deep nets
- Performance evaluation
    - Other Metrics (e.g., error, precision, recall)
    - Confusion matrix
    - Cross-validation
        - Parameter tuning (grid/random search, via cross-validation)
- Overview of reinforcement learning
- Two or more applications of machine learning algorithms
    - E.g., medicine and health, economics, vision, natural language, robotics, game play
- Ethics for Machine Learning
    - (Note: although this is related to Ethics in AI/Fundamental Issues, ML ethics differs substantially from general AI ethics discussions (e.g., what's in Russell and Norvig).

Learning Outcomes:

1. Describe the differences among the three main styles of learning: supervised, reinforcement, and unsupervised.
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 (such as k-nearest-neighbors and logistic regression) to a classification task and measure the classifiers' accuracy.
5. Identify over-fitting 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. Describe the neural network training process and resulting learned representations
8. 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.
9. Implement and compare two machine learning algorithms on a dataset, preprocessing it from scratch.

AI/Applications and Societal Impact

[1 CS Core hour; 2 additional KA hours]

(Crosslist with SP)

Note: There is substantial benefit to studying applications and ethics/fairness/trust/explainability in a curriculum alongside the methods and theory that it applies 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.

For the CS core, cover at least one application and an overview of the societal issues of AI/ML. The KA core should go more in-depth with one or more additional applications, and an analysis and discussion of the social issues.

Topics:

- Applications of AI to a broad set of problems and diverse fields, such as medicine, health, sustainability, social media, economics, robotics, etc. (choose one for CS Core, at least one additional for KA core)
    - Formulating and evaluating a specific application as an AI problem
    - Data availability and cleanliness
        - Basic data cleaning and preprocessing
        - Data set bias
    - Algorithmic bias
    - Evaluation bias
- Societal impact of AI
    - Ethics
    - Fairness
    - Trust / explainability

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.

AI/Advanced Search *[Elective]*

Note that the general topics of Branch-and-bound and Dynamic Programming are listed in (AL/Algorithmic Strategies).

Topics:

- Understanding the search space
    - Constructing search trees
    - Dynamic search spaces
    - Combinatorial explosion of search space
    - Search space topology (ridges, saddle points, local minima, etc.)
- Local search and constraint satisfaction
- Tabu search
- Variations on A* (IDA*, SMA*, RBFS)

- Two-player adversarial games
    - Alpha-beta pruning
        - Ply cutoff
    - The horizon effect
    - Opening playbooks / endgame solutions
- Implementation of minimax search, beam search
- Expectimax search (MDP-solving) and chance nodes
- Stochastic search
    - Simulated annealing
    - Genetic algorithms
    - Monte-Carlo tree search

Learning Outcomes:

1. Design and implement a genetic algorithm solution to a problem.
2. Design and implement a simulated annealing schedule to avoid local minima in a problem.
3. 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.
4. Apply minimax search with alpha-beta pruning to prune search space in a two-player adversarial game (e.g., connect four).
5. 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).
6. Compare and contrast various heuristic searches vis-a-vis applicability to a given problem.

AI/Advanced Representation and Reasoning *[Elective]*

Topics:

- Review of propositional and predicate logic (cross-reference DS/Basic Logic)
- Resolution and theorem proving (propositional logic only)
    - Forward chaining, backward chaining
- Knowledge representation issues
    - Description logics
    - Ontology engineering
- Semantic web
- Non-monotonic reasoning (e.g., non-classical logics, default reasoning)
- Argumentation
- Reasoning about action and change (e.g., situation and event calculus)
- Temporal and spatial reasoning
- Logic programming
    - Prolog, Answer Set Programming
- Rule-based Expert Systems
- Semantic networks
- Model-based and Case-based reasoning
- Planning:
    - Partial and totally ordered planning
    - Plan graphs
    - Hierarchical planning
    - Planning and execution including conditional planning and continuous planning
    - Mobile agent/Multi-agent planning

Learning Outcomes:

1. Translate a natural language (e.g., English) sentence into predicate logic statement.
2. Convert a logic statement into clause 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/Probabilistic Representation and Reasoning *[Elective]*

Topics:

