15-381: AI
Introduction

Instructors: Manuela Veloso and Luis von Ahn

TAs: Sue Ann Hong, Gabriel Levi,
Mary McGlohon, and Abe Othman

http://www.cs.cmu.edu/afs/andrew/course/15/381-f09/www/

Carnegie Mellon

Grading

- **6 Problem sets - 50%**
- **Midterm - 20%**
- **Final - 30%**

- Problem sets can be done in groups of up to 2 people – no need to have the same group for all homeworks.

- 8 "mercy" days (no penalty) for late homeworks, cannot use more than 2 mercy days in a single homework. No credit for late homeworks with no mercy days.

Resources

- Lectures
  - Presentation and discussion in class
  - Lecture slides annotated and enriched by TAs with examples and further details
- Instructors – office hours by appointment
- TAs – office hours will be announced

What is *Artificial Intelligence*?

What is "intelligence" ?

Can we emulate

intelligent behavior in

machines ?

How far can we take it ?

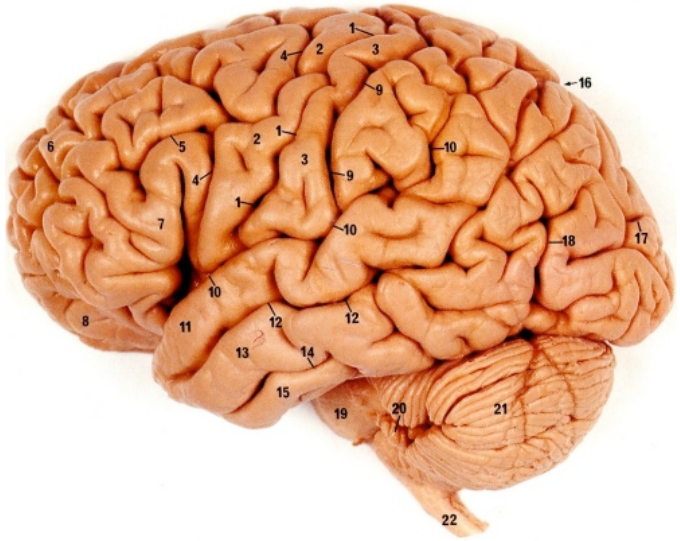

Intelligent Systems

Three key steps (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. internal representations are translated into action

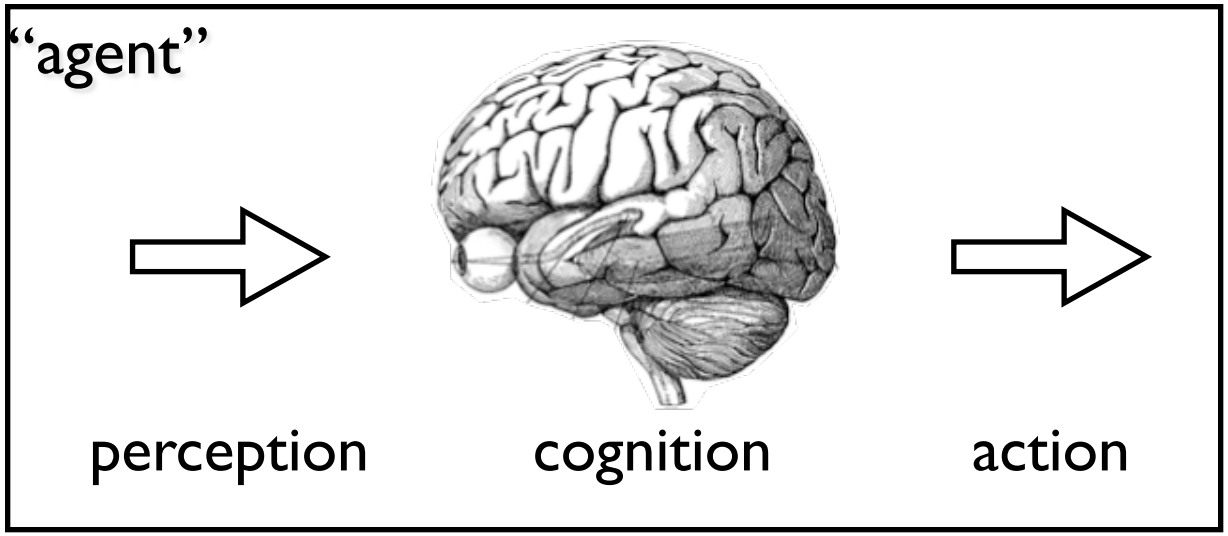

Views of AI

Think like humans	**Think rationally**
Cognitive Science	Formalize inference into laws of thought
Act like humans	**Act rationally**
Turing test	Act according to laws

