15-381: Artificial Intelligence

Introduction and Overview

Course data

- All up-to-date info is on the course web page:
  - http://www.cs.cmu.edu/afs/cs.cmu.edu/academic/class/15381-s07/www/
- Instructors:
  - Martial Hebert
  - Mike Lewicki
- TAs:
  - Rebecca Hutchinson
  - Gil Jones
  - Ellie Lin
  - Einat Minkov
  - Arthur Tu
- See web page for contact info, office hours, etc.

Intelligence

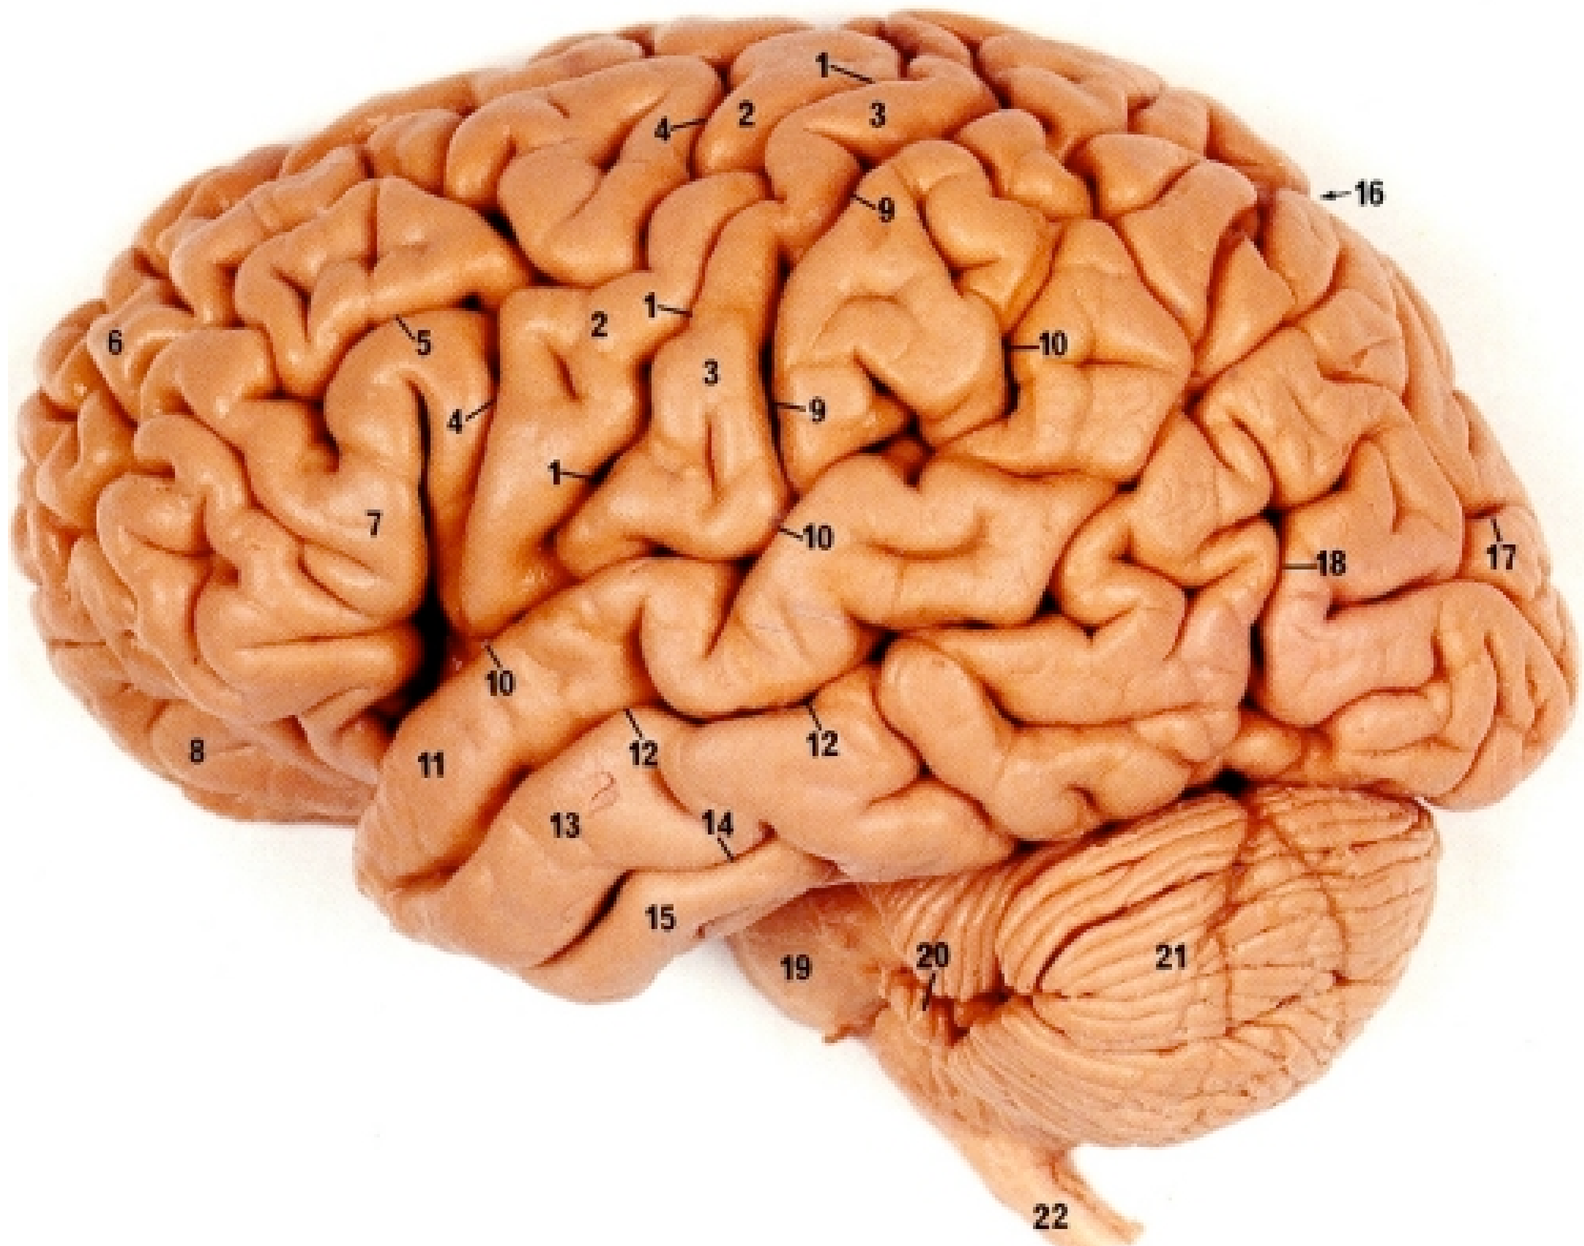

What is "intelligence" ?