- Conditional Independence review
- Knowledge representations
    - Bayesian Networks
        - Exact inference and its complexity
        - Markov blankets and d-separation
        - Randomized sampling (Monte Carlo) methods (e.g. Gibbs sampling)
    - Markov Networks
    - Relational probability models
    - Hidden Markov Models
- Decision Theory
    - Preferences and utility functions
    - Maximizing expected utility

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 *[Elective]*

Topics:

- Review of propositional and first-order logic
- Planning operators and state representations
- Partial-order planning
- GraphPlan
- Planning languages and representations
    - PDDL
- Multi-agent planning

Learning Outcomes:

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

AI/Agents *[Elective]*

(Cross-reference HCI/Collaboration and Communication)

Topics:

- Definitions of agents
- Agent architectures (e.g., reactive, layered, cognitive)
- Agent theory
- Rationality, Game Theory
    - Decision-theoretic agents
    - Markov decision processes (MDP)
- Software agents, personal assistants, and information access
    - Collaborative agents
    - Information-gathering agents
    - Believable agents (synthetic characters, modeling emotions in agents)
- Learning agents
- Multi-agent systems
    - Collaborating agents
    - Agent teams
    - Competitive agents (e.g., auctions, voting)
    - Swarm systems and biologically inspired models
    - Multi-agent learning
- Human-agent interaction
    - Communication methodologies
    - Practical issues
    - Applications
        - Trading agents, supply chain management

Learning Outcomes:

4. List the defining characteristics of an intelligent agent.
5. Characterize and contrast the standard agent architectures.
6. Describe the applications of agent theory to domains such as software agents, personal assistants, and believable agents.

7. Describe the primary paradigms used by learning agents.
8. Demonstrate using appropriate examples how multi-agent systems support agent interaction.
9. Construct an intelligent agent using a well-established cognitive architecture (ACT-R, Soar) for solving a specific problem.

AI/Natural Language Processing *[Elective]*

(Cross-reference HCI/New Interactive Technologies)

Topics:

- Deterministic and stochastic grammars
- Parsing algorithms
    - CFGs and chart parsers (e.g. CYK)
    - Probabilistic CFGs and weighted CYK
- Representing meaning / Semantics
    - Logic-based knowledge representations
    - Semantic roles
    - Temporal representations
    - Beliefs, desires, and intentions
- Corpus-based methods
- N-grams and HMMs
- Smoothing and backoff
- Examples of use: POS tagging and morphology
- Information retrieval (Cross-reference IM/Information Storage and Retrieval)
    - Vector space model
        - TF & IDF
    - Precision and recall
- Information extraction
- Language translation
- Text classification, categorization
    - Bag of words model
- Deep learning for NLP (Cross-reference AI/Advanced Machine Learning)
    - RNNs
    - Transformers
    - Multi-modal embeddings (e.g., images + text)

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/Advanced Machine Learning *[Elective]*

(See also AI/Basic Machine Learning)
Topics:

- Definition and examples of broad variety of machine learning tasks
- General statistical-based learning, parameter estimation (maximum likelihood)
- Inductive logic programming (ILP)
- Supervised learning
    - Decision trees
    - Learning simple neural networks / multi-layer perceptrons
    - Linear regression
    - Logistic regression
    - Support vector machines (SVMs) and kernels
    - Gaussian Processes
- Overfitting
    - The curse of dimensionality
    - Regularization (math computations, L_2 and L_1 regularization)
- Ensembles
    - Weighted majority
    - Boosting/Bagging
    - Random forest
    - Gated ensemble
- Nearest-neighbor classification and regression
- Bayesian learning (Cross-Reference AI/Reasoning Under Uncertainty)
    - Naive Bayes and its relationship to linear models
    - Bayesian networks
    - Prior/posterior
    - Generative models
- Deep learning
    - Deep feed-forward networks
    - Neural tangent kernel and understanding neural network training
    - Convolutional neural networks
    - Auto-encoders
    - Recurrent networks
    - Representation transfer
    - Adversarial training and generative adversarial networks
- Representations
    - Manually crafted representations
    - Basis expansion
    - Learned representations (e.g., deep neural networks)
- Unsupervised learning and clustering
    - K-means
    - Gaussian mixture models
    - Expectation maximization (EM)
    - Self-organizing maps
- Semi-supervised learning
- Graphical models (Cross-reference AI/Probabilistic Representation and Reasoning)
- Performance evaluation
    - Metrics
        - error, precision, recall, confusion matrix
    - Cross-validation
    - Receiver operating characteristic (ROC) curve and area under ROC curve
    - Parameter tuning (grid/random search, via cross-validation)
- Learning theory

    - General overview of learning theory / why learning works
    - VC dimension
    - Generalization bounds
- Reinforcement learning
    - Exploration vs. exploitation trade-off
    - Markov decision processes
    - Value and policy iteration
    - Policy gradient methods
    - Deep reinforcement learning
- Application of machine learning algorithms to:
    - Medicine and health
    - Economics
    - Vision
    - Natural language
    - Robotics
    - Game play
    - Data Mining (Cross-reference IM/Data Mining)
- Ethics for Machine Learning
    - (Note: although this is related to Ethics in AI/Fundamental Issues, ML ethics differs substantially from general AI ethics discussions (e.g., what's in Russell and Norvig).

Learning Outcomes:

1. Explain the differences among the three main styles of learning: supervised, reinforcement, and unsupervised.
2. Implement simple algorithms for supervised learning, reinforcement learning, and unsupervised learning.
3. Determine which of the three learning styles is appropriate to a particular problem domain.
4. Compare and contrast each of the following techniques, providing examples of when each strategy is superior: decision trees, logistic regression, naive Bayes, neural networks, and belief networks.
5. Evaluate the performance of a simple learning system on a real-world dataset.
6. Characterize the state of the art in learning theory, including its achievements and its shortcomings.
7. Explain the problem of overfitting, along with techniques for detecting and managing the problem.

AI/Robotics *[Elective]*

Topics:

- Overview: problems and progress
    - State-of-the-art robot systems, including their sensors and an overview of their sensor processing
    - Robot control architectures, e.g., deliberative vs. reactive control and Braitenberg vehicles
    - World modeling and world models
    - Inherent uncertainty in sensing and in control
- Configuration space and environmental maps
- Interpreting uncertain sensor data
- Localization and mapping
- Navigation and control
- Motion planning and trajectory optimization
- Multiple-robot coordination
- Human-robot collaboration

Learning Outcomes:

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.

AI/Perception and Computer Vision *[Elective]*

Topics:

- Computer vision
  - Image acquisition, representation, processing and properties
  - Shape representation, object recognition, and segmentation
  - Motion analysis
  - Generative models
- Audio and speech recognition
- Modularity in recognition
- Approaches to pattern recognition. [cross-reference AI/Advanced Machine Learning]
  - Classification algorithms and measures of classification quality
  - Statistical techniques
  - Deep learning techniques

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.

Desirable Professional Dispositions

The most desirable professional dispositions for this knowledge area are:

- **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 efforts. These call for persistence on the part of the student.
- **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 choice for when and how to apply AI techniques.

Necessary and Desirable Math

The Math necessary for this knowledge area includes:

- Discrete Math:
  - sets, relations, functions
- Linear Algebra:
  - Matrix operations, matrix algebra
- Probability and 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

The Math desirable for this knowledge area includes:

- Calculus-based probability and statistics
- Other topics in probability and statistics
  - Hypothesis testing, data resampling, experimental design techniques
- Optimization

Shared and Crosscutting Concepts

This knowledge area shares the following concepts with other knowledge areas:

- Search algorithms (AI/Basic Search) with Algorithms and Complexity (AL/Search)
- Data management (for data science) with TBD
- Ethics with TBD
- Robotics with SPD/Industrial Platforms

Crosscutting concepts that apply to this knowledge area include:

- Efficiency
- Ethics
- Modeling
- Programming

Subcommittee

Chair: Eric Eaton, University of Pennsylvania
Subcommittee members:

- Zachary Dodds, Harvey Mudd College
- Susan Epstein, Hunter College, CUNY
- Laura Hiatt, US Naval Research Lab
- Amruth N. Kumar, Ramapo College of New Jersey
- Peter Norvig, Google
- Meinolf Sellmann, GE Research
- Reid Simmons, Carnegie Mellon University

Other contributors: *(meant for one-time or occasional contributors)*

- Claudia Schulz, Thomson Reuters
-