

15-381 Artificial Intelligence

Admin.

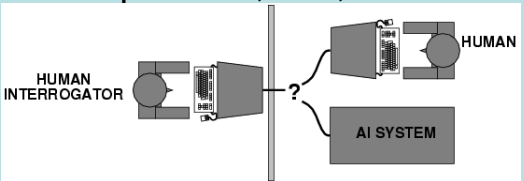
- **Instructor:**
 - Martial Hebert, NSH 4101, x8-2585
- **Textbook:**
 - Recommended (optional) textbook: [Russell and Norvig's "Artificial Intelligence: A Modern Approach"](#) (2nd edition)
 - Recommended (optional) second textbook: [Pattern Classification \(2nd Edition\)](#), Duda, Hart and Stork
- **Other resources:**
 - <http://aima.cs.berkeley.edu/>
 - <http://www.autonlab.org/tutorials/>
- **TAs:**
 - [Sonia Chernova](#), WeH 1302, x8-2601
 - [Sajid Siddiqi](#), NSH 3112, x8-6014
 - [Vaibhav Mehta](#), WeH 8303, x8-2993
 - [Rong Yan](#), NSH 4533, x8-9515
- **Grading:**
 - Midterm, Final, (probably) 6 homeworks

What is AI?

- Many different definitions and approaches

Think Like Humans	Think Rationally
Act Like Humans	Act Rationally

<h3>Think Like Humans</h3> <p>“The exciting new effort to make computers think ... machines with minds, in the full and literal sense” (Haugeland, 1985)</p> <p>“The automation of activities that we associate with human thinking, activities such as decision-making, problem solving, learning ...” (Bellman, 1978)</p>	<h3>Think Rationally</h3> <p>“The study of mental faculties through the use of computational models” (Charniak and McDermott, 1985)</p> <p>“The study of the computations that make it possible to perceive, reason, and act” (Winston, 1992)</p>
<h3>Act Like Humans</h3> <p>“The art of creating machines that perform functions that require intelligence when performed by people” (Kurzweil, 1990)</p> <p>“The study of how to make computers do things at which, at the moment, people are better” (Rich and Knight, 1991)</p>	<h3>Act Rationally</h3> <p>“A field of study that seeks to explain and emulate intelligent behavior in terms of computational processes” (Schalkoff, 1990)</p> <p>“The branch of computer science that is concerned with the automation of intelligent behavior” (Luger and Stubblefield, 1993)</p>

<p>Think Like</p> <p>“The exciting new effort to make computers think ... minds, in the full and literal sense” (Haugeland, 1985)</p> <p>“The automation of activities that we associate with human thinking, activities such as decision-making, problem solving, learning ...” (Bellman, 1978)</p>	<ul style="list-style-type: none"> • Turing test (1950) • Insight into major components: knowledge, reasoning, perception, ... <p>But difficult to reproduce, test, and analyze</p> 
<p>Act Like Humans</p> <p>“The art of creating machines that perform functions that require intelligence when performed by people” (Kurzweil, 1990)</p> <p>“The study of how to make computers do things at which, at the moment, people are better” (Rich and Knight, 1991)</p>	<p>Act Rationally</p> <p>“A field of study that seeks to explain and emulate intelligent behavior in terms of computational processes” (Schalkoff, 1990)</p> <p>“The branch of computer science that is concerned with the automation of intelligent behavior” (Luger and Stubblefield, 1993)</p>

<p>Think Like Humans</p> <p>“The exciting new effort to make computers think ... machines with minds, in the full and literal sense” (Haugeland, 1985)</p> <p>“The automation of activities that we associate with human thinking, activities such as decision-making, problem solving, learning ...” (Bellman, 1978)</p>	<p>Think Rationally</p> <p>“The study of mental faculties through the use of computational models” (Charniak and McDermott, 1985)</p> <p>“The study of the computations that make it possible to perceive, reason, and act” (Winston, 1992)</p>
<p>Act Like Humans</p> <p>“The art of creating machines that perform functions that require intelligence when performed by people” (Kurzweil, 1990)</p> <p>“The study of how to make computers do things at which, at the moment, people are better” (Rich and Knight, 1991)</p>	<ul style="list-style-type: none"> • Requires scientific theories of internal activities of the brain • Addressed in cognitive science <p>is concerned with the automation of intelligent behavior” (Luger and Stubblefield, 1993)</p>