Can we emulate intelligent behavior in machines ?

How far can we take it ?

Brains vs computers

Brains (adult cortex)

- surface area: 2500 cm^2
- squishy

- neurons: 20 billion
- synapses: 240 trillion

- neuron size: 15 um
- synapse size: 1 um

- synaptic OPS: 30 trillion

Computers (Intel Core 2)

- surface area: 90 mm^2
- crystalline

- transistors: 291 million

- transistor size: 65 nm

- FLOPS: 25 billion

Deep Blue: 512 processors, 1 TFLOP

Intelligent systems

Three key steps of a knowledge-based agent (Craik, 1943):

1. the stimulus must be translated into an internal representation
2. the representation is manipulated by cognitive processes to derive new internal representations
3. these in turn are translated into action

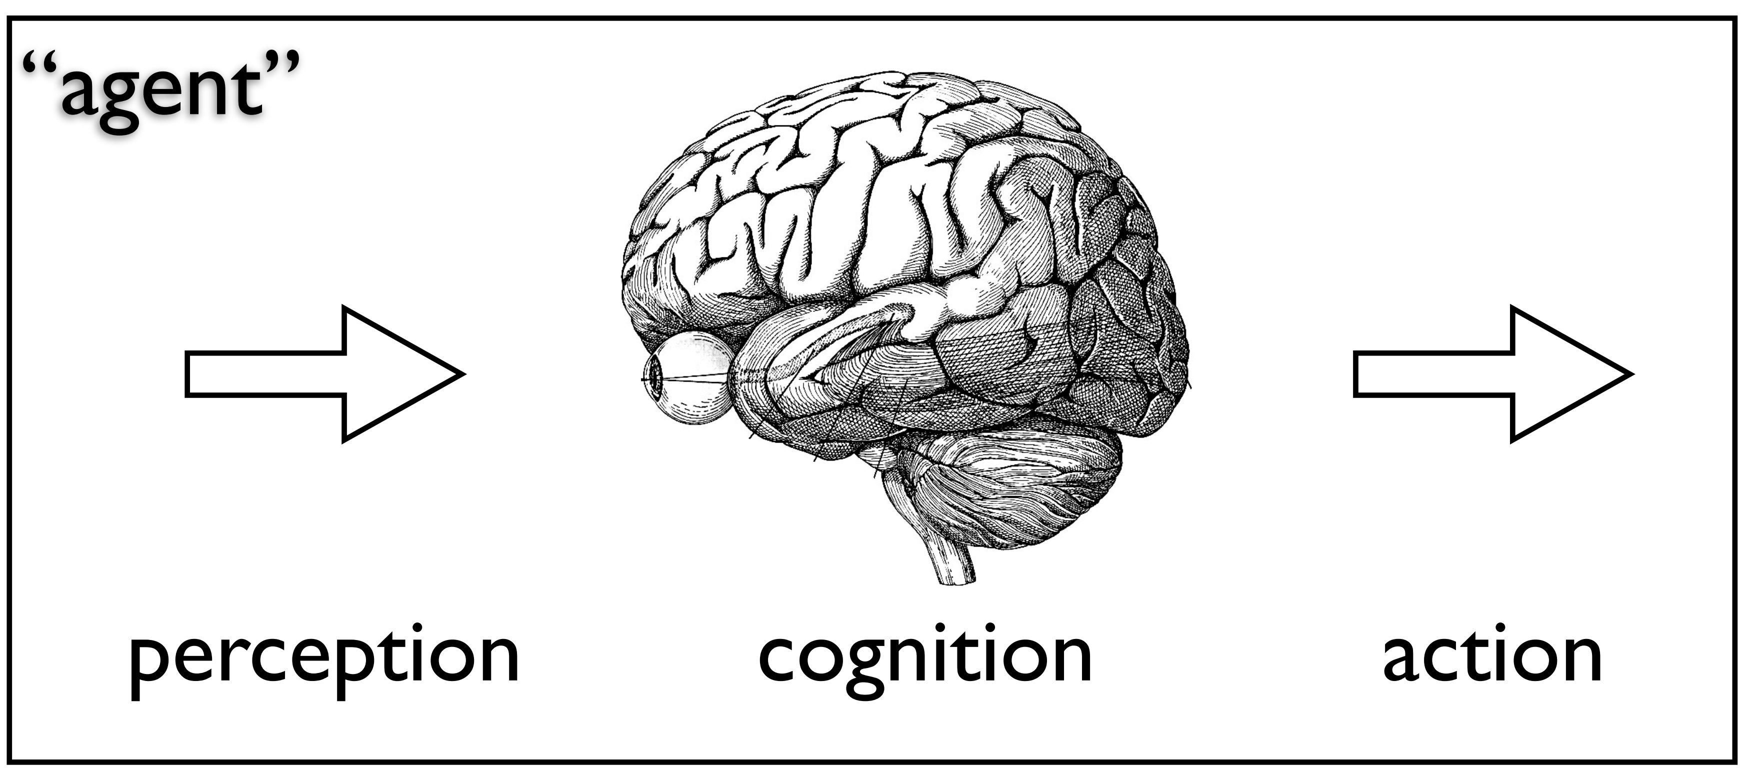

Representation

All AI problems require some form of representation.

- chess board
- maze
- text
- object
- room
- sound
- visual scene

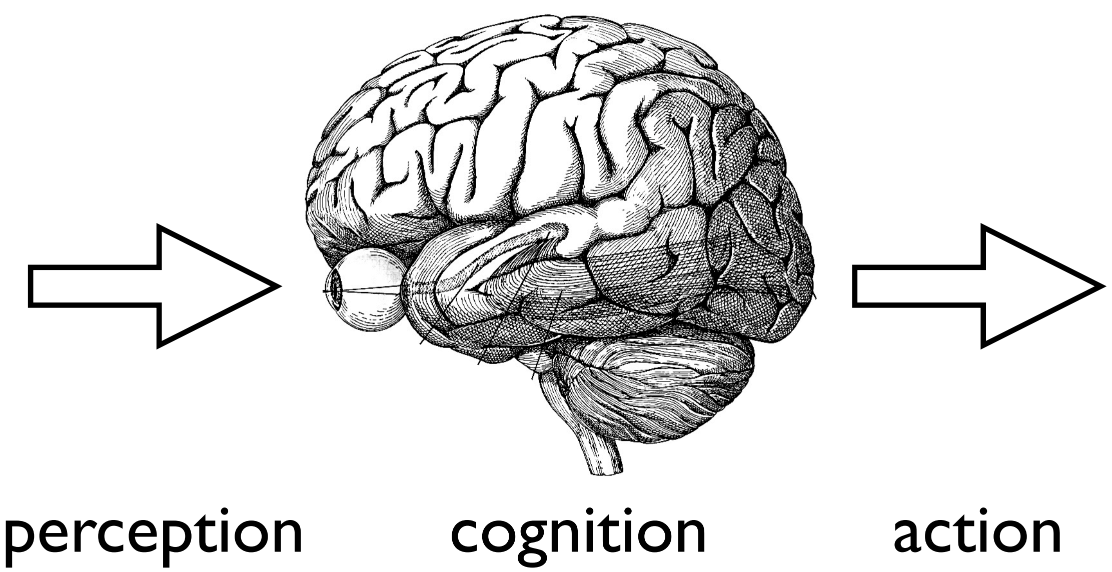

A major part AI is representing the problem space so as to allow efficient search for the best solution(s).