Perception

Cognition

Action

AGENTS

Wean Hall 5409
Carnegie Mellon University
early 90s

Allen Newell *d.1992*

Artificial Intelligence

- Computer Science:
  - "*The study of computers and the phenomena that surround them.*"
    - Alan Perlis, Allen Newell, Herb Simon
- Ambitious scientific pursuits:
  - What is the nature of human intelligence?
  - How does the brain work?
  - How to solve problems effectively?
  - How do humans and machines learn?
  - How do we create intelligent creatures?

The Dartmouth Conference

"We propose that a two-month, ten-man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, NH. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it."

The Proponents

- John McCarthy, assistant professor of mathematics at Dartmouth (Stanford)
- Marvin Minsky, Harvard junior fellow in mathematics and neurology (MIT)
- Nathaniel Rochester, manager of information research at IBM, NY (?)
- Claude Shannon, information theory, mathematician at Bell Labs (2001)

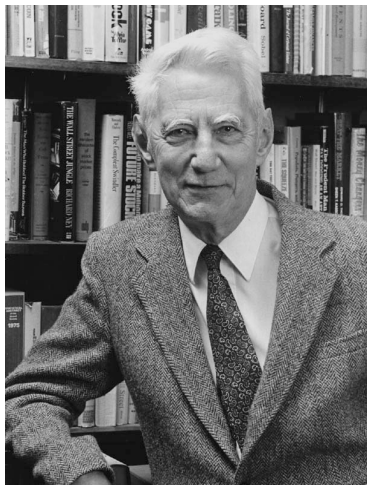

The Invited

- Trenchard More, IBM
- Arthur Samuel, IBM
- Oliver Selfridge, Lincoln Labs, MIT
- Ray Solomoff, MIT

And "*two vaguely known persons from RAND and Carnegie Tech… a significant afterthought.*" (Pamela McCorduck, "Machines Who Think", page 94)

Herbert A. Simon and Allen Newell

Problem Solving

- Allen Newell and Herb Simon – 1950s
  - Given:
    - an initial ***state***
    - a set of ***actions***
    - a ***goal*** statement
  - Find a ***plan***, a sequence of actions that transform the initial state into a state where the goal is satisfied

Search

- Find a sequence of states from current state to state that satisfies goal statement

