Intelligent Systems (IS)

The field of artificial intelligence (AI) is concerned with the design and analysis of autonomous agents. These are software systems and/or physical machines, with sensors and actuators, embodied for example within a robot or an autonomous spacecraft. An intelligent system has to perceive its environment, to act rationally towards its assigned tasks, to interact with other agents and with human beings.

These capabilities are covered by topics such as computer vision, planning and acting, robotics, multiagents systems, speech recognition, and natural language understanding. They rely on a broad set of general and specialized knowledge representations and reasoning mechanisms, on problem solving and search algorithms, and on machine learning techniques.

Furthermore, artificial intelligence provides a set of tools for solving problems that are difficult or impractical to solve with other methods. These include heuristic search and planning algorithms, formalisms for knowledge representation and reasoning, machine learning techniques, and methods applicable to sensing and action problems such as speech and language understanding, computer vision, and robotics, among others. The student needs to be able to determine when an AI approach is appropriate for a given problem, and to be able to select and implement a suitable AI method.

IS. Intelligent Systems (10 core hours)

IS/FundamntalIssues [core]
IS/BasicSearchStrategies [core]
IS/KnowledgeBasedReasoning [core]
IS/AdvancedSearch [elective]
IS/AdvancedReasoning [elective]
IS/Agents [elective]
IS/NaturaLanguageProcessing [elective]
IS/MachineLearning [elective]
IS/PlanningSystems [elective]
IS/Robotics [elective]
IS/Perception [elective]

IS/FundamentalIssues [core]

Minimum core *coverage time: 1 hour*

Topics:

- History of artificial intelligence
- Philosophical questions
- The Turing test
- Searle's "Chinese Room" thought experiment
- Ethical issues in AI
- Fundamental definitions
- Optimal vs. human-like reasoning
- Optimal vs. human-like behavior
- Philosophical questions
- Modeling the world
- The role of heuristics

Learning Objectives:

1. Describe the Turing test and the "Chinese Room" thought experiment.
2. Differentiate the concepts of optimal reasoning and human-like reasoning.
3. Differentiate the concepts of optimal behavior and human-like behavior.
4. List examples of intelligent systems that depend on models of the world.
5. Describe the role of heuristics and the need for tradeoffs between optimality and efficiency.

IS/BasicSearchStrategies [core]

Minimum core *coverage time: 5 hours*

Topics:

- Problem spaces; problem solving by search
- Brute-force search (breadth-first, depth-first, depth-first with iterative deepening)
- Best-first search (generic best-first, Dijkstra's algorithm, A*, admissibility of A*)
- Two-player games (minimax search, alpha-beta pruning
- Constraint satisfaction (backtracking and local search methods)

Learning Objectives:

1. Formulate an efficient problem space for a problem expressed in English by expressing that problem space in terms of states, operators, an initial state, and a description of a goal state.
2. Describe the problem of combinatorial explosion and its consequences.
3. Select an appropriate brute-force search algorithm for a problem, implement it, and characterize its time and space complexities.
4. Select an appropriate heuristic search algorithm for a problem and implement it by designing the necessary heuristic evaluation function.
5. Describe under what conditions heuristic algorithms guarantee optimal solution.
6. Implement minimax search with alpha-beta pruning for some two-player game.
7. Formulate a problem specified in English as a constraint-satisfaction problem and implement it using a chronological backtracking algorithm.

IS/KnowledgeBasedReasoning [core]

Minimum core *coverage time: 4 hours*

Topics:

- Review of propositional and predicate logic
- Resolution and theorem proving
- Nonmonotonic inference; unification and lifting, forward chaining, backward chaining, resolution
- Probabilistic reasoning
- Bayes theorem

Learning Objectives:

1. Explain the operation of the resolution technique for theorem proving.
2. Explain the distinction between monotonic and non-monotonic inference.
3. Discuss the advantages and shortcomings of probabilistic reasoning.
4. Apply Bayes theorem to determine conditional probabilities.

IS/AdvancedSearch [elective]

Topics:

- Heuristics
- Local search and optimization
- Hill climbing
- Genetic algorithms
- Simulated annealing
- Local beam search
- Adversarial search for games

Learning Objectives:

1. Explain what genetic algorithms are and constrastcontrast their effectiveness with the classic problem-solving and search techniques.
2. Explain how simulated annealing can be used to reduce search complexity and contrast its operation with classic search techniques.
3. Apply local search techniques to a classic domain.