Sometimes the representation is the output.
E.g., discovering "patterns".

Output

The output action can also be complex.

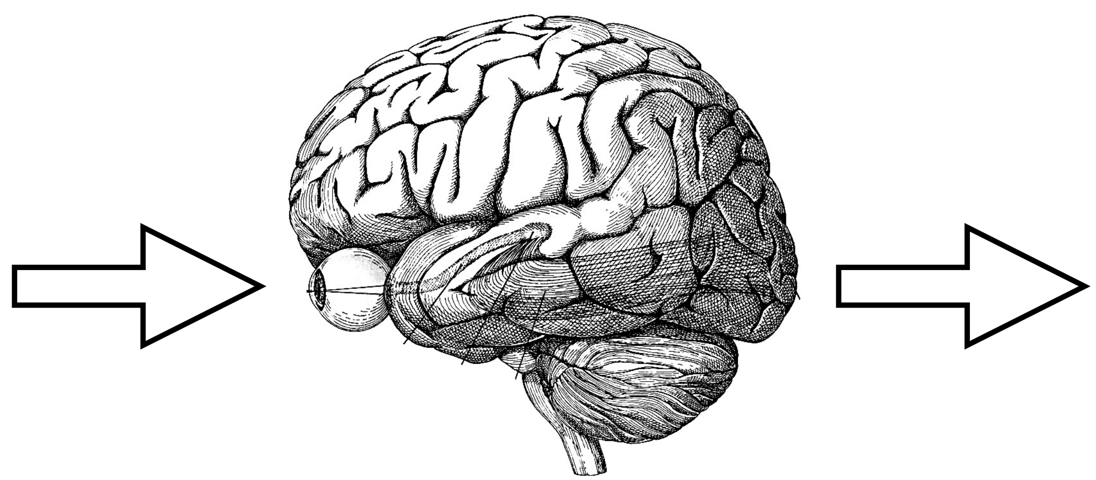

perception cognition action

- next move
- text
- label
- actuator
- movement

From a simple chess move to a motor sequence to grasp an object.

Russel and Norvig question 1.8

- Is AI's traditional focus on higher-level cognitive abilities misplaced?
  - Some authors have claimed that perception and motor skills are the most important part of intelligence.
  - "higher level" capacities are necessarily parasitic - simple add-ons
  - Most of evolution and the brain have been devoted to perception and motor skills
  - AI has found tasks such as game playing and logical inference easier than perceiving and acting in the real world.

Thinking

What do you do once you have a representation? This requires a goal.

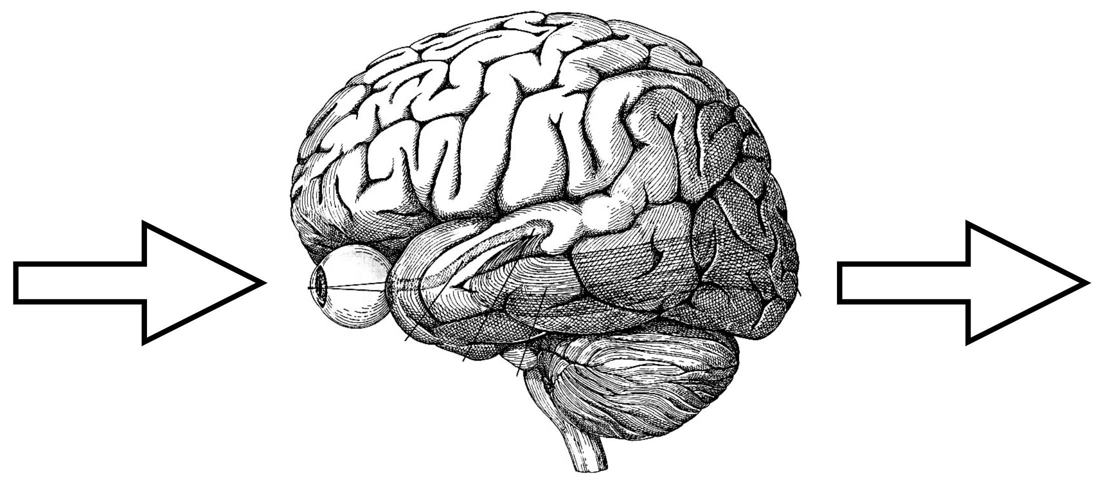

perception cognition action

- chess board
- maze
- text
- object
- room
- sound
- visual scene

- find best move
- shortest path
- semantic parsing
- recognition
- object localization
- speech recognition
- path navigation

Rational behavior:
choose actions that maximize goal achievement given available information

The Turing Test

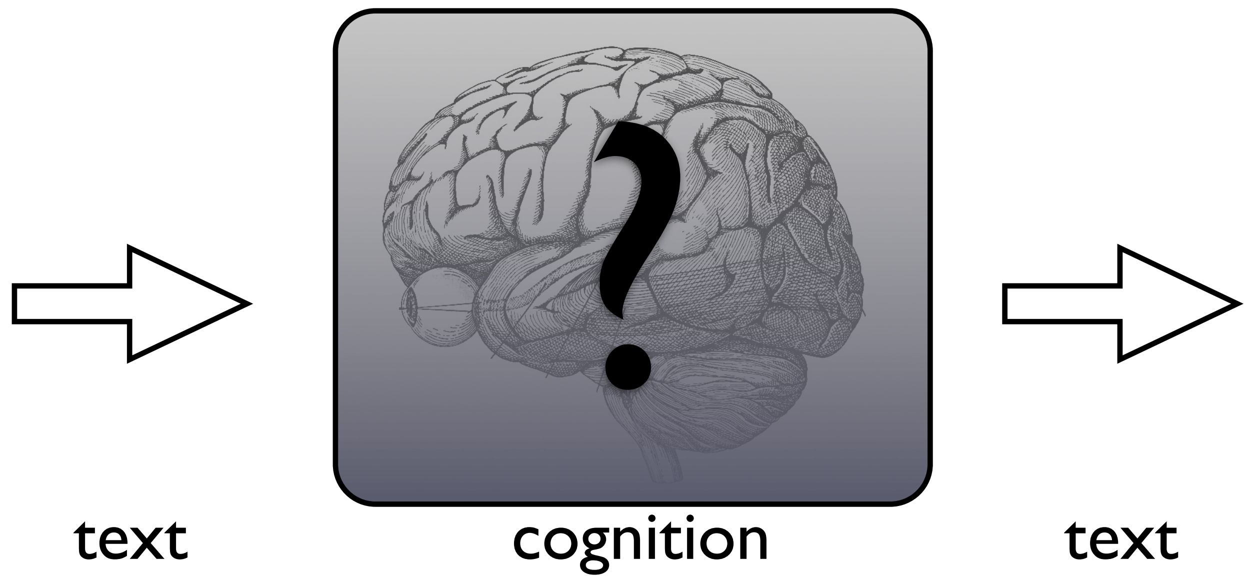

Strategy

What if your world includes another agent?