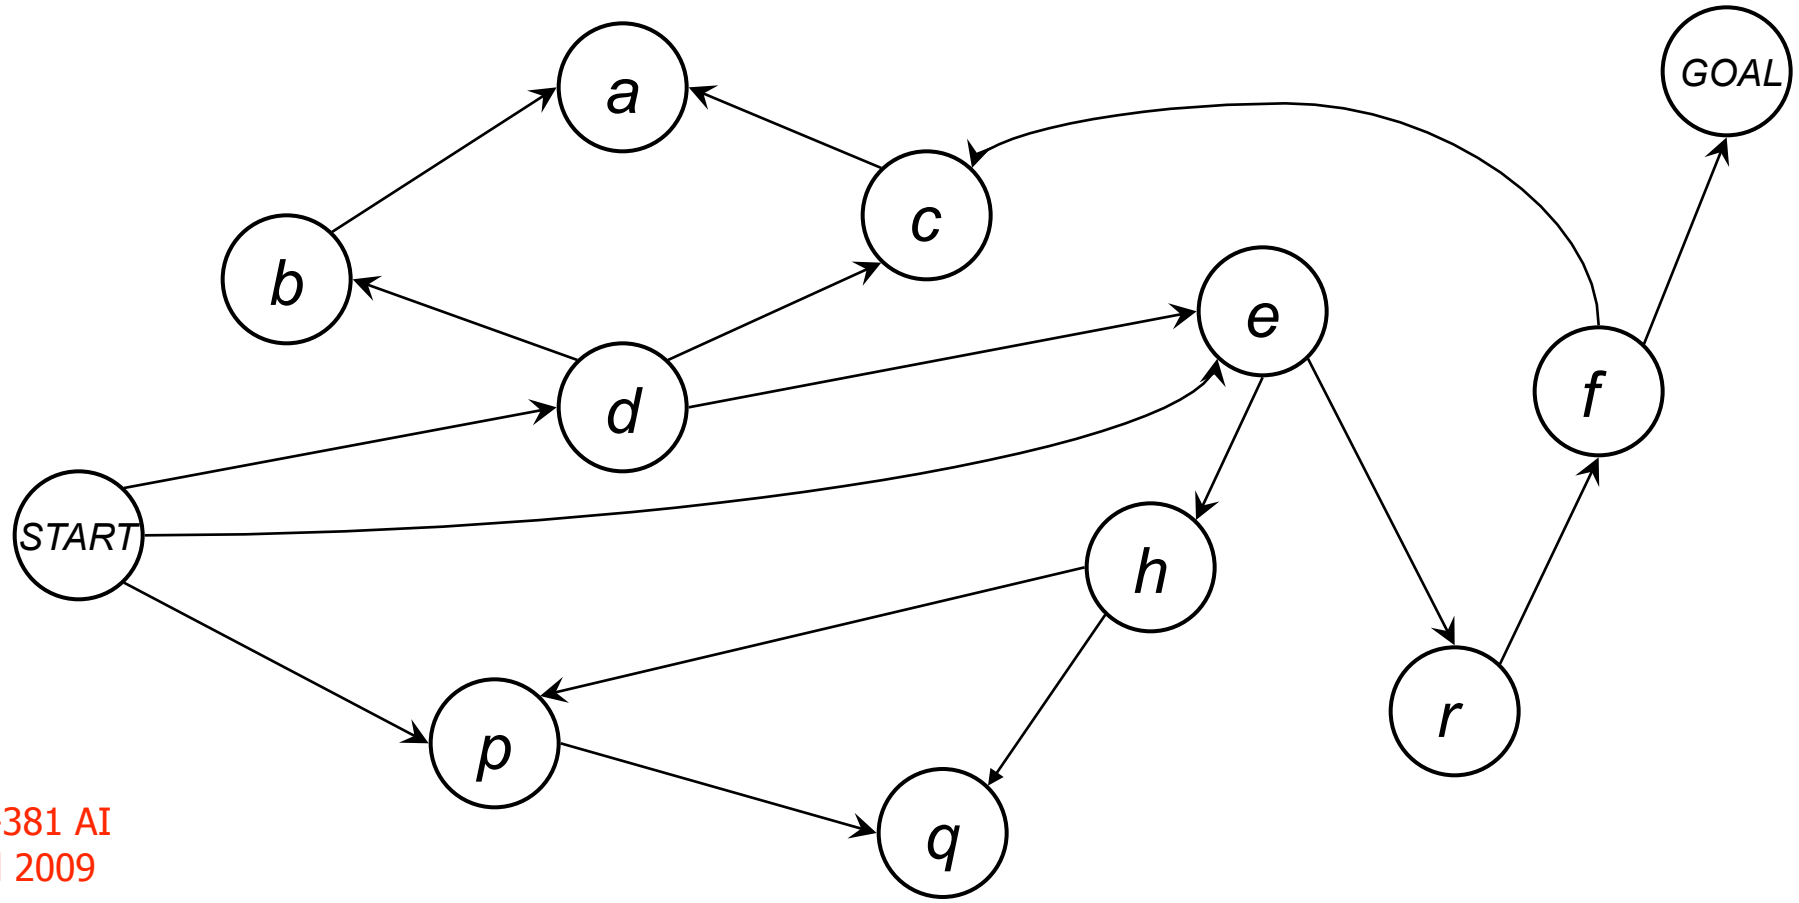

Schedule

- M Aug 24 – **Introduction**
- W Aug 26 – Uninformed **search** methods
- M Aug 31 – Informed search
- W Sep 2 – Stochastic search - HMW1 out
- M Sep 7 – No class, Labor's Day
- W Sep 9 – More search
- M Sep 14 – Constraint satisfaction problems
- W Sep 16 - CSPs - HMW1 due, HMW2 out

Problem Solving Components

- Given the *actions* available in a task domain.
- Given a problem specified as:
  - an initial *state* of the world,
  - a set of *goals* to be achieved.

