Artificial Intelligence CPSC 470/570

Brian Scassellati

scaz@cs.yale.edu

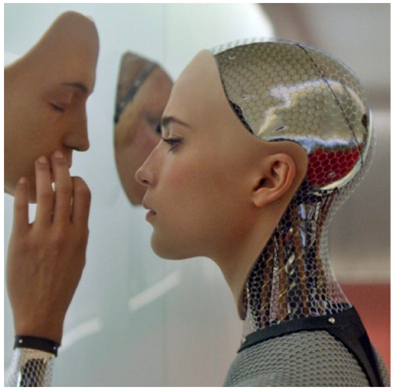

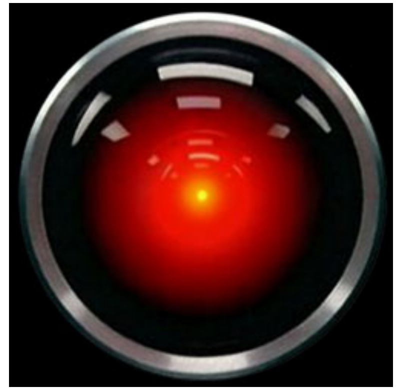

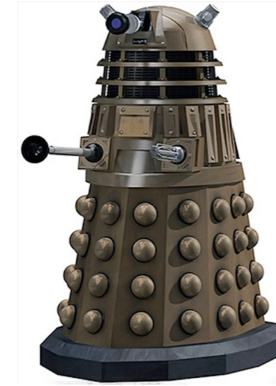

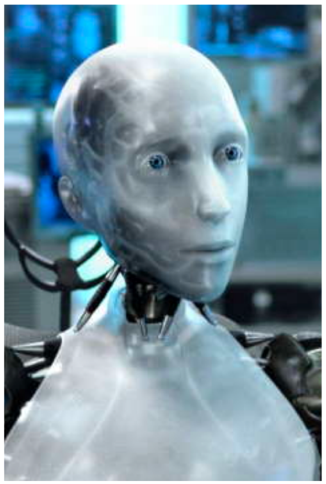
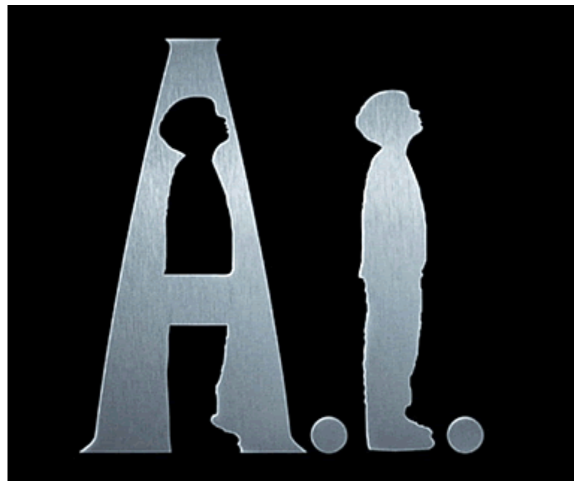
A.I.



REJECT
CHAPPIE

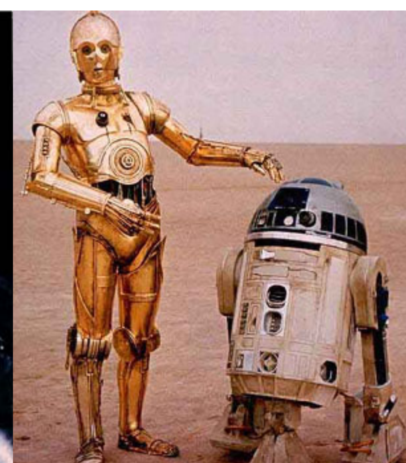
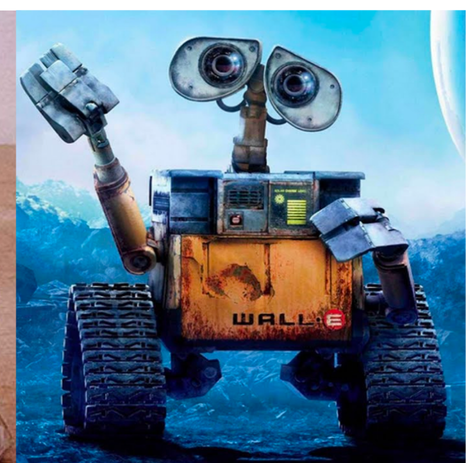
WALL·E
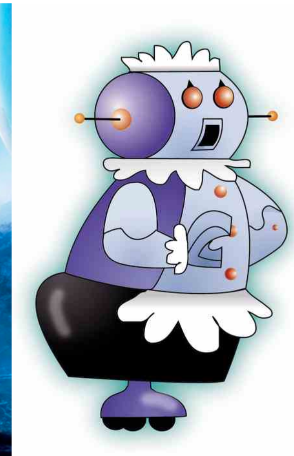

AI as Game Playing

Checkers: Chinook vs. Tinsley (1994)

Chess: IBM Deep Blue vs. Kasparov (1997)

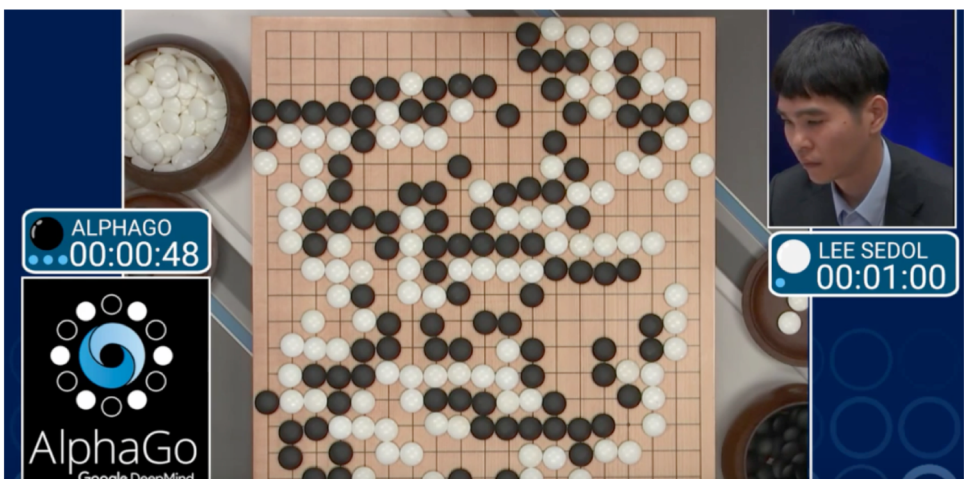

Go: Google AlphaGo vs. Lee Sedol (2016)

Poker (No-limit hold'em): CMU Libratus (2017)

AI is also

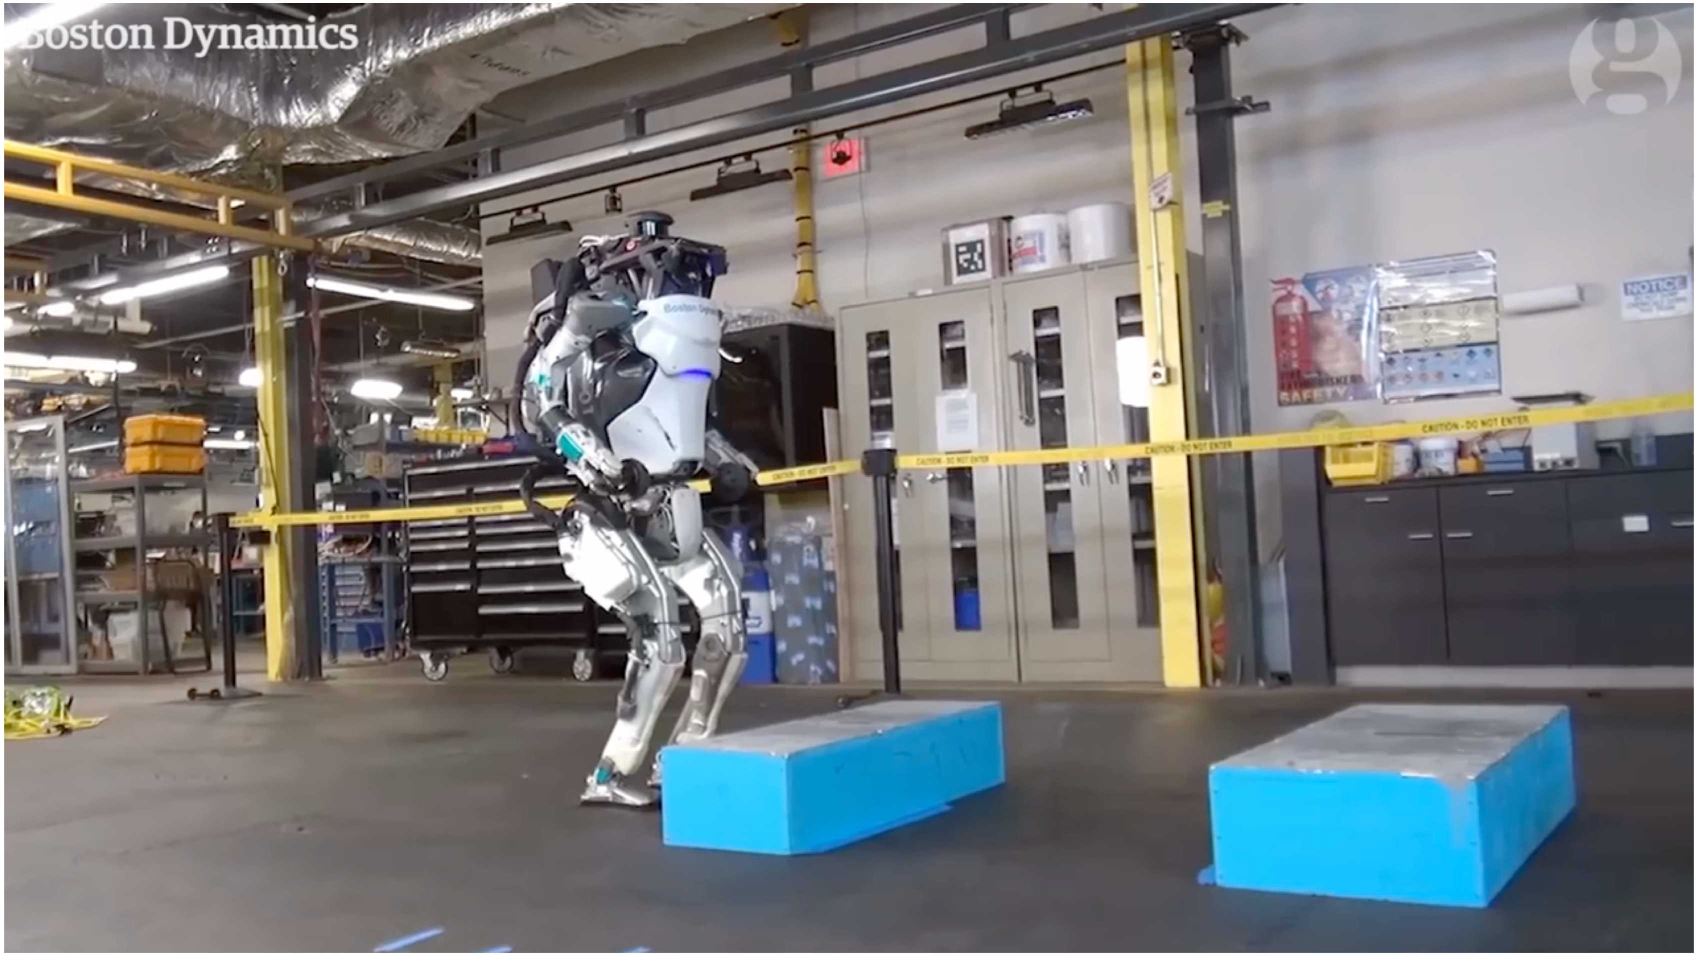

Source: Boston Dynamics. https://www.youtube.com/watch?v=WcbGRBPkrps

Definitions of AI

Think like Humans

"The automation of activities that we associate with human thinking, activities such as decision-making, problem solving, learning…" – Bellman, 1978