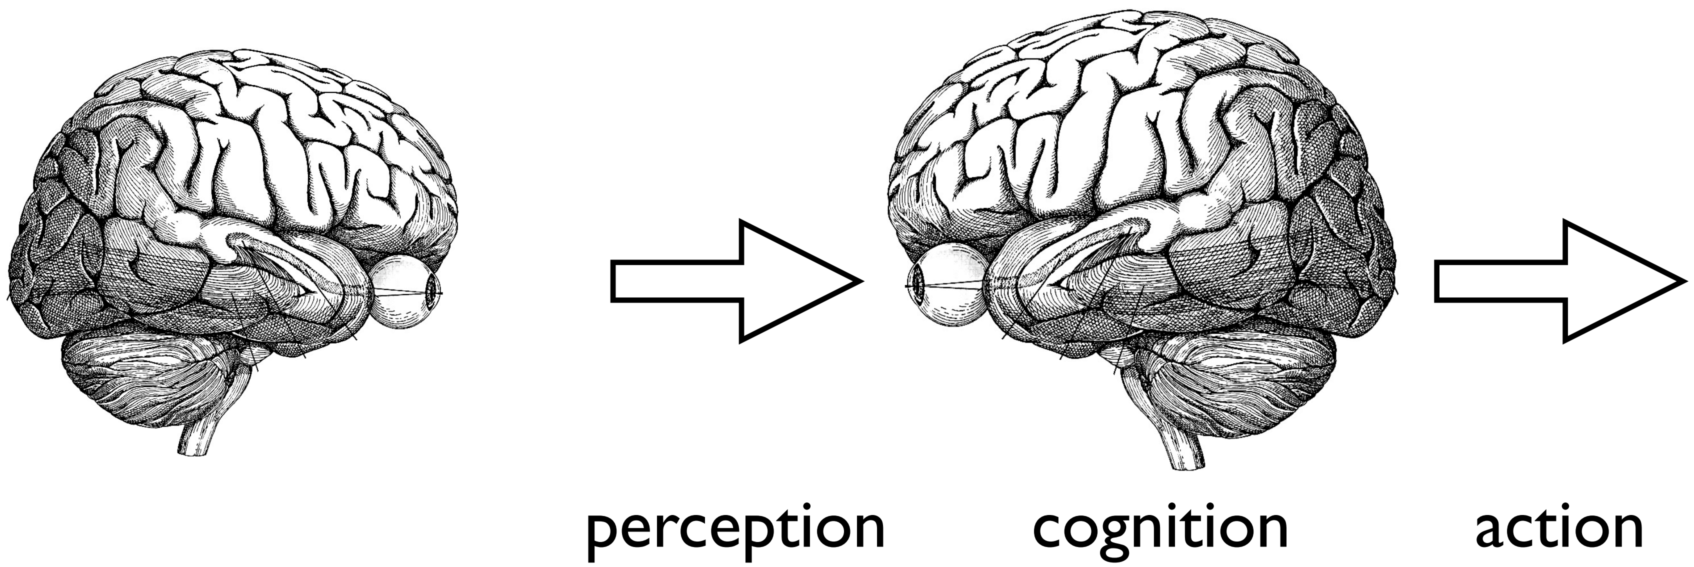

- strategic game play
- auctions
- modeling other agents
- uncertainty: chance and future actions

Rational behavior:
How do we choose moves/actions to win?

Or guarantee fairest outcome?