Action Model, State, Goals

Actions, States, Goals

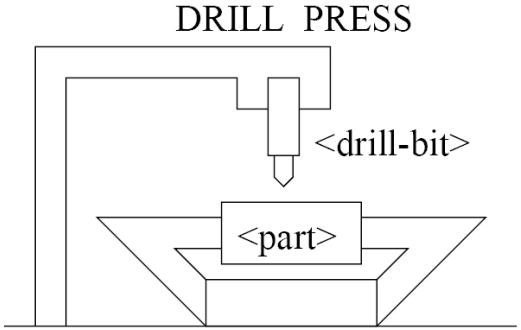

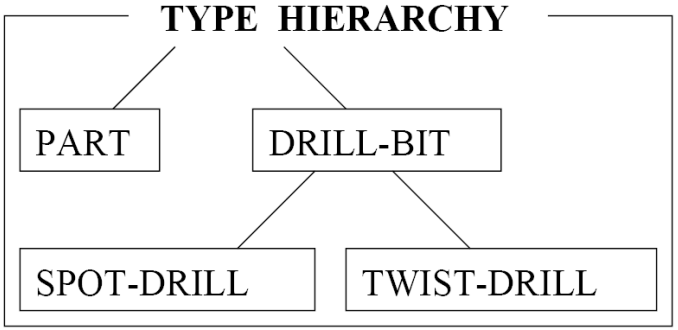

drill-spot (<part>, <drill-bit>)
<part>: type PART
<drill-bit>: type SPOT-DRILL
Pre: (holding-tool <drill-bit>)
(holding-part <part>)
Add: (has-spot <part>)

put-drill-bit (<drill-bit>)
<drill-bit>: type DRILL-BIT
Pre: tool-holder-empty
Add: (holding-tool <drill-bit>)
Del: tool-holder-empty

put-part(<part>)
<part>: type PART
Pre: part-holder-empty
Add: (holding-part <drill-bit>)
Del: part-holder-empty

drill-hole(<part>, <drill-bit>)
<part>: type PART
<drill-bit>: type TWIST-DRILL
Pre: (has-spot <part>)
(holding-tool <drill-bit>)
(holding-part <part>)
Add: (has-hole <part>)

remove-drill-bit(<drill-bit>)
<drill-bit>: type DRILL-BIT
Pre: (holding-tool <drill-bit>)
Add: tool-holder-empty
Del: (holding-tool <drill-bit>)

remove-part(<part>)
<part>: type PART
Pre: (holding-part <drill-bit>)
Add: part-holder-empty
Del: (holding-part <drill-bit>)

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).

Intelligent Agents

- **Sensing**: vision, hearing, touch, smell, taste, …
- **Cognition**: think, reason, plan, learn, …
- **Action**: motion, speak, manipulation, …
- **Interaction with other agents:** negotiation, strategic behavior, speculation, …