<h3>Think Like Humans</h3> <p>“The exciting new effort to make computers think ... machines with minds, in the full and literal sense” (Haugeland, 1985)</p> <p>“The automation of activities that you associate with human thinking.”</p>	<h3>Think Rationally</h3> <p>“The study of mental faculties through the use of computational models” (Harniak and McDermott, 1985)</p> <p>“The study of the computations that make it possible to perceive, reason, and act” (Winston, 1992)</p>
<ul style="list-style-type: none"> • Aristotle: what are correct arguments/thought processes? • Development of several forms of logic: notation and rules of derivation for thoughts • Not all intelligent behavior is mediated by logical deliberation • Very difficult to represent uncertain knowledge • Rules formalize thought process but are not necessarily the best ones for goal achievement 	<h3>Rationally</h3> <p>“A field of study that seeks to explain intelligent behavior in terms of computational processes” (Schalkoff, 1990)</p> <p>“The branch of computer science that is concerned with the automation of intelligent behavior” (Luger and Stubblefield, 1993)</p>

<h3>Think Like Humans</h3> <p>“The exciting new effort to make computers think ... machines with minds, in the full and literal sense” (Haugeland, 1985)</p> <p>“The automation of activities that you associate with human thinking.”</p>	<h3>Think Rationally</h3> <p>“The study of mental faculties through the use of computational models” (Harniak and McDermott, 1985)</p> <p>“The study of the computations that make it possible to perceive, reason, and act” (Winston, 1992)</p>
<h3>Act Like Humans</h3> <p>“The art of creating machines that perform functions that require intelligence when performed by people” (Kurzweil, 1990)</p> <p>“The study of how to make computers do things at which, at the moment, people are better” (Rich and Knight, 1991)</p>	<h3>Act Rationally</h3> <p>“A field of study that seeks to explain and emulate intelligent behavior in terms of computational processes” (Schalkoff, 1990)</p> <p>“The branch of computer science that is concerned with the automation of intelligent behavior” (Luger and Stubblefield, 1993)</p>

• Aristotle: “Every art and every inquiry, and similarly every action and pursuit, is thought to aim at some good”

• Rational behavior = choose actions/decisions that is expected to maximize goal achievement, given the available information

<p>“The art of creating machines that perform functions that require intelligence when performed by people” (Kurzweil, 1990)</p> <p>“The study of how to make computers do things at which, at the moment, people are better” (Rich and Knight, 1991)</p>	<p>“A field of study that seeks to explain and emulate intelligent behavior in terms of computational processes” (Schalkoff, 1990)</p> <p>“The branch of computer science that is concerned with the automation of intelligent behavior” (Luger and Stubblefield, 1993)</p>
<p>Act Like Humans</p>	<p>Act Rationally</p>

• Aristotle: “Every art and every inquiry, and similarly every action and pursuit, is thought to aim at some good”

• Rational behavior = choose actions/decisions that is expected to maximize goal achievement, given the available information

Rational Agents

- Agent perceives and acts
- Maps percepts to actions/decisions
- Designed such that, for any environment and task, the agent achieves the best possible performance
- Computation limitation → Often impossible to find “optimal” agent → Find the best approximation given computational resources → “perfect” rationality in not possible

Search

- For a single agent,
- Find an “optimal” sequence of states between current state and goal state

“Games”

- Multiple agents maybe competing or cooperating to achieve a task
- Capabilities for finding strategies, equilibrium between agents, auctioning, bargaining, negotiating...



Planning and Reasoning

- Infer statements from a knowledge base
- Assess consistency of a knowledge base

KB =

Person \Rightarrow Mortal	If it's a person, it's a mortal
Socrates \Rightarrow Person	If it's Socrates, it's a person
Socrates \wedge Mortal \Rightarrow False	It can't be both Socrates and Mortal
True \Rightarrow Socrates	It is Socrates

Learning

- Automatically generate strategies to classify or predict from training examples

mpg	cylinders	displacement	horsepower	weight	acceleration	modelyear	maker
good	4	low	low	low	high	75to78	asia
bad	6	medium	medium	medium	medium	70to74	america
bad	4	medium	medium	medium	low	75to78	europa
bad	8	high	high	high	low	70to74	america
bad	6	medium	medium	medium	medium	70to74	america
bad	4	low	medium	low	medium	70to74	asia
bad	4	low	medium	low	low	70to74	asia
bad	8	high	high	high	low	75to78	america
:	:	:	:	:	:	:	:
:	:	:	:	:	:	:	:
:	:	:	:	:	:	:	:
bad	8	high	high	high	low	70to74	america
good	8	high	medium	high	high	79to83	america
bad	8	high	high	high	low	75to78	america
good	4	low	low	low	low	79to83	america
bad	6	medium	medium	medium	high	75to78	america
good	4	medium	low	low	low	79to83	america
good	4	low	low	medium	high	79to83	america
bad	8	high	high	high	low	70to74	america
good	4	low	medium	low	medium	75to78	europa

europa
75to78
medium
medium
medium
5 medium


➔ Mpg good/bad

Predict mpg on new data


Training data: good/bad
mpg for example cars

Learning

- Automatically generate strategies to classify or predict from training examples