Team Play

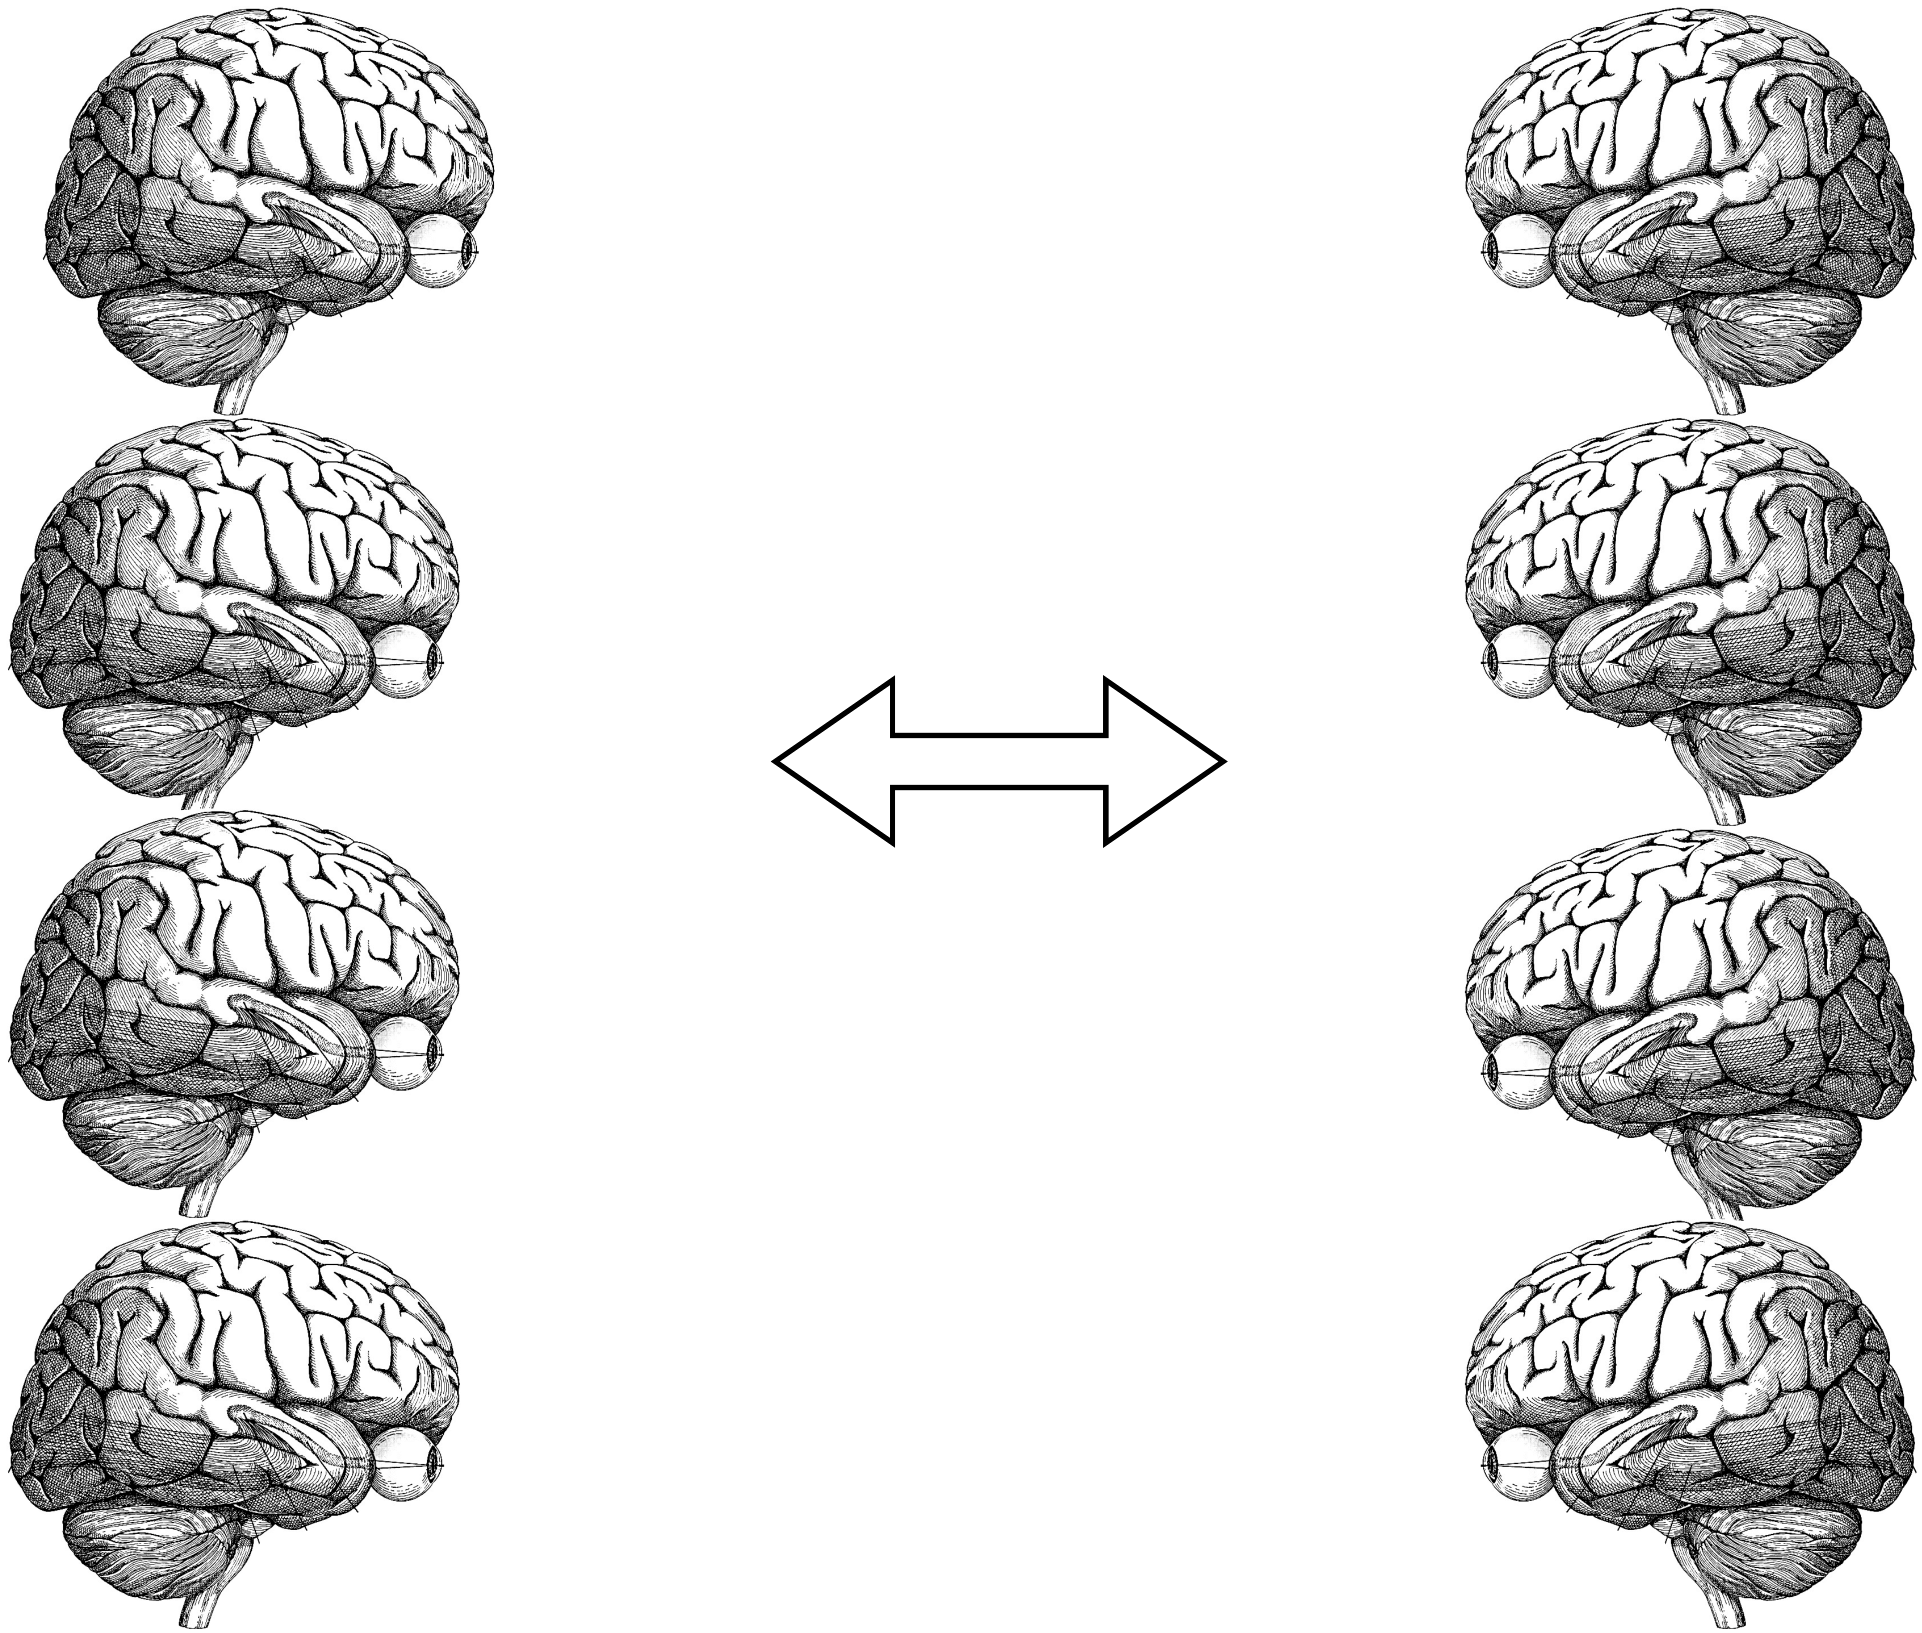

Reasoning

Reasoning can be thought of as constructing an accurate world model.