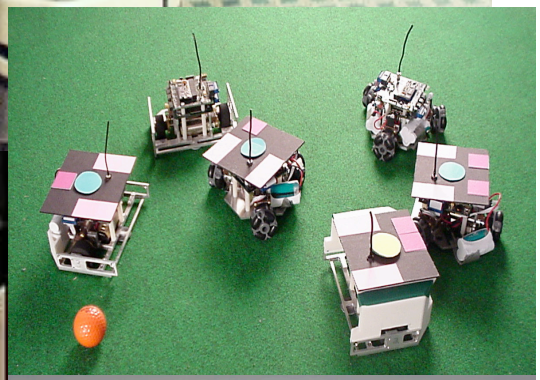

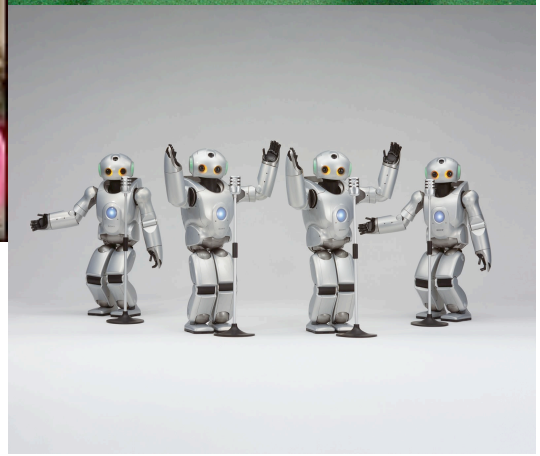

Perception – Sensors to State

- Sensors – "signal" (data) collectors from the physical world:
  - Vision, sound, touch, sonar, laser, infrared, GPS, temperature,….
- *Signal-to-symbol* challenge:
  - Recognize the *state* of the environment
  - …wall at 2m… door on the left… green light… person in front… personX entering the room… ball at 1m and 30^{o} East…

Reasoning with uncertain information

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

- 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?

Reasoning with Uncertainty

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

Probability(Flu|TravelSubway)