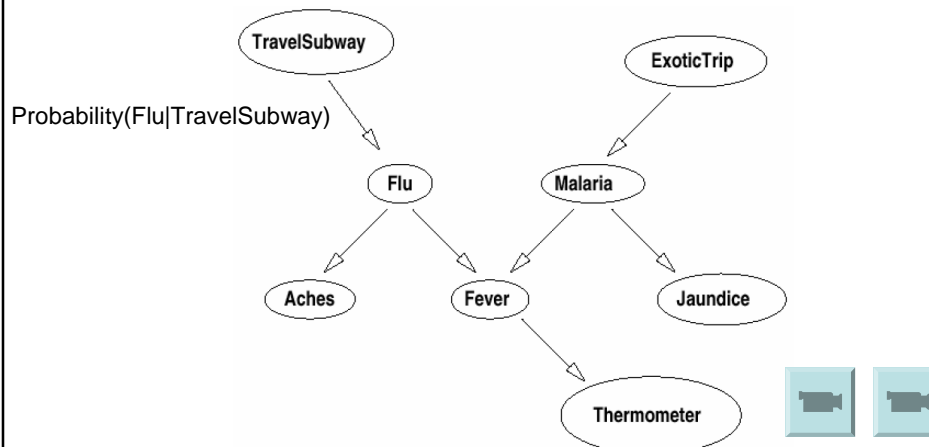
Training data: Example images of object



Classification: Is the object present in the input image, yes/no?

Reasoning with Uncertainty

- Reason (infer, make decisions, etc.) based on uncertain models, observations, knowledge



Applications

- Don't be fooled by the (sometimes) toyish examples used in the class. The AI techniques are used in a huge array of applications
 - Robotics
 - Scheduling
 - Diagnosis
 - HCI
 - Games
 - Data mining
 - Logistics
 -

Tentative schedule;
subject to frequent
changes

Jan. 17	Intro		
	SEARCH		
Jan. 19	Search	3	
Jan. 24	Search	3,4	HW1 out
Jan. 26	Search: Hill Climbing, Stochastic Search, Simulated Annealing	3,4	
Jan. 31	Search: Hill Climbing, Stochastic Search, Simulated Annealing	3,4	HW1 Review
Feb. 2	Constraint Satisfaction Problems	5	
Feb. 7	Constraint Satisfaction Problems	5	HW1 due, HW2 out
Feb. 9	Robot Motion Planning	25	
	GAMES		
Feb. 14	Algorithms for Playing and Solving Games	6	HW2 Review
Feb. 16	Games with Hidden Information	6	
Feb. 21	Non-Zero-Sum Games	6	HW2 due, HW3 out
Feb. 23	Game Theory, continued	6	
Feb. 28	Auctions and Negotiations	6	HW3 Review
	REASONING		
Mar. 2	Automated Theorem Proving with Propositional Logic	8,9	
Mar. 7	Reasoning, Cont.	8,9	HW3 due, HW4 out
Mar. 9	Midterm Exam		Midterm
Mar. 14/16	Midterm Break		
	LEARNING		
Mar. 21	Intro + Decision Trees	18	HW4 Review
Mar. 23	Decision Trees, Cont.	18	
Mar. 28	Probabilistic Learning and Naive Bayes	20	HW4 due, HW5 out
Mar. 30	Neural Networks	20	
Apr. 4	Cross-Validation	20	HW5 Review
Apr. 6	K-Means	20	
Apr. 11	Instance-Based Learning	20	HW5 due, HW6 out
	REASONING WITH UNCERTAINTY		
Apr. 13	Bayes Nets	13	
Apr. 18	Bayes Nets	13	HW6 Review
Apr. 20			
Apr. 25	Markov Decision Processes	16,17	
Apr. 27	Markov Decision Processes	16,17	HW6 due
May. 2	Reinforcement Learning	21	
May. 4	Reinforcement Learning	21	