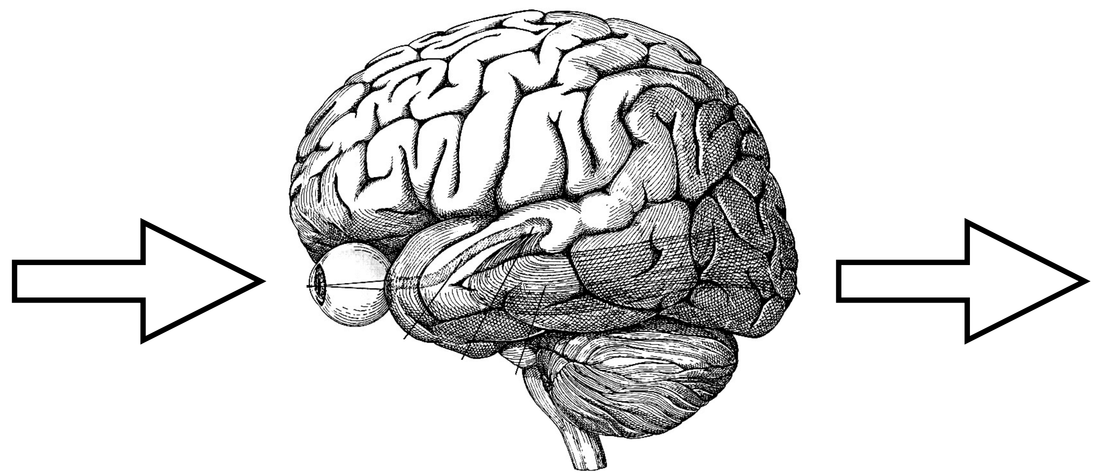

perception cognition action

- facts
- observations
- "wet ground"

- logical consequences
- inferences
- "it rained" or "sprinkler" ?

Rational inference:
What can be logically inferred give available information?

Reasoning with uncertain information

Most facts are not concrete and are not known with certainty.

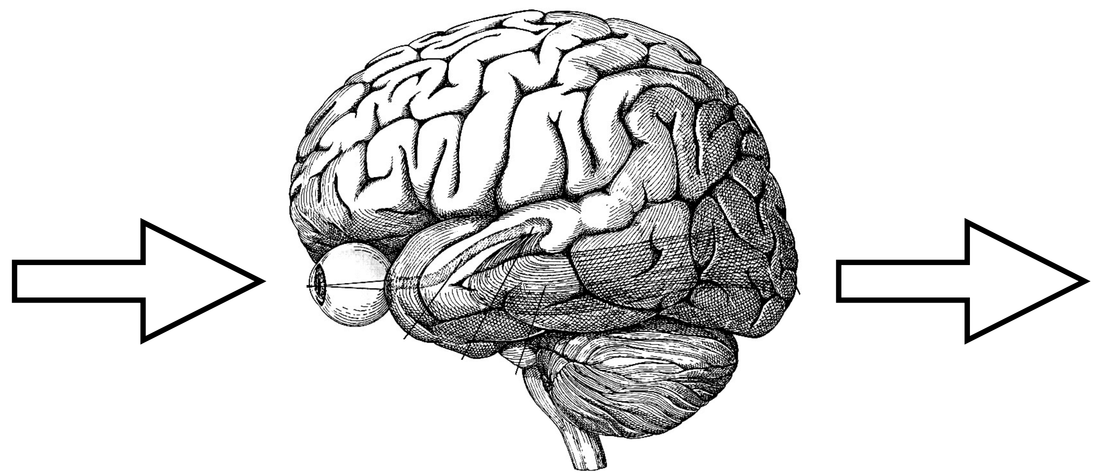

perception cognition action

- facts
- observations
- “fever”
- “aches”
- platelet count=N

- inferences
- What disease?
- What causes?

Probabilistic inference:
How do we give the proper weight to each observation?

What is ideal?

Learning

What if your world is changing? How do we maintain an accurate model?

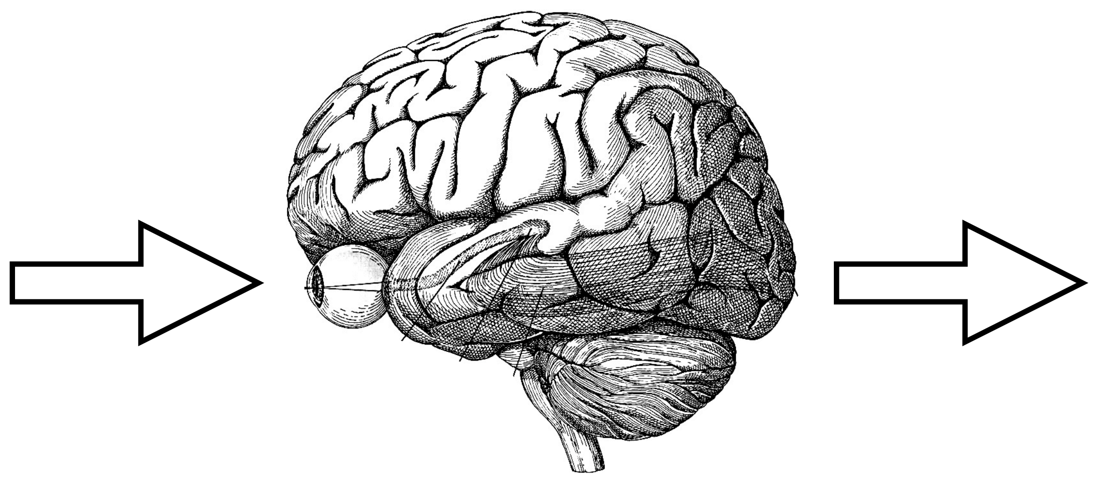

- chess board
- maze
- text
- object
- room
- sound
- visual scene

Learning:
adapt internal representation so that it is as accurate as possible.

Can also adapt our models of other agents.

Where can this go?

- Robotics
- Internet search
- Scheduling
- Planing
- Logistics
- HCI
- Games
- Auction design
- Diagnosis
- General reasoning

In class, we will focus on the AI fundamentals.