Think Rationally

"The study of mental faculties through the use of computational models" – Charniak and McDermott, 1985

Thought

Act like Humans

"The art of creating machines that perform functions that require intelligence when performed by people." – Kurzweil, 1990

Act Rationally

"A field of study that seeks to explain and emulate intelligent behavior in terms of computational processes." – Schalkoff, 1990

Action

Defined in terms of Humans

Defined in terms of Logic

What AI was…

- Spreadsheets
- Graphical interfaces
- Icon-oriented interfaces
- Object-oriented programming languages
- Sketching software
- Automated theorem provers

and every robotics, vision, natural language, sound processing and reasoning project…

AI is a Moving Target

Hype

AI Headlines from today (1/14/19)

AI beats expert doctors at finding cervical pre-cancers - Tech News

The Star Online • today

IBM's AI Machine Makes A Convincing Case That It's Mastering The Human Art Of Persuasion

Forbes • today

The Future of Artificial Intelligence In The Workplace

Forbes • 2 days ago

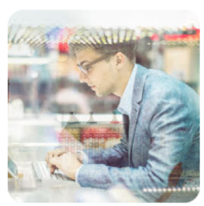

How AI is making business travel better

CNN • 5 days ago

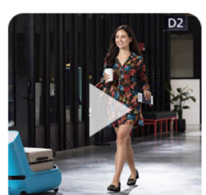

Most Kiwi staff see AI as a threat rather than an opportunity: survey

CIO New Zealand • today

Remember Elon Musk's Scary Warning Against AI? Here's More Reason to Worry.

Entrepreneur • 3 days ago

Commentary: Bad news. Artificial intelligence is biased

Channel NewsAsia • 2 days ago • Opinion

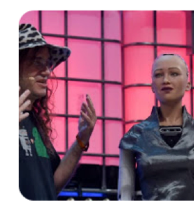

Never mind killer robots—here are six real AI dangers to watch out for in 2019

MIT Technology Review • 6 days ago

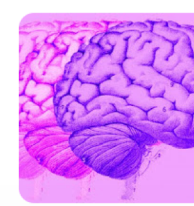

The good, the bad, and the ugly

- AI is the new electricity!
  – Andrew Ng, Chief scientist Baidu
- Will robots take our children's jobs?
  – NYT, Dec 11, 2017
- Bill Gates: AI taking everyone's jobs will be a good thing
  – Business Insider, Jan 25, 2018
- AI is more dangerous than nuclear weapons
  – Elon Musk at SXSW, Mar 13, 2018
- Stephen Hawking: AI could destroy civilization!
  – Newsweek, Nov 7, 2017

Growth of AI

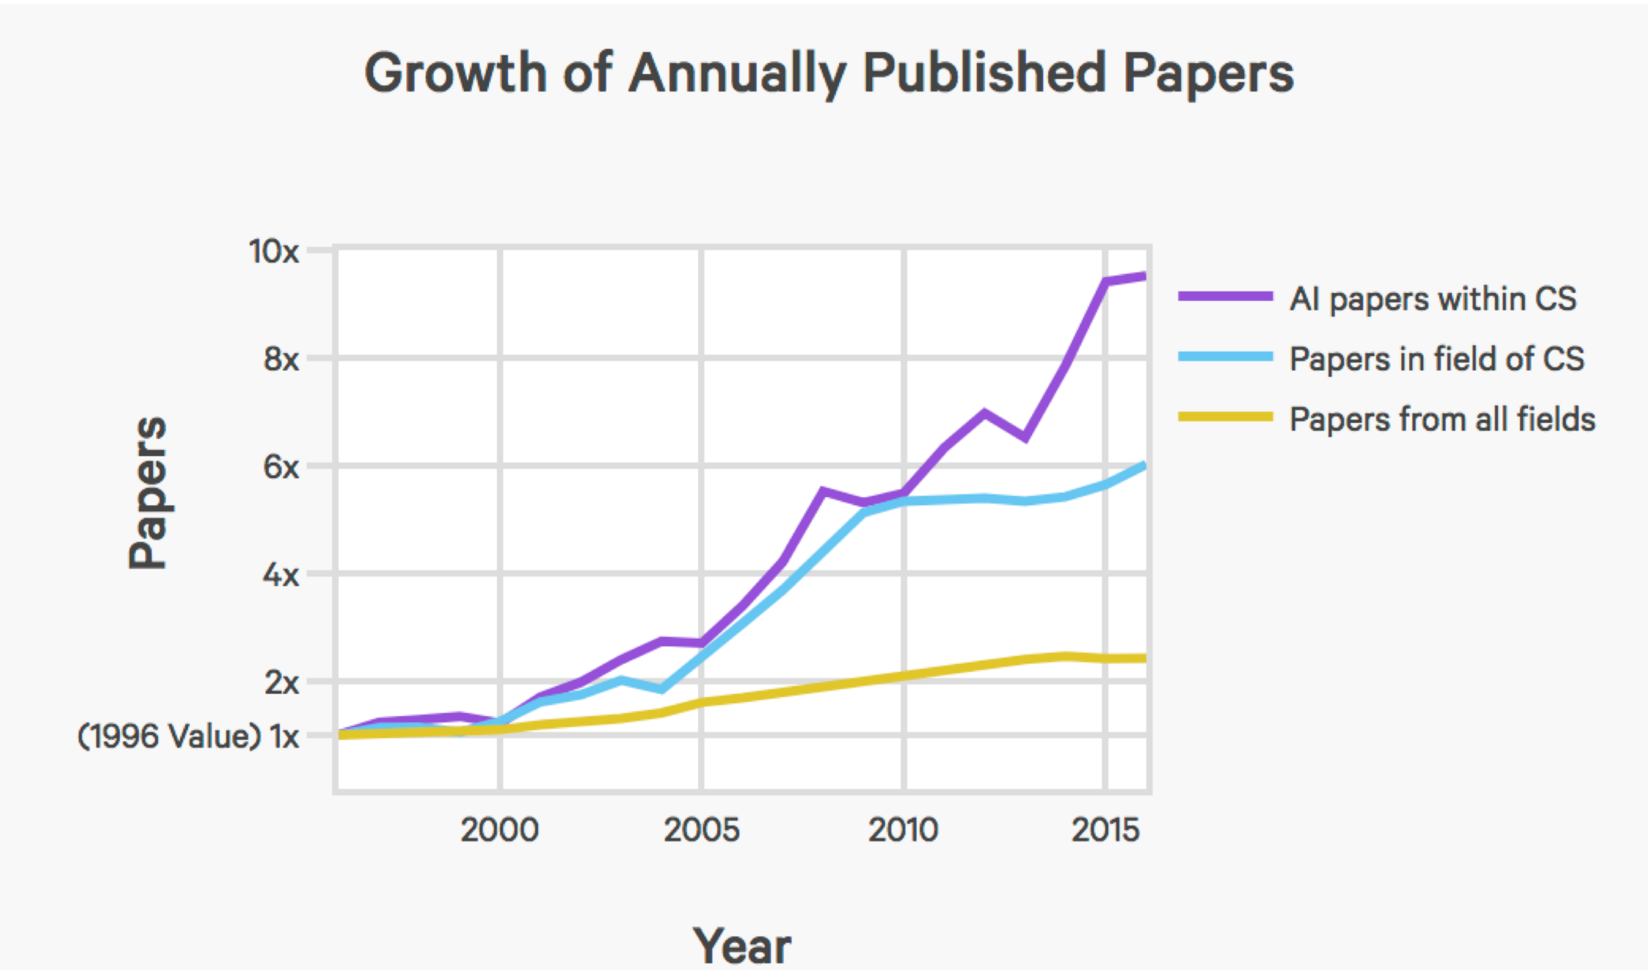

Source: Scopus.com, via 2017 AI index (http://aiindex.org/)

Growth of AI: Large Corporate AI Investments

- Late 2015: Toyota announces $1b USD investment in AI
- Hired leadership:
  - CEO Gil Pratt, former DARPA PM
  - CTO James Kuffner, former Google autonomous vehicle lead
- Feb 2017: two systems announced
  - Chauffeur (level 4/5 autonomy)
  - Guardian (level 1/2 driver assist)

Source: http://pressroom.toyota.com/releases/tri+autonomous+test+vehicle+sonoma+raceway+prius+challenge.htm