Bayes Nets

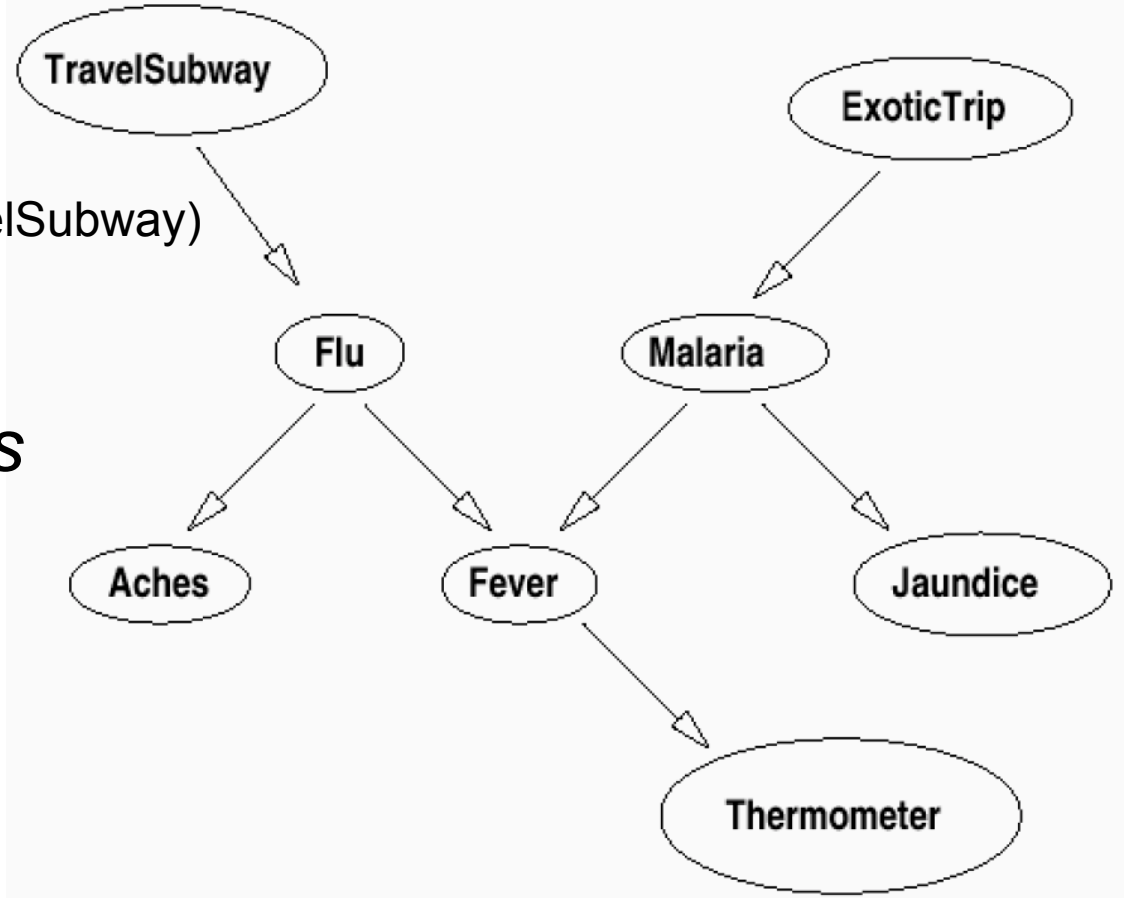

Schedule

- M Sep 21 – Deterministic reasoning, planning
- W Sep 23 – Uncertainty, robot motion planning
- M Sep 28 – Probability
- W Sep 30 – Bayesian networks - HMW2 due, HMW3 out
- M Oct 5 – Probabilistic reasoning
- W Oct 7 – Uncertainty HWM3 due, HMW4 out
- M Oct 12 – Review
- W Oct 14 – **MIDTERM**

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

bad	8	high	high	high	low	75to78	america
bad	6	medium	medium	medium	medium	70to74	america
good	4	low	low	low	low	79to83	america
bad	4	low	medium	low	medium	70to74	asia
bad	6	medium	medium	medium	high	75to78	america

?	8	high	high	high	low	75to78	america

➡ Mpg good/bad

Predict mpg on new data

Learning

- Automatically generate strategies to classify or predict from training examples