Brains vs computers revisited

Brains (adult cortex)	Computers (Intel Core 2)
• surface area: 2500 cm^2	• surface area: 90 mm^2
• squishy	• crystalline
• neurons: 20 billion	• transistors: 291 million
• synapses: 240 trillion	
• neuron size: 15 um	• transistor size: 65 nm
• synapse size: 1 um	
• synaptic OPS: 30 trillion	• FLOPS: 25 billion
• power usage: 12 W	• power usage: 60 W
• operations per joule: 2.5 trillion	• operations per joule: 0.4 billion

15-381 Artificial Intelligence

Martial Hebert
Mike Lewicki

Admin.

- **Instructor:**
  - Martial Hebert, NSH 4101, x8-2585
- **Textbook:**
  - Recommended (optional) textbook: Russell and Norvig's "Artificial Intelligence: A Modern Approach" (2^{nd} 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:**
  - **Rebecca Hutchinson (rah@cs.cmu.edu), WeH 3708, x8-8184**
  - **Gil Jones (egjones+@cs.cmu.edu), NSH 2201, x8-7413**
  - **Ellie Lin (elliel+15381@cs.cmu.edu), EDSH 223, x8-4858**
  - **Einat Minkov (einat@cs.cmu.edu), NSH 3612, x8-6591**
- **Grading:**
  - Midterm, Final, 6 homeworks

Admin.

- Class page:

http://www.cs.cmu.edu/afs/cs.cmu.edu/academic/class/15381-s07/www/

- Review sessions (look for announcements):