Growth of AI: Startup Funding Soaring

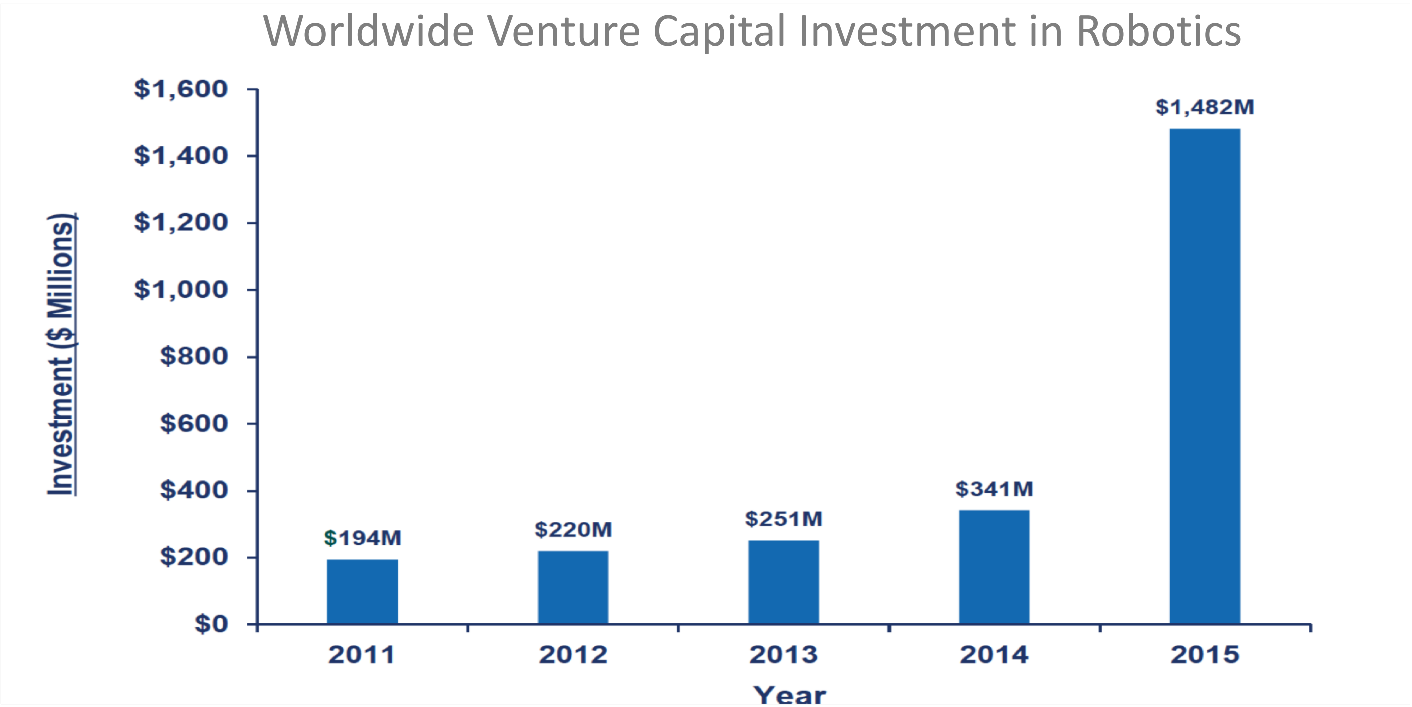

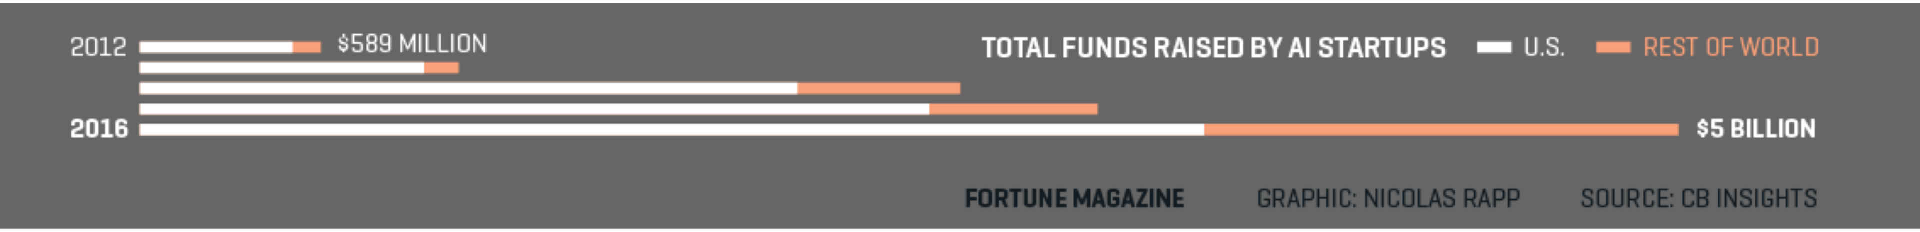

Source: (top) Hizook, Jan 12 2016, Funding for Robotics in 2015 (bottom) Fortune, Mar 1 2017. “Betting on AI”.

Growth of AI: Unprecedented Hiring

- "Universities' AI Talent Poached by Tech Giants"
  - WSJ, 11/24/16
- "Giant Corporations are Hoarding the World's AI Talent"
  - Wired, 11/16/16
- "Over 4,000 Artificial Intelligence job roles vacant on shortage of talent
  - Forbes, 12/18/18

- Median annual salary (source: NSF)
  - $55,000 post-doc in academia
  - $110,000 in industry labs

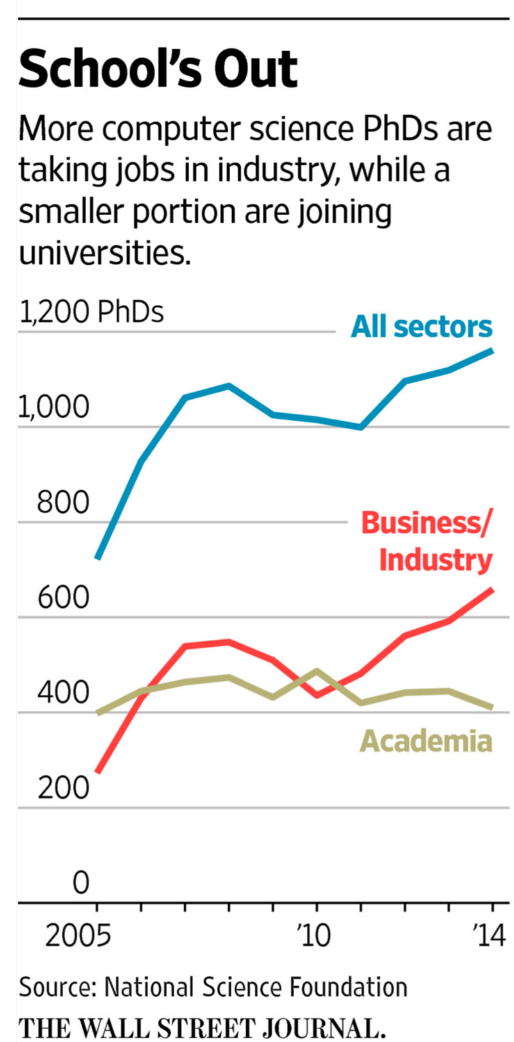

Why now?

- Access to massive amounts of data
- Access to powerful computing platforms
  - Multicore chips
  - Ubiquitous cellphones and tablets
  - Cloud computing
- Maturity of robotics hardware

Syllabus

- Approximately one week for each of these topics:
  - Search
  - Game Playing
  - Logical Formalisms
  - Inference
  - Planning
  - Dealing with Uncertainty
  - Machine Learning
  - Communication and Language
  - Perception
  - Robotics

Agents as a Unifying Design

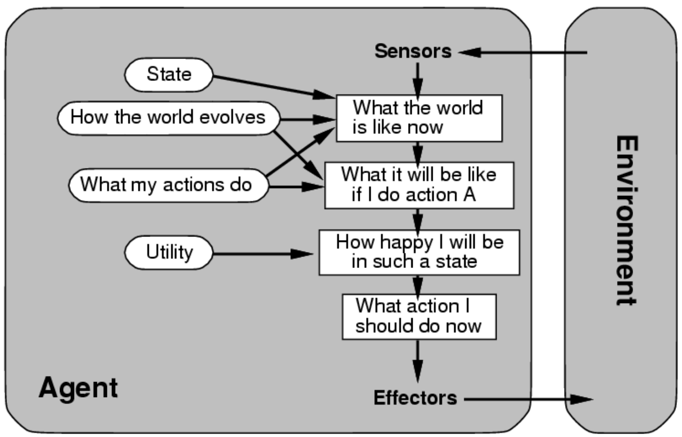

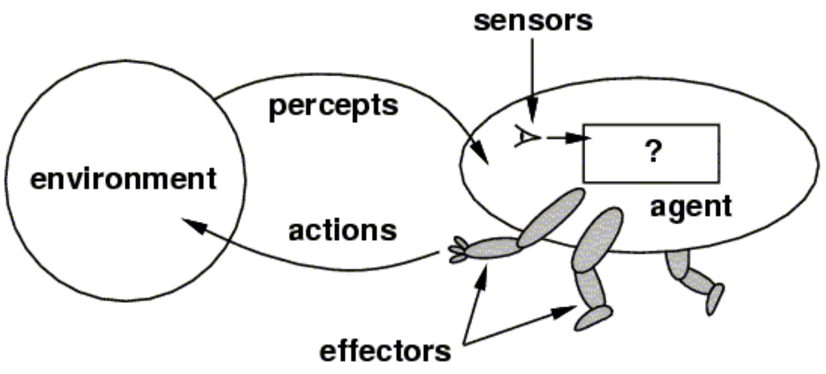

Environment	Accessible	Deterministic	Episodic	Static	Discrete
	Do sensors give complete world state?	Can next state be determined by current state and action?	Does quality of an action depend only on current state?	Does the env. stay the same while the agent thinks?	Are the number of percepts and actions limited?
Chess (no clock)	Yes	Yes	No	Yes	Yes
Poker	No	No	No	Yes	Yes
Taxi driving	No	No	No	No	No
Image analysis	Yes	Yes	Yes	Semi	No
Part-picking robot	No	No	Yes	No	No
Refinery controller	No	No	No	No	Yes

Basic Search

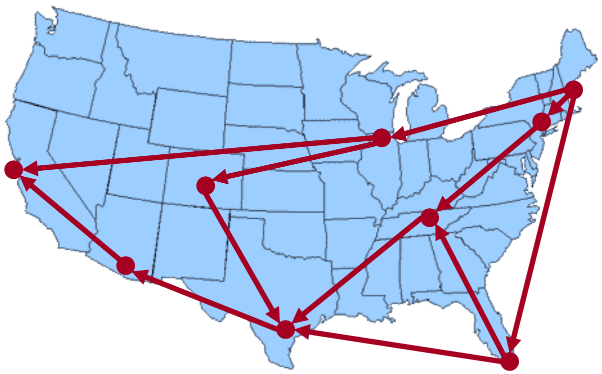

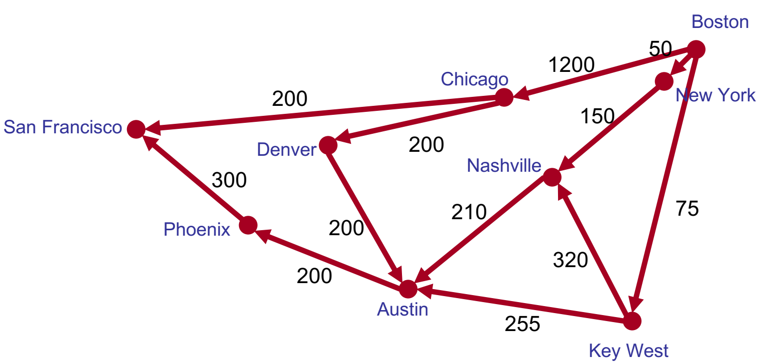

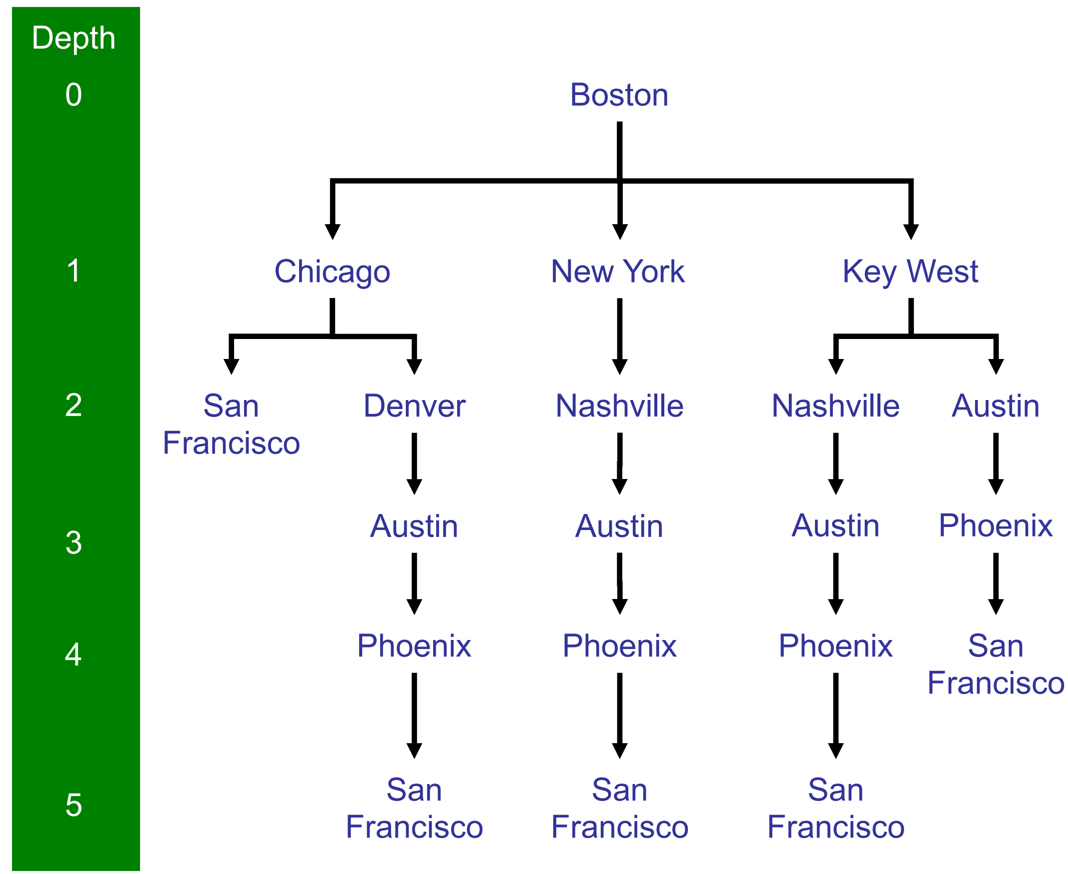

Branching Factor b=3

Heuristic Search

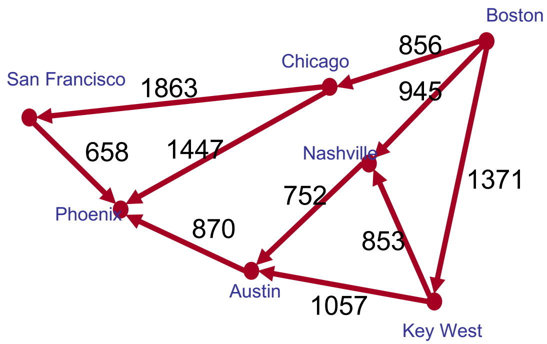

	Distance to Phoenix
Boston	2299
Chicago	1447
Nashville	1444
Key West	1927
Austin	870
San Francisco	658