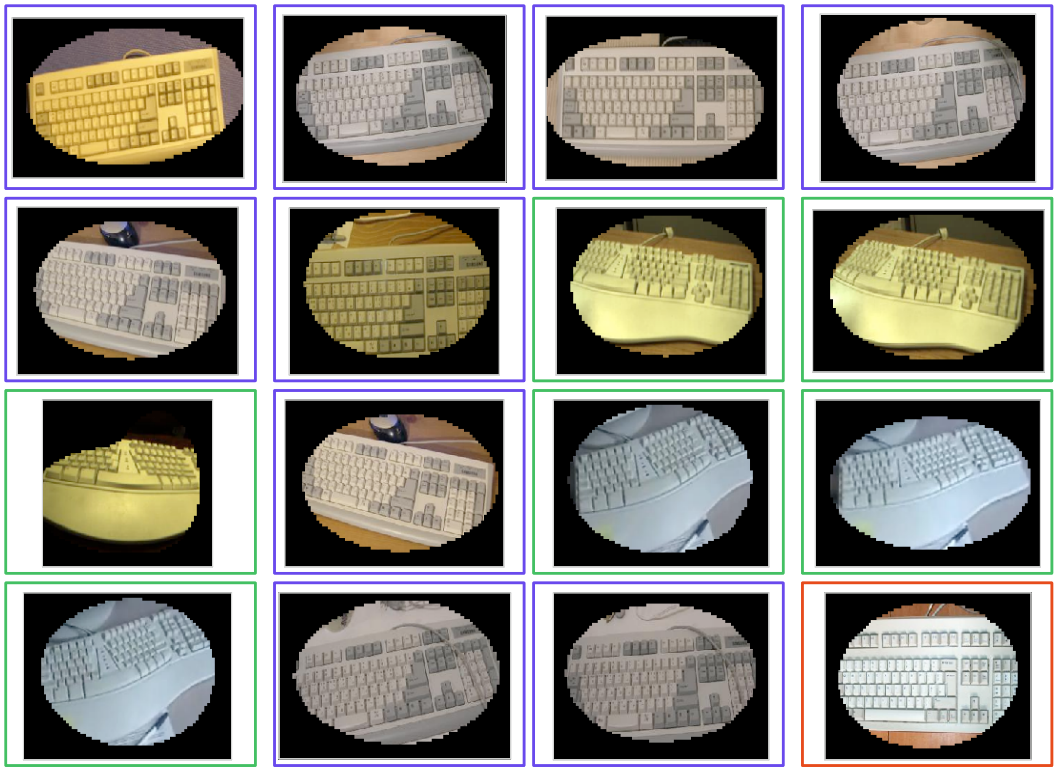

Training data: Example images of object

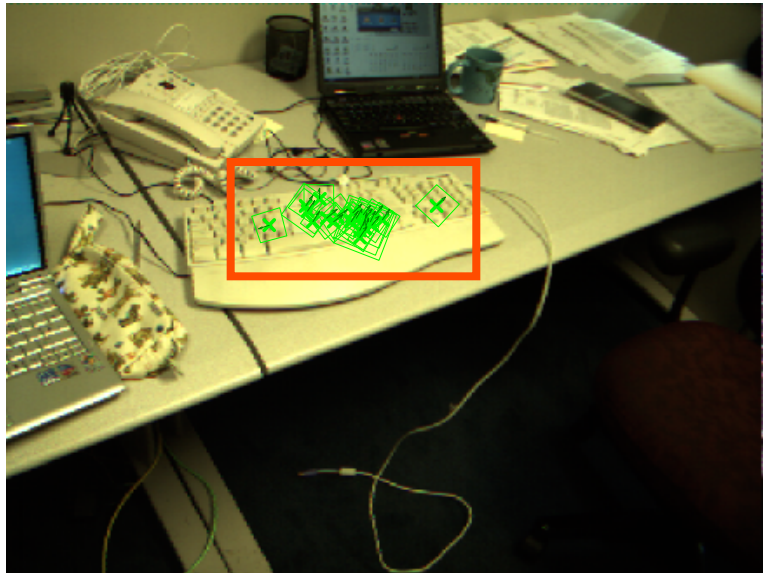

Classification: Is the object present in the input image, yes/no?

"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
- .....

Multiagent Systems and Learning

- How can an agent learn from experience in a world that contains other agents too ?
  - Other agents' learning makes the world nonstationary for the former agent
- Games
  - Learn to play Nash equilibrium
  - Learn to play optimally against static opponents

Schedule

- M Oct 19 – Decision Trees
- W Oct 21 – Decision Trees
- M Oct 26 – Neural Nets
- W Oct 28 – Robot Learning, HMW4 due, HMW5 out
- M Nov 2 – Classification
- W Nov 4 – Clustering
- M Nov 9 – Support Vector Machines
- W Nov 11– Markov Decision Processe, HMW5 due, HMW6 out
- M Nov 16 – MDPs
- W Nov 18 – Reinforcement learning
- M Nov 23 – Game theory, multiagent systems
- W Nov 24 – No class, Thanksgiving
- M Nov 30 – Multi-robot systems
- W Dec 2 – Review – WrapUp

- Final Exam – TBA

- Mon, Aug 24: Introduction, Search
- Wed, Aug 26: Uninformed search methods
- Mon, Aug 31: Search - informed methods
- Wed, Sep 2: Search, hill climbing, Homework 1 out
- Mon, Sep 7: NO CLASS - Labor day
- Wed, Sep 9: Search
- Mon, Sep 14: Constraint satisfaction problems (CSPs)
- Wed, Sep 16: Homework 1 due: Constraint satisfaction problems (CSPs) , Homework 2 out
- Mon, Sep 21: Symbolic reasoning, planning
- Wed, Sep 23: Uncertainty, robot motion planning
- Mon, Sep 28: Probability
- Wed, Sep 30: Bayesian networks, Homework 2 due. Homework 3 out
- Mon, Oct 5: Uncertainty
- Wed, Oct 7: Probability, Homework 3 due, Homework 4 out
- Mon, Oct 12: Midterm review
- Wed, Oct 14:**Midterm Exam**
- Mon, Oct 19: Decision trees, neural networks
- Wed, Oct 21: Decision Trees, cont.
- Mon, Oct 26:: Neural Networks
- Wed, Oct 28: Robot learning, Homework 4 due, Homework 5 out
- Mon, Nov 2: Clustering
- Wed, Nov 4: Support Vector Machines
- Mon, Nov 9: Markov Decision Processes (MDPs)
- Wed, Nov 11: Markov Decision Processes (MDPs), Homework 5 due, Homework 6 out
- Mon., Nov 16:Reinforcement Learning
- Wed, Nov 18: Reinforcement Learning
- Mon, Nov 23: Game Theory
- Wed, Nov 25: NO CLASS - Thanksgiving
- Mon, Nov 30; Game theory, multi-agent, multi-robot systems
- Wed, Dec 2: Final review, wrap-up