Tuesday 6:00pm-8:00pm in WeH 4623

Search

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

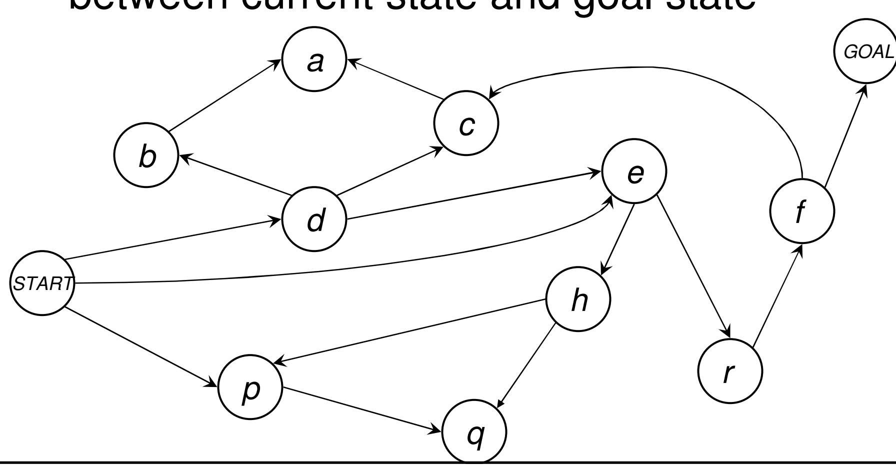

Search

- Uninformed search
- Informed search
- Constraint satisfaction

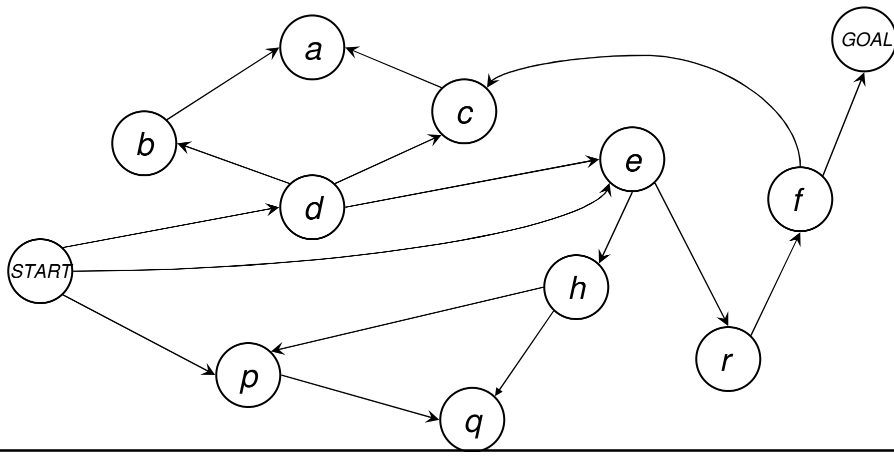

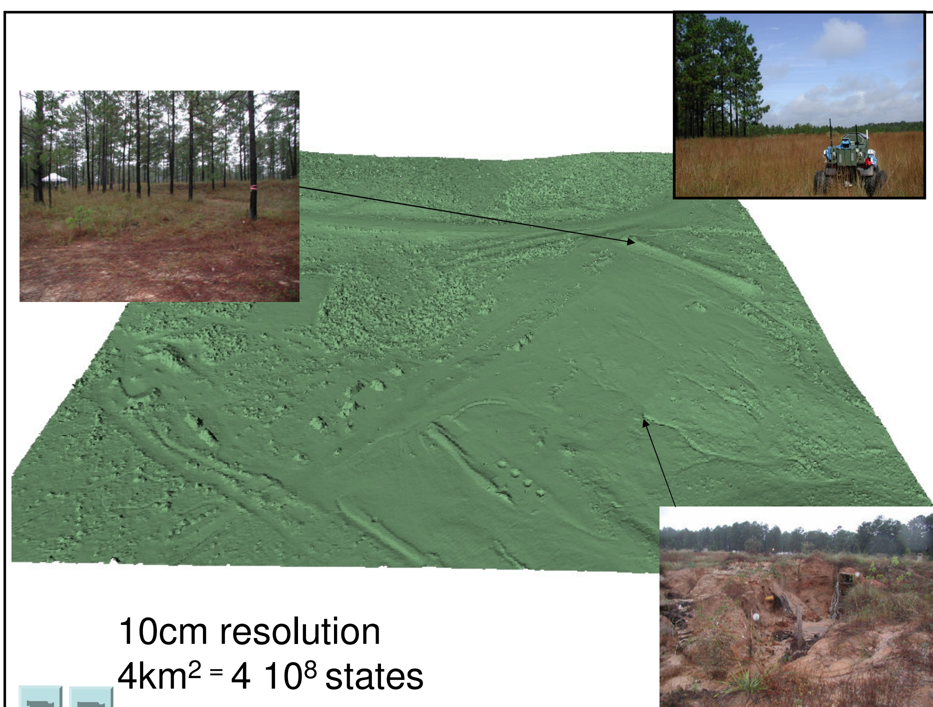

10cm resolution
$4km^2 = 4\ 10^8$ states

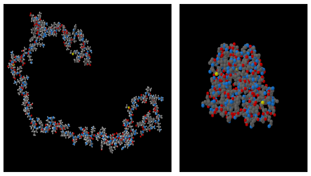

Protein design

http://www.blueprint.org/proteinfolding/trades/trades_problem.html

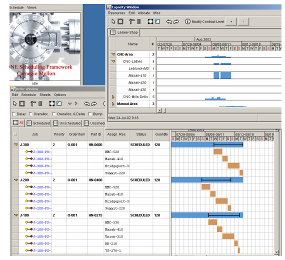

Scheduling/Manufacturing

http://www.ozone.ri.cmu.edu/projects/dms/dmsmain.html

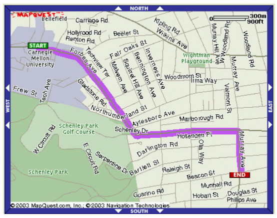

Route planning

Robot navigation

http://www.frc.ri.cmu.edu/projects/mars/dstar.html

Scheduling/Science

http://www.ozone.ri.cmu.edu/projects/hsts/hstsmain.html

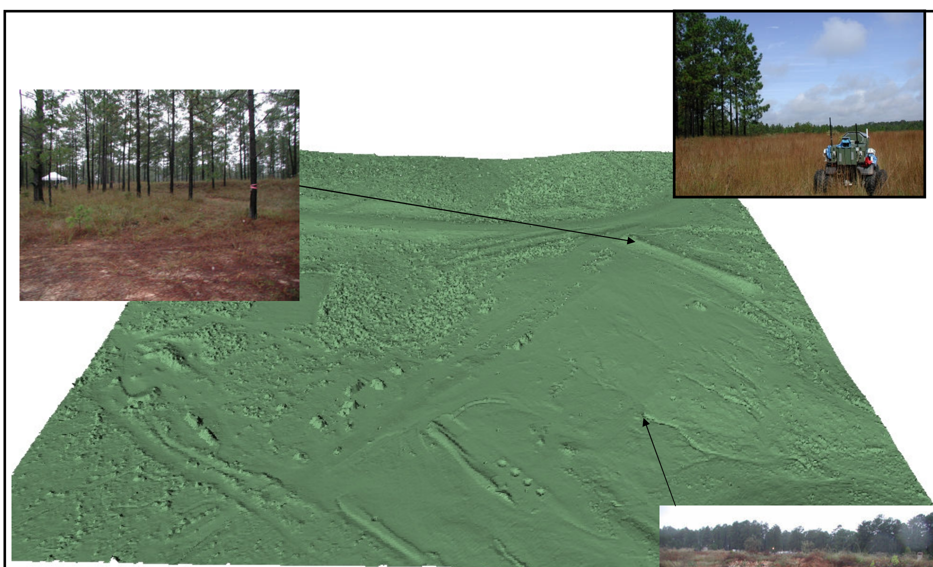

10cm resolution

$4km^2 = 4\ 10^8$ states

“Games”

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

- Business
- E-commerce
- Robotics
- Investment management
- …..

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

Reasoning with Uncertainty

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

Probability(Flu|TravelSubway)

Bayes Nets

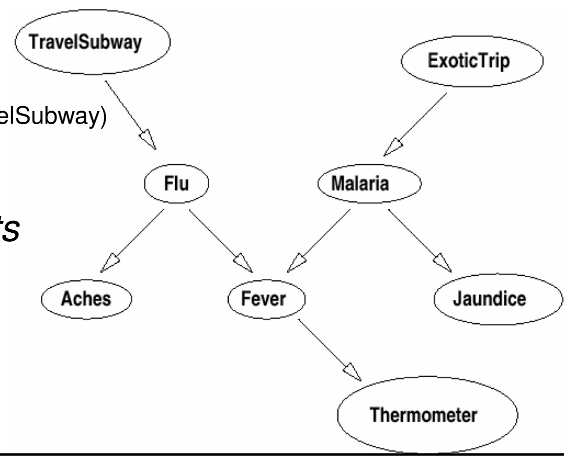

Learning

- 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	europe
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	europe

Training data: good/bad mpg for example cars

Mpg good/bad

Predict mpg on new data

Learning

- Automatically generate strategies to classify or predict from training examples

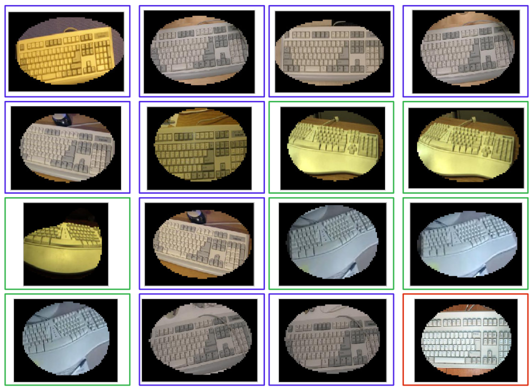

Training data: Example images of object

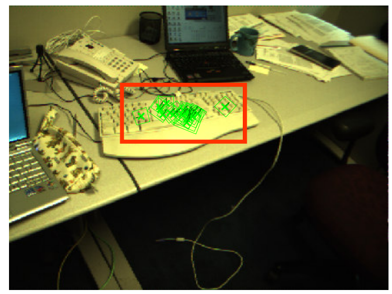

Classification: Is the object present in the input image, yes/no?

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

Date	Topic	Chapter	Notes
Jan. 16	Intro		
	SEARCH		
Jan. 18	Search	3	
Jan. 23	Search	3	HW1 out
Jan. 25	Search: Hill Climbing, Stochastic Search, Simulated Annealing	3,4	
Jan. 30	Search: Hill Climbing, Stochastic Search, Simulated Annealing	3,4	
Feb. 1	Constraint Satisfaction Problems	5	
Feb. 6	Constraint Satisfaction Problems	5	HW1 Due, HW2 out
Feb. 8	Robot Motion Planning	25	
	Game Theory		
Feb. 13	Algorithms for Playing and Solving Games	6	HW2 Due, HW3 out
Feb. 15	Games with Hidden Information	6	
Feb. 20	Non-Zero-Sum Games	6	
Feb. 22	Game Theory, continued	6	
Feb. 27	Auctions and Negotiations	6	
	SYMBOLIC REASONING		
Mar. 1	Automated Theorem Proving with Propositional Logic	8,9	
Mar. 6	Reasoning, Continued	11	
Mar. 8	MIDTERM		
Mar. 13	SPRING BREAK		
Mar. 15	SPRING BREAK		
	PROBABILISTIC REASONING		
Mar. 20	Probability and Uncertainty		
Mar. 22	Probability and Uncertainty		
Mar. 27	Bayes Nets	14	
Mar. 29	Bayes Nets	14	
Apr. 3	Markov Decision Processes	16,17	HW4 Due, HW5 out
Apr. 5	Markov Decision Processes	16,17	
	LEARNING		
Apr. 10	Intro + Decision Trees	18	
Apr. 12	Decision Trees, Cont.	18	
Apr. 17	Probabilistic Learning and Naive Bayes	20	HW5 Due, HW6 out
Apr. 19	Neural Networks	20	
Apr. 24	Cross-Validation	14	
Apr. 26	Nearest Neighbors	14	
May. 1	Reinforcement Learning	21	HW6 Due
May. 3	Reinforcement Learning	21	