Greedy Search

Heuristic function gives an estimate of the distance to the goal

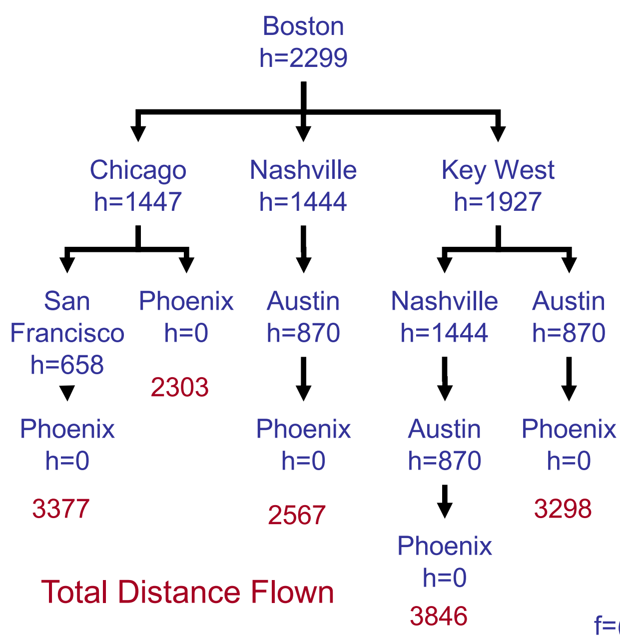

A* Search

Minimize the total path cost (f) =

actual path so far (g) +

estimate of future path to goal (h)

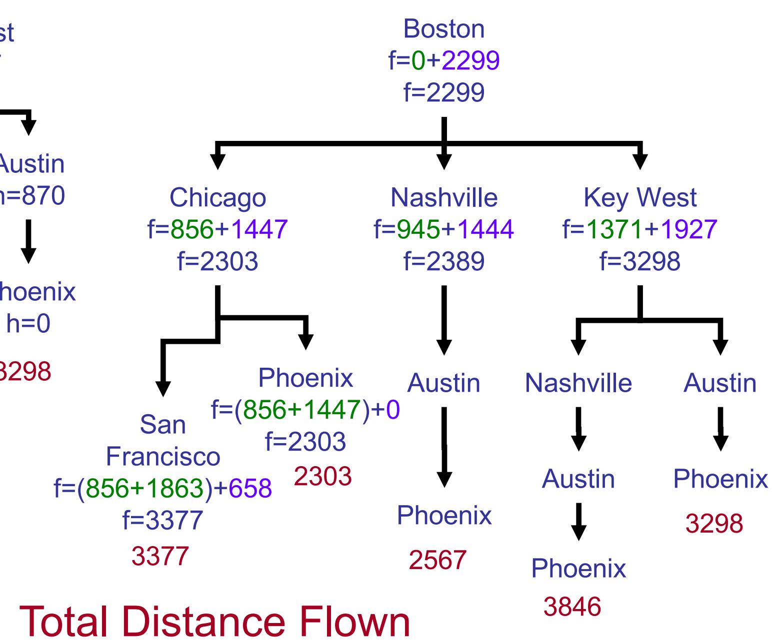

Search and Game Playing

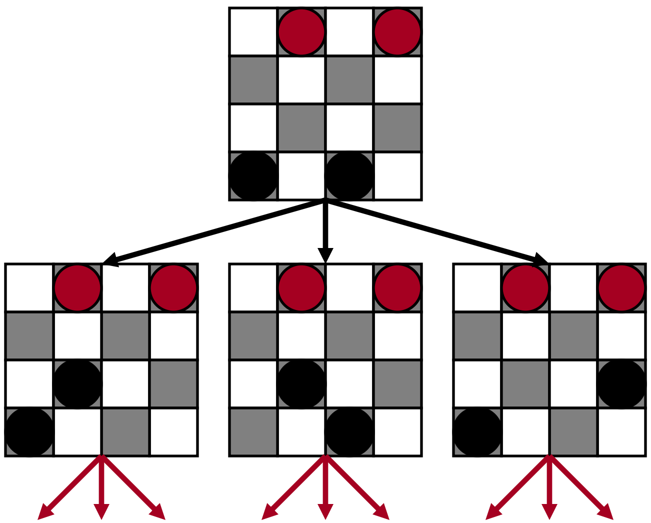

Kasparov vs. Deep Blue

Minimax Search with Alpha-Beta Pruning

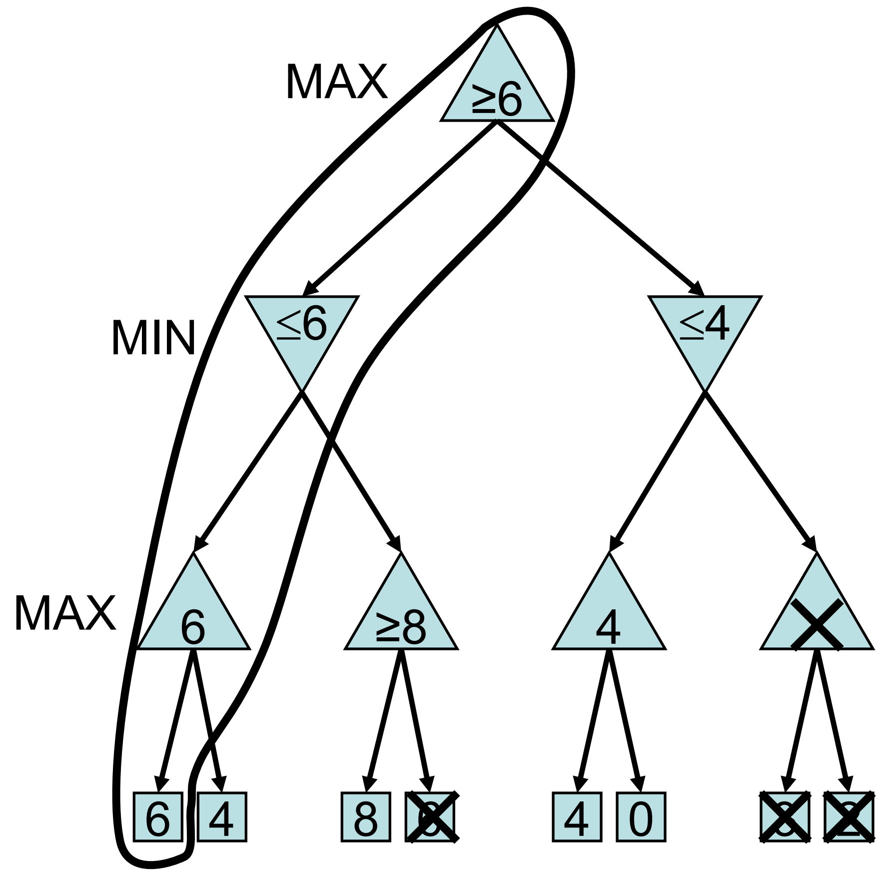

Knowledge Representation

Propositional Logic Syntax

$$Sentence \rightarrow AtomicSentence \mid ComplexSentence$$

$$AtomicSentence \rightarrow True \mid False \mid P \mid Q \mid \ldots$$

$$ComplexSentence \rightarrow (Sentence) \mid$$

$$Sentence\ Connective\ Sentence \mid$$

$$\neg Sentence$$

$$Connective \rightarrow \wedge \mid \vee \mid \Rightarrow \mid \Leftrightarrow$$

Inference Rules

$$\frac{\alpha \Rightarrow \beta\ ,\ \alpha}{\beta} \qquad \frac{\neg\neg\alpha}{\alpha} \qquad \frac{\alpha_1, \alpha_2, \alpha_3, \ldots \alpha_n}{\alpha_1 \wedge \alpha_2 \wedge \ldots \wedge \alpha_n}$$

$$\frac{\alpha_1 \wedge \alpha_2 \wedge \alpha_3 \wedge \ldots \wedge \alpha_n}{\alpha_i} \qquad \frac{\alpha \vee \beta\ ,\ \neg\beta}{\alpha} \qquad \frac{\alpha_i}{\alpha_1 \vee \alpha_2 \vee \ldots \vee \alpha_n}$$

$$\frac{\neg\alpha \Rightarrow \beta\ ,\ \beta \Rightarrow \gamma}{\neg\alpha \Rightarrow \gamma} \qquad \frac{\alpha \vee \beta\ ,\ \neg\beta \vee \gamma}{\alpha \vee \gamma}$$

Wumpus World

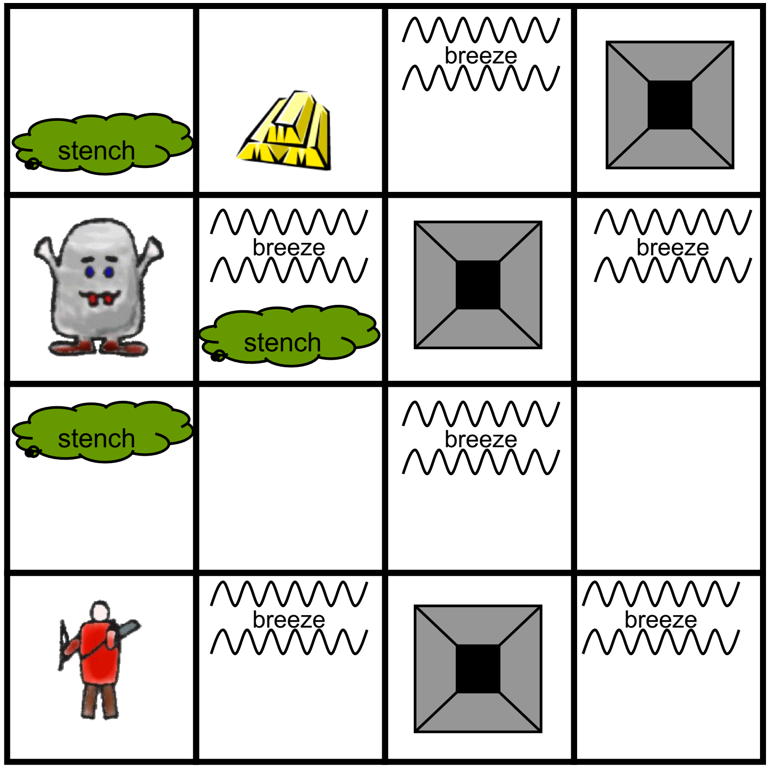

First-Order Logic

- Existential and Universal Quantifiers

$$
\begin{aligned}
Sentence &\rightarrow AtomicSentence \\
&\quad \mid Sentence\ Connective\ Sentence \\
&\quad \mid Quantifier\ Variable,\ldots Sentence \\
&\quad \mid \neg Sentence \\
&\quad \mid (Sentence) \\
AtomicSentence &\rightarrow Predicate(Term,\ldots) \\
&\quad \mid Term = Term \\
Term &\rightarrow Function(Term,\ldots) \\
&\quad \mid Constant \\
&\quad \mid Variable \\
Connective &\rightarrow \Rightarrow \mid \wedge \mid \vee \mid \Leftrightarrow \\
Quantifier &\rightarrow \forall \mid \exists \\
Variable &\rightarrow a \mid b \mid c \mid \ldots \\
Function &\rightarrow Mother \mid LeftLegOf \mid \ldots \\
Predicate &\rightarrow Before \mid HasColor \mid Raining \mid \ldots \\
Constant &\rightarrow A \mid X_1 \mid John \mid \ldots
\end{aligned}
$$

- Situation Calculus

  At(Agent,[1,1], S_0) $\wedge$
  At(Agent,[1 2], S_1)

- Changes from one situation to the next

  Result(Forward, S_0)$\Rightarrow$$S_1$

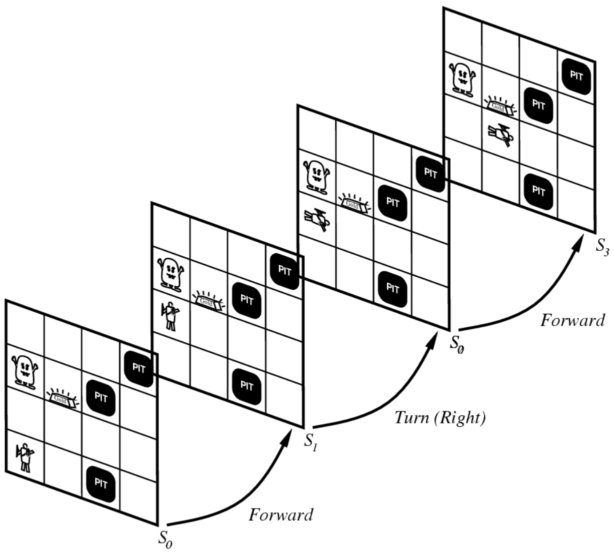

Building a Knowledge Base

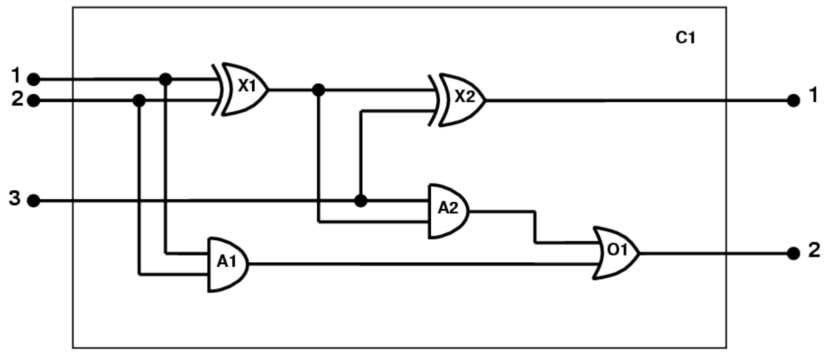

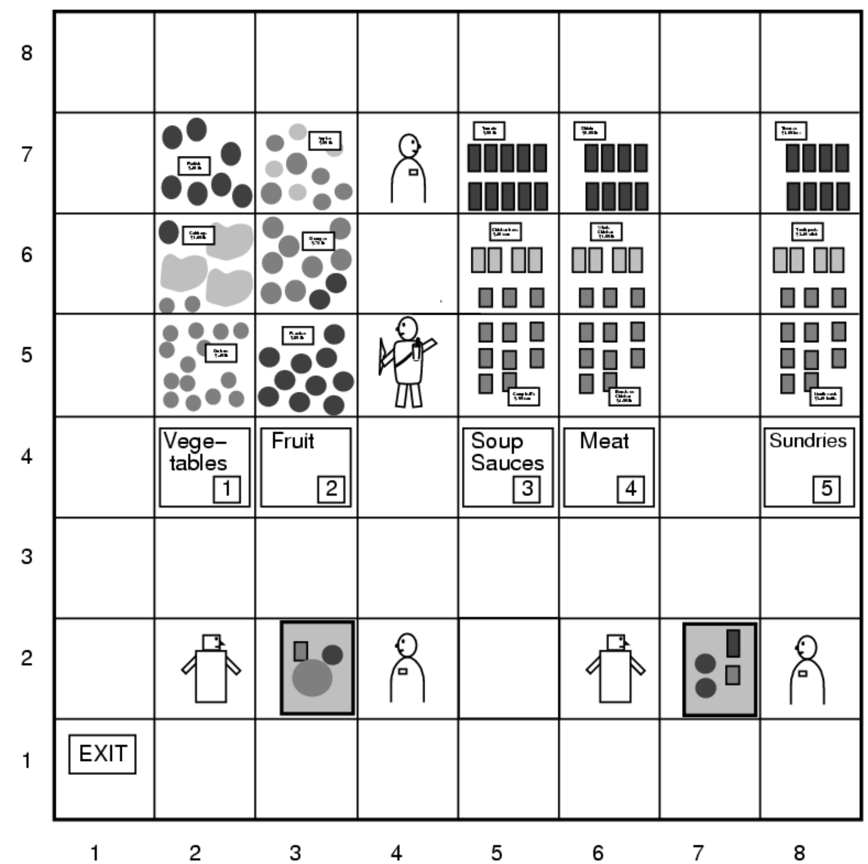

- Decide what to talk about
- Decide on a vocabulary of predicates, functions, and constants
  - Ontology
- Encode general knowledge within the domain
  - Limiting errors
- Encode a description of the specific problem
- Pose queries and get answers

Inference

Resolution

- American(x) ∧ Alcohol(y) ∧ Minor(z) ∧ Sells(x,y,z) ⇒ Criminal(x)
- Minor(Jimmy)
- Owns(Jimmy,B1)
- Beer(B1)
- Owns(Jimmy,x) ∧ Beer(x) ⇒ Sells(Nathan,x,Jimmy)
- American(Nathan)
- Beer(x)⇒Alcohol(x)
- Using 4, 7 and modus ponens
  Alcohol(B1)
- Using 5, 3, 4 and modus ponens
  Sells(Nathan,B1,Jimmy)
- Using 1, 6, 8, 2, 9 and modus ponens
  Criminal(Nathan)

Proof by Refutation

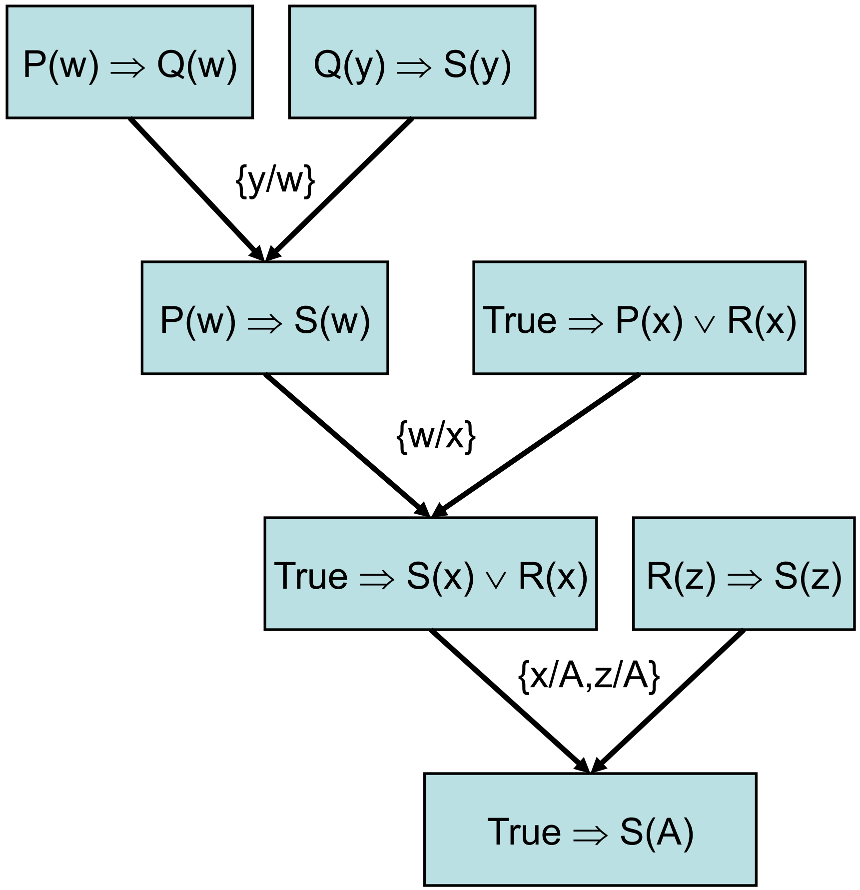

Expert Systems

SAINT

$$\int \frac{x^4}{(1-x^2)^{\frac{5}{2}}} dx$$

Try $y = \arcsin x$, yielding:

$$\int \frac{\sin^4 y}{\cos^4 y} dy$$

DENDRAL

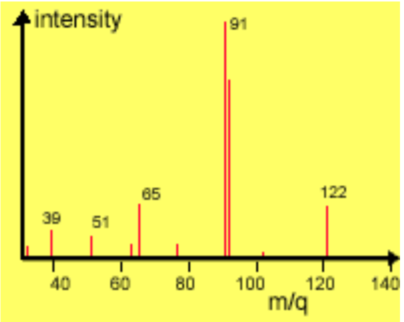

Mass spectrogram for $C_8H_{10}O$

XCON (R1)

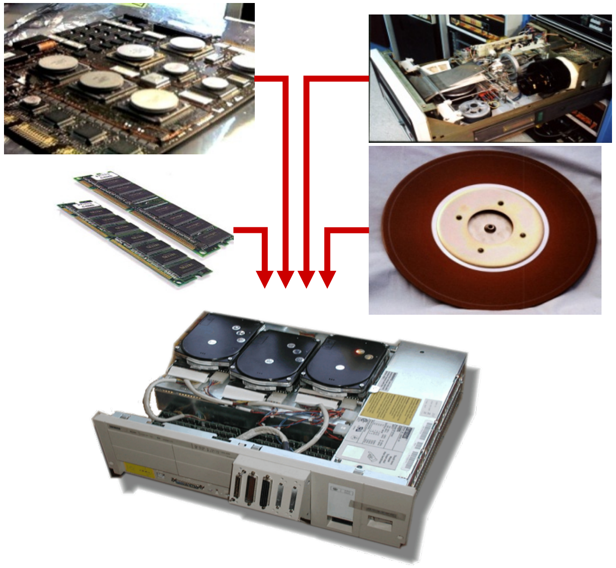

CYC

CYCORP

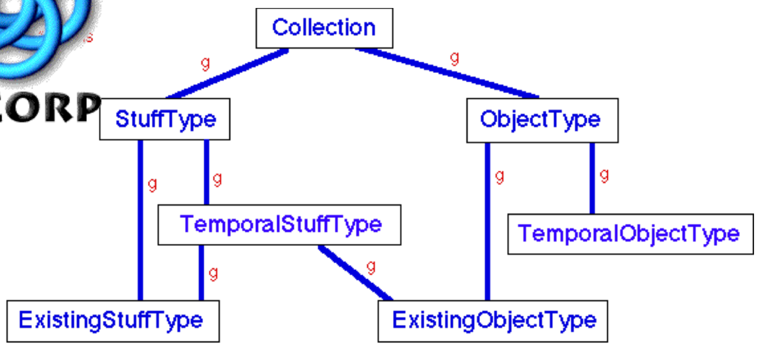

Planning

Representing World State and Change in a Logical Language

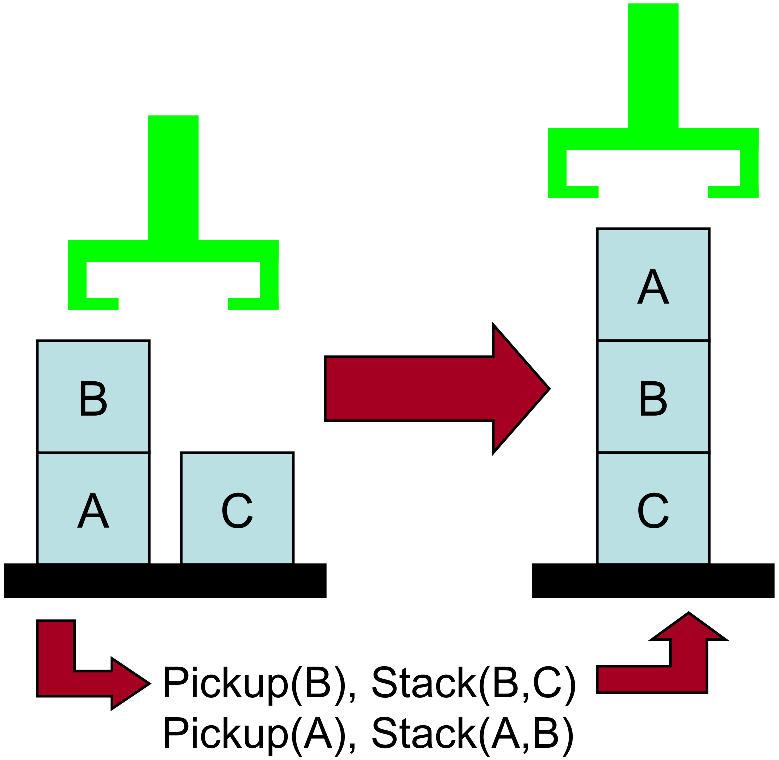

Partial-Order Planning

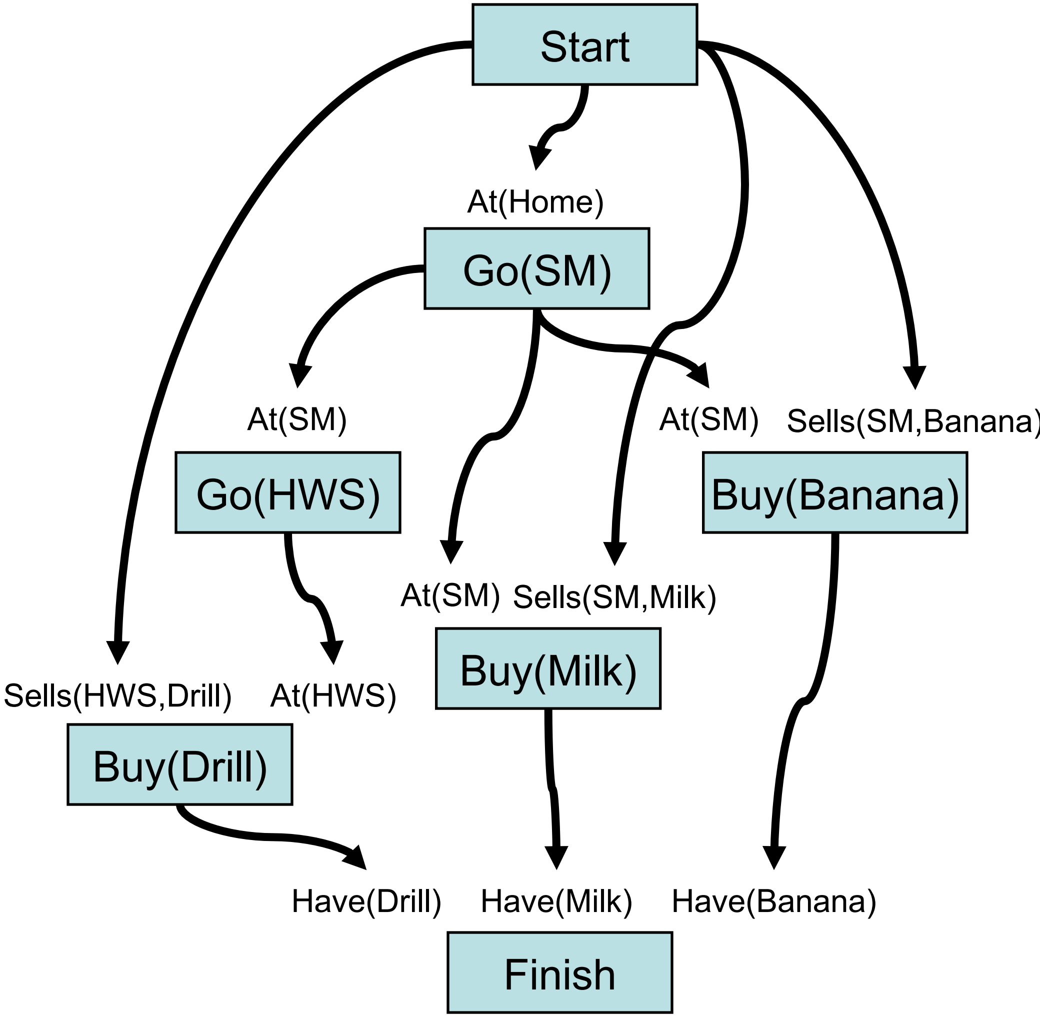

Planning in the Real World: Robot path planning

Configuration Spaces

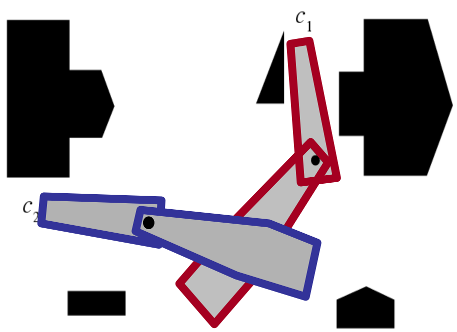

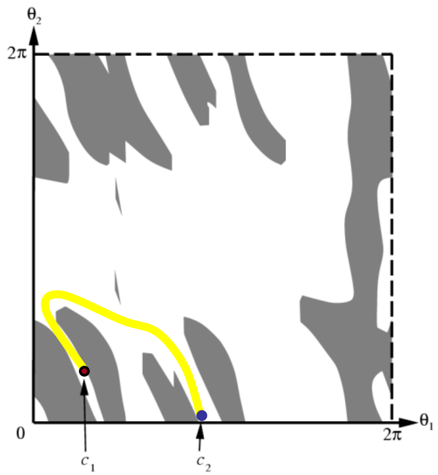

Probabilistic Roadmap

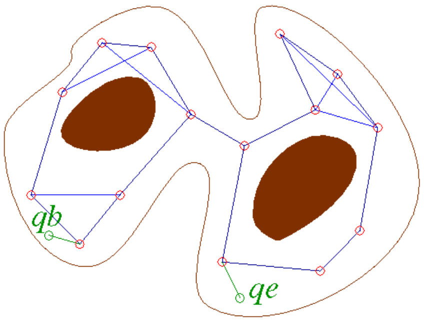

Cell Decomposition

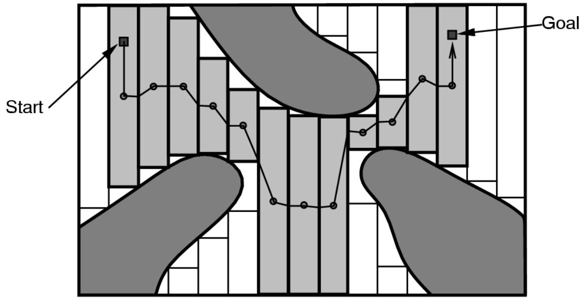

Visibility Graphs

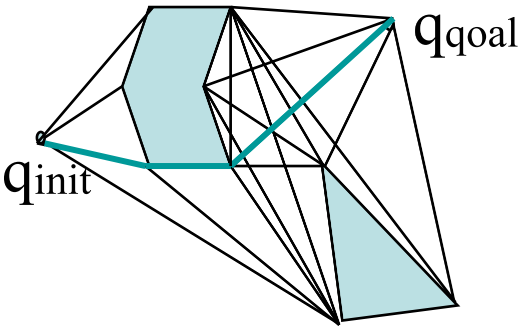

Potential Fields

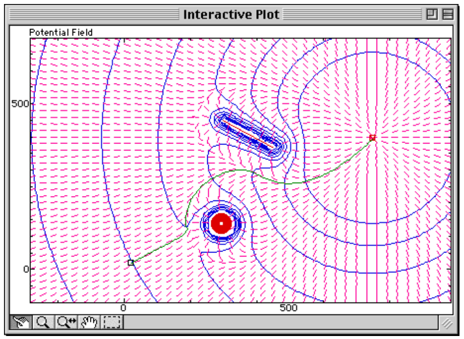

Planning in Real-World Systems

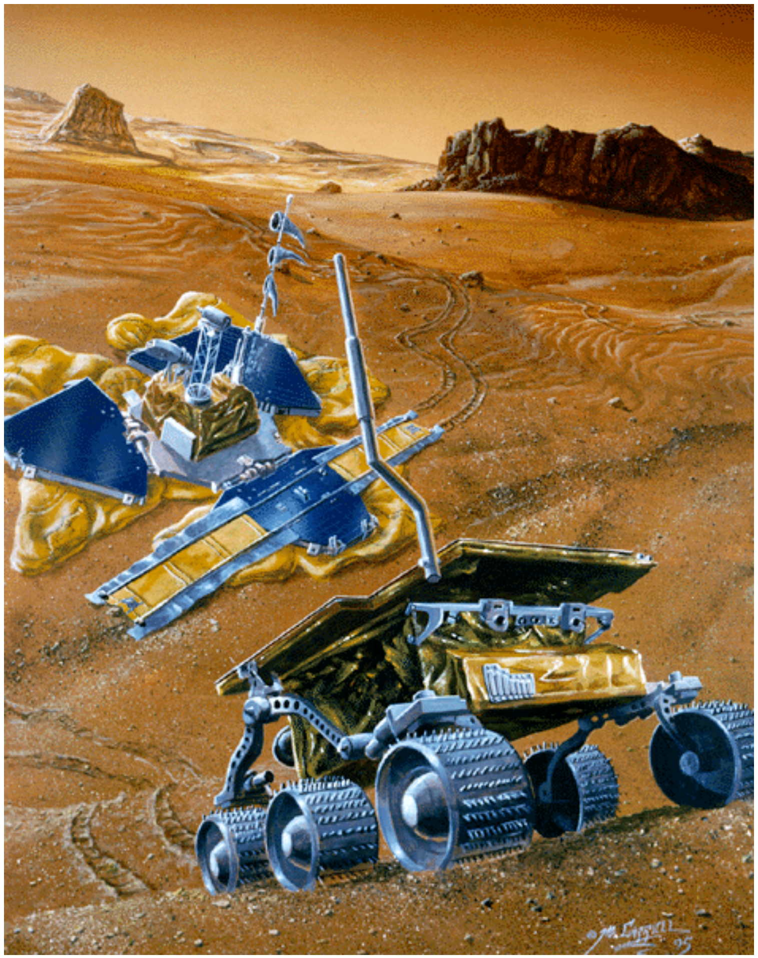

Conditional Planning

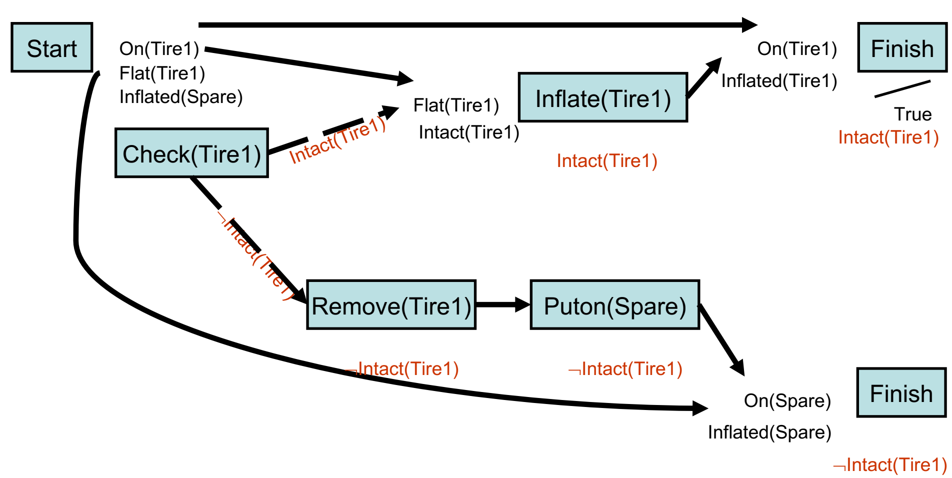

Execution Monitoring

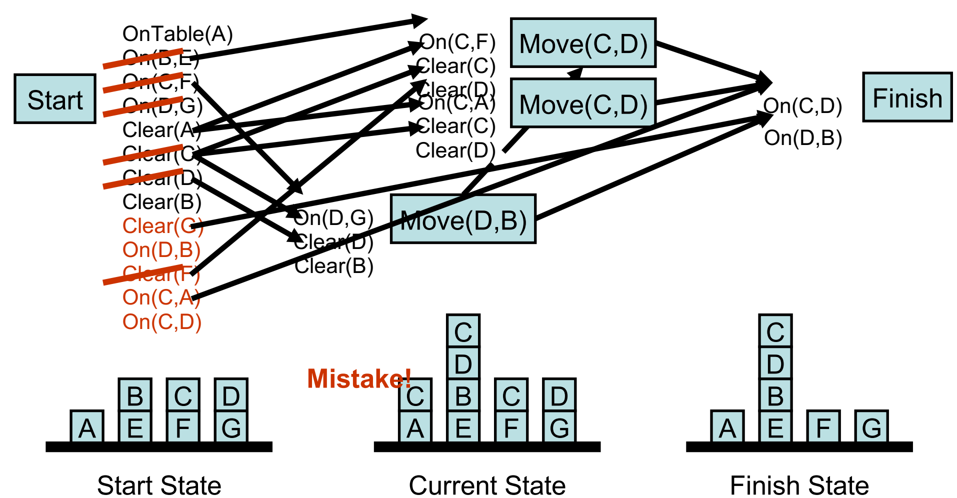

Dealing with Uncertainty

Belief Networks

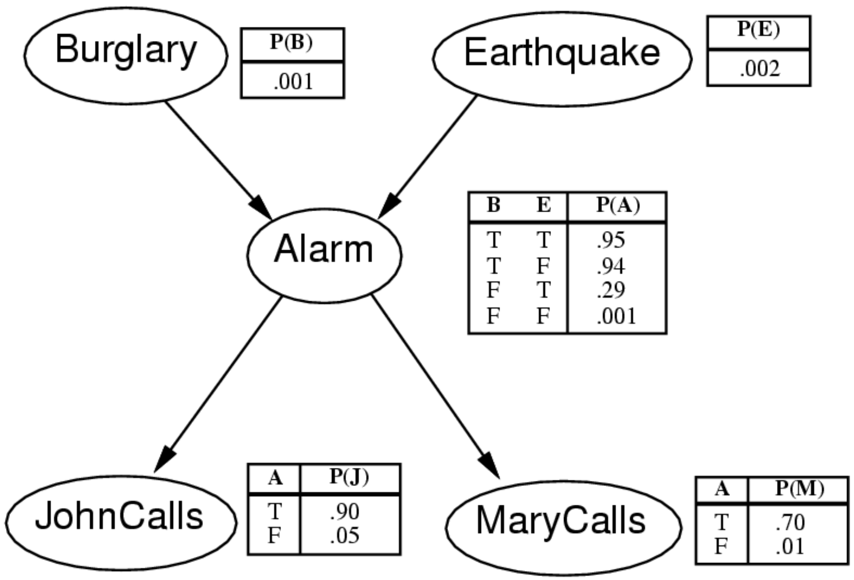

A conditional probability table gives the likelihood of a particular combination of values

Incremental Construction

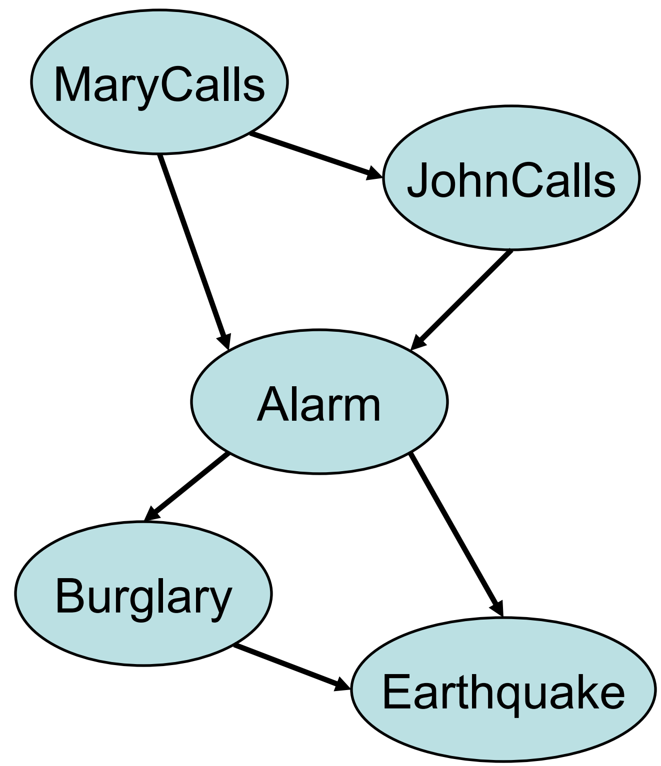

Learning from Observations

Learning Optimal Decision Trees

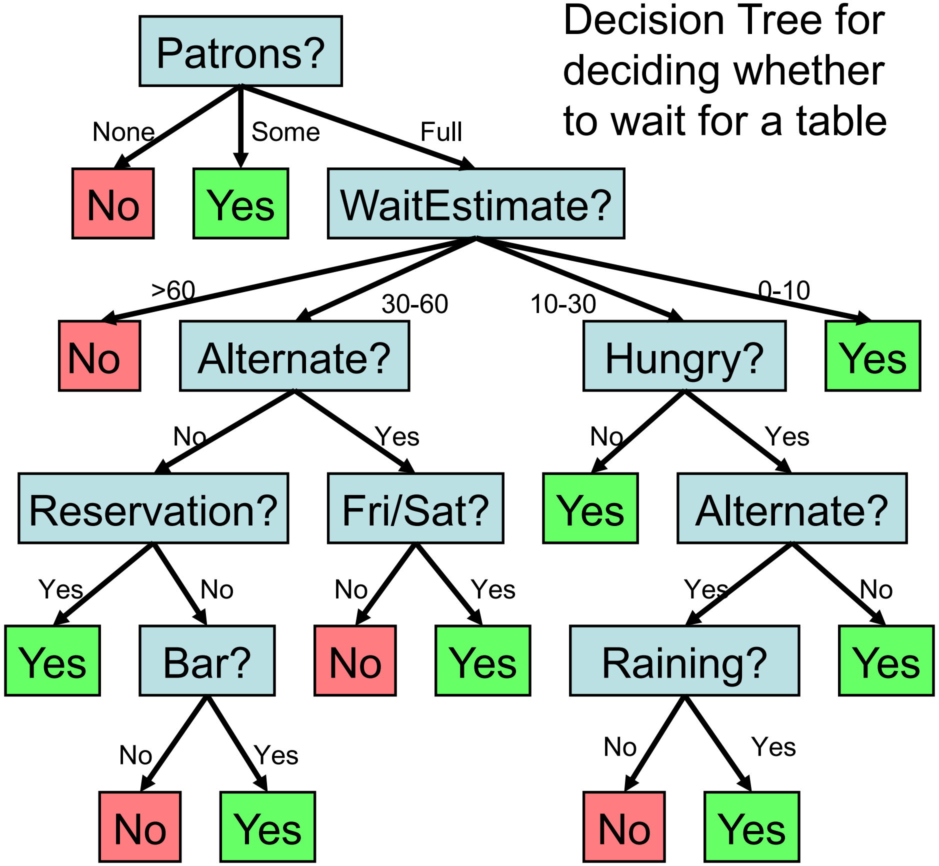

Decision Tree for deciding whether to wait for a table

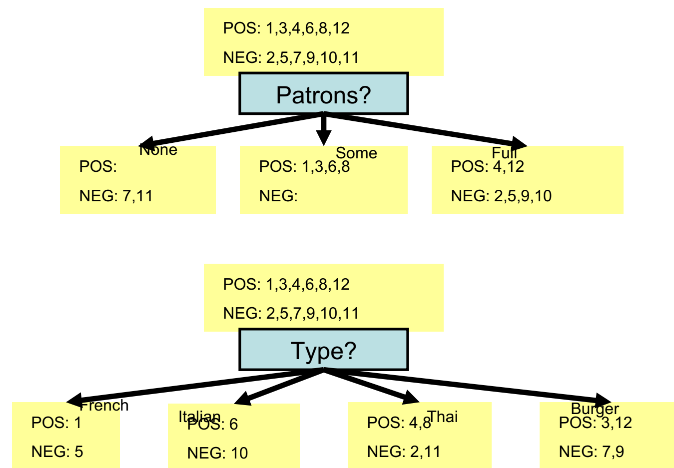

$$Remainder(A) = \sum_{i=1}^{v} \frac{p_i+n_i}{p+n} I\left(\frac{p_i}{p_i+n_i}, \frac{n_i}{p_i+n_i}\right)$$

$$Remainder(Patrons) = \tfrac{2}{12} I(0,1) + \tfrac{4}{12} I(1,0) + \tfrac{6}{12} I(\tfrac{2}{6}, \tfrac{4}{6})$$

$$Remainder(Patrons) \approx 0 + 0 + \tfrac{6}{12}\left(-\tfrac{2}{6}\log\tfrac{2}{6} - \tfrac{4}{6}\log\tfrac{4}{6}\right)$$

$$Remainder(Patrons) \approx 0.459 \text{ bits}$$

Supervised Learning Using Version Spaces

Most general boundaries (G)

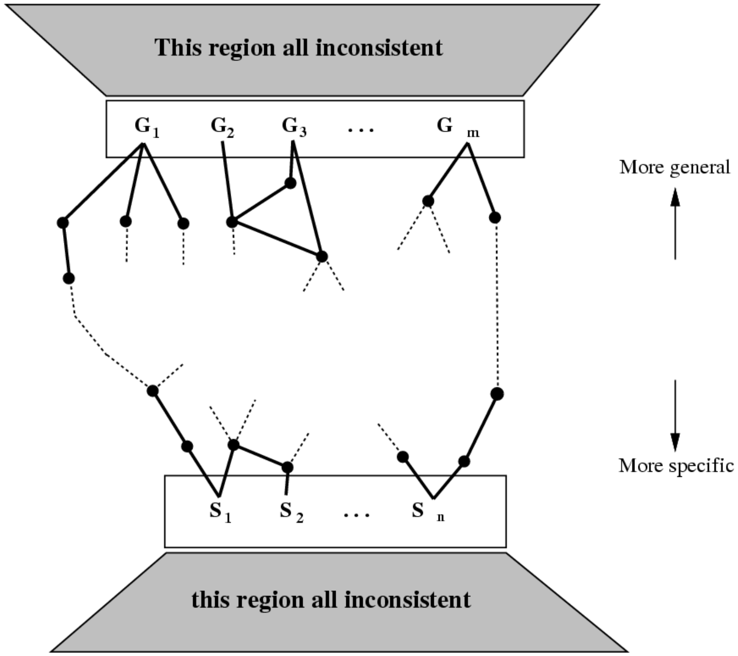

Most specific boundaries (S)

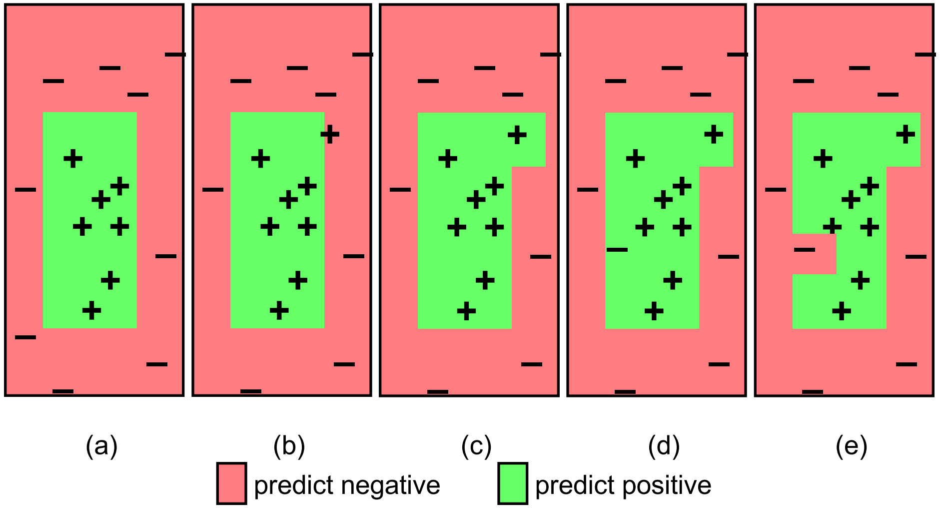

(a) consistent

(b) false negative

(c) generalization includes the false negative example

(d) false positive

(e) specialization removes the false positive example

Genetic Algorithms

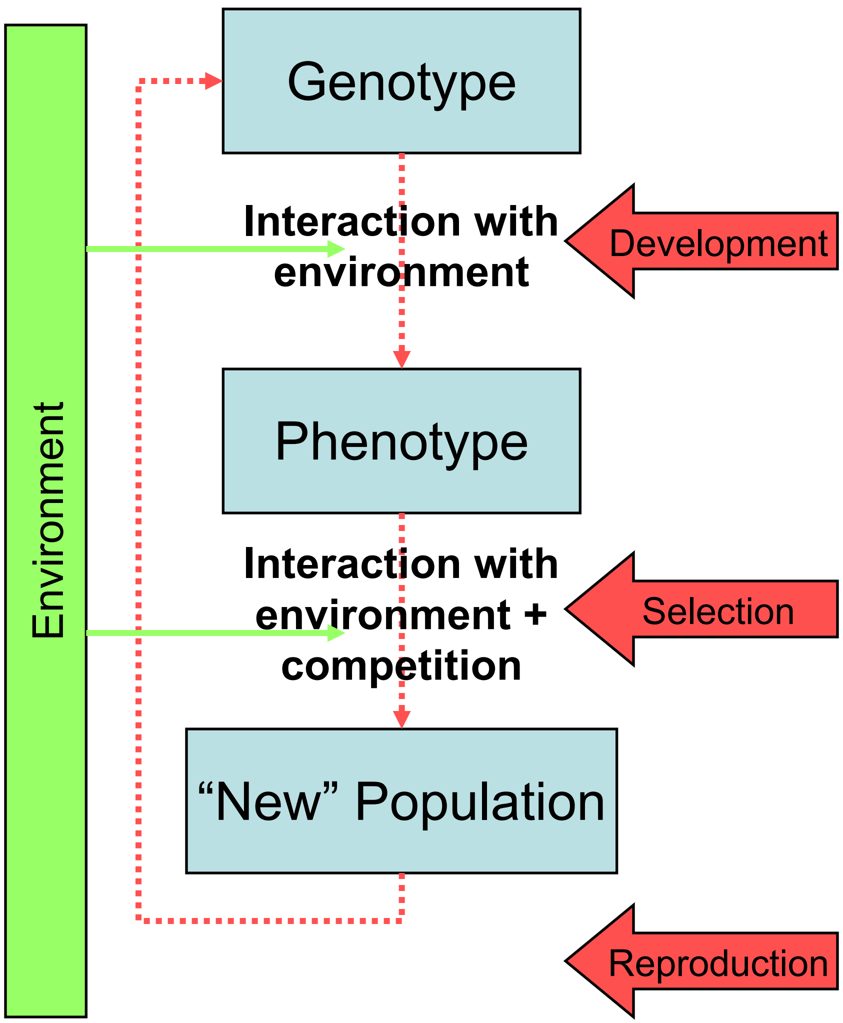

Evolving physical morphology and control: Karl Sims

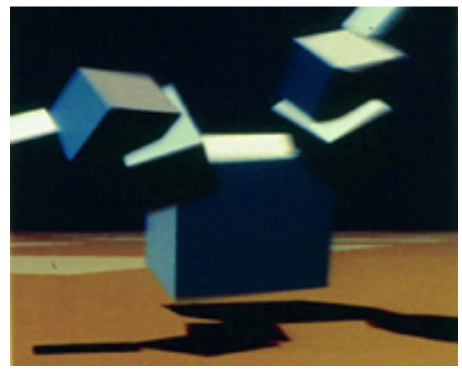

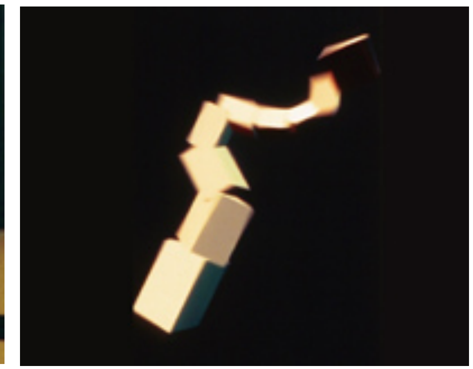

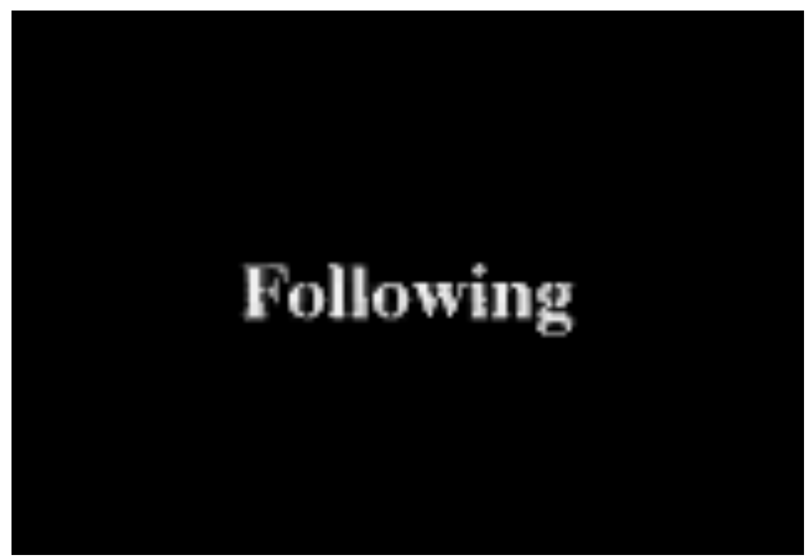

Learning Using Neural Nets

Perceptrons

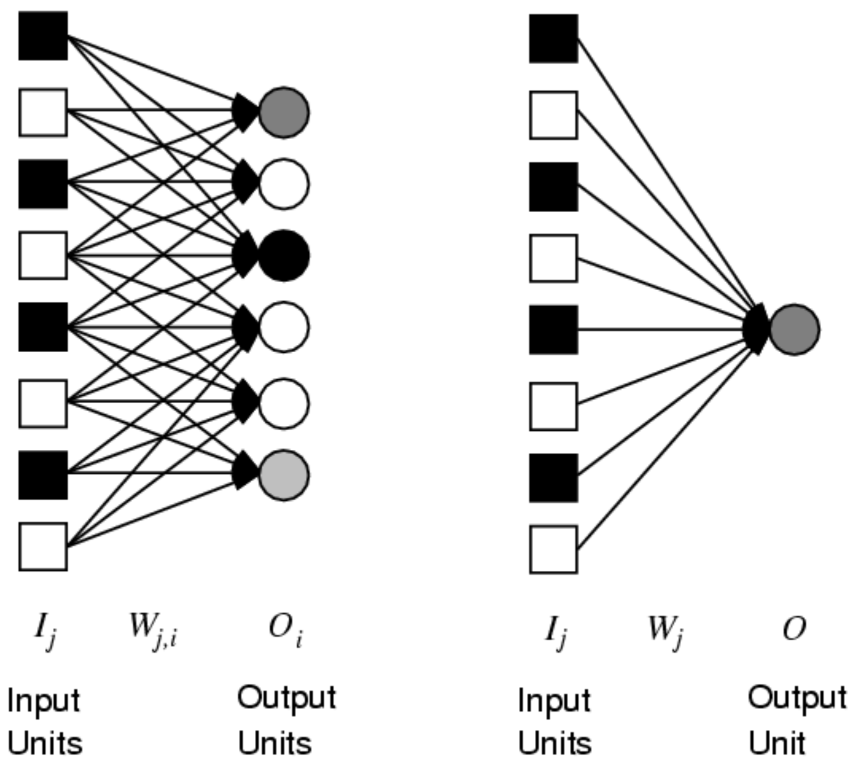

Perceptron Network **Single Perceptron**

Multi-Layer Networks

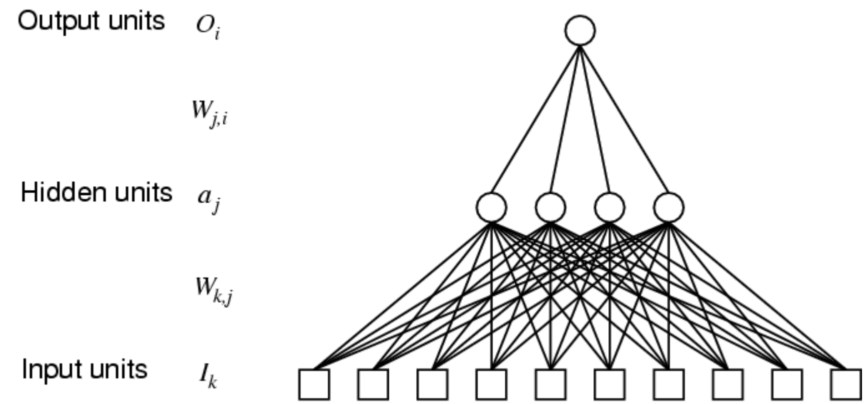