IS/AdvancedReasoning [elective]

Topics:

- Structured representation
  - Frames and objects
  - Description logics
  - Inheritance systems
- Non-monotonic reasoning
  - Nonclassical logics
  - Default reasoning
  - Belief revision
  - Preference logics
  - Integration of knowledge sources
  - Aggregation of conflicting belief
- Reasoning on action and change
  - Situation calculus
  - Event calculus
  - Ramification problems
- Temporal and spatial reasoning
- Uncertainty
  - Probabilistic reasoning
  - Bayesian nets
  - Decision theory
- Knowledge representation for diagnosis, qualitative representation
- Ontology engineering
- Semantic networks

Learning Objectives:

1. Compare and contrast the most common models used for structured knowledge representation, highlighting their strengths and weaknesses.
2. Characterize the components of nonmonotonic reasoning and its usefulness as a representational mechanisms for belief systems.
3. Apply situation and event calculus to problems of action and change.
4. Articulate the distinction between temporal and spatial reasoning, explaining how they interrelate.
5. Describe and contrast the basic techniques for representing uncertainty.
6. Describe and contrast the basic techniques for diagnosis and qualitative representation.

IS/Agents [elective]

Topics:

- Definition of agents
- Successful applications and state-of-the-art agent-based systems
- Agent architectures
  - Simple reactive agents
  - Reactive planners
  - Layered architectures
  - Example architectures and applications
- Agent theory
  - Commitments
  - Intentions
  - 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
  - Economically inspired multi-agent systems
  - Collaborating agents

    - Agent teams
    - Agent modeling
    - Multi-agent learning
- Introduction to robotic agents
- Mobile agents

Learning Objectives:
1. Explain how an agent differs from other categories of intelligent systems.
2. Characterize and contrast the standard agent architectures.
3. Describe the applications of agent theory, to domains such as software agents, personal assistants, and believable agents.
4. Describe the distinction between agents that learn and those that don't.
5. Demonstrate using appropriate examples how multi-agent systems support agent interaction.
6. Describe and contrast robotic and mobile agents.

IS/NaturalLanguageProcessing [elective]

Topics:
- Deterministic and stochastic grammar
- Parsing algorithms
- Corpus-based methods
- Information retrieval and information extraction
- Language translation
- Speech recognition

Learning Objectives:
1. Define and contrast deterministic and stochastic grammars, providing examples to show the adequacy of each.
2. Identify the classic parsing algorithms for parsing natural language.
3. Defend the need for an established corpus.
4. Give examples of catalog and look up procedures in a corpus-based approach.
5. Articulate the distinction between techniques for information retrieval, language translation, and speech recognition.

IS/MachineLearning [elective]

Topics:
- Definition and examples of machine learning
- Inductive learning, statistical based learning, reinforcement learning
- Supervised learning
- Learning decision trees
- Learning neural networks
- Learning belief networks
- The nearest neighbor algorithm
- Learning theory
- The problem of overfitting
- Unsupervised learning
- Reinforcement learning

Learning Objectives:
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, neural networks, and belief networks..
5. Implement a simple learning system using decision trees, neural networks and/or belief networks, as appropriate.
6. Characterize the state of the art in learning theory, including its achievements and its shortcomings.

7. Explain the nearest neighbor algorithm and its place within learning theory.
8. Explain the problem of overfitting, along with techniques for detecting and managing the problem.

IS/PlanningSystems [elective]

Topics:
- Definition and examples of planning systems<
- Planning as search
- Operator-based planning
- Planning graphs
- Propositional planning
- Extending planning systems (case-based, learning, and probabilistic systems)
- Static world planning systems
- Planning and execution including conditional planning and continuous planning
- Mobile agent planning
- Planning and robotics

Learning Objectives:
1. Define the concept of a planning system.
2. Explain how planning systems differ from classical search techniques.
3. Articulate the differences between planning as search, operator-based planning, and propositional planning, providing examples of domains where each is most applicable.
4. Define and provide examples for each of the following techniques: case-based, learning, and probabilistic planning.
5. Compare and contrast static world planning systems with those need dynamic execution.
6. Explain the impact of dynamic planning on robotics.

IS/Robotics [elective]

Topics:
- Overview
- State-of-the-art robot systems
- Planning vs. reactive control
- Uncertainty in control
- Sensing
- World models
- Configuration space<
- Planning
- Sensing
- Robot programming
- Navigation and control
- Robotic software and its architecture

Learning Objectives:
1. Outline the potential and limitations of today's state-of-the-art robot systems.
2. Implement configuration space algorithms for a 2D robot and complex polygons.
3. Implement simple motion planning algorithms.
4. Explain the uncertainties associated with sensors and how to deal with those uncertainties.
5. Design a simple control architecture.
6. Describe various strategies for navigation in unknown environments, including the strengths and shortcomings of each.
7. Describe various strategies for navigation with the aid of landmarks, including the strengths and shortcomings of each.

IS/Perception [elective]

Topics:

- Perception: role and applications
- Image formation: light, colour, shades
- Image and object detection: feature recognition, object recognition
- Technologies
- Software characteristics

Learning Objectives:

1. Describe the importance of image and object recognition in Ai and indicate significant applications of this technology.
2. Outline the main approaches to object recognition
3. Describe the distinguishing characteristics of the technologies used for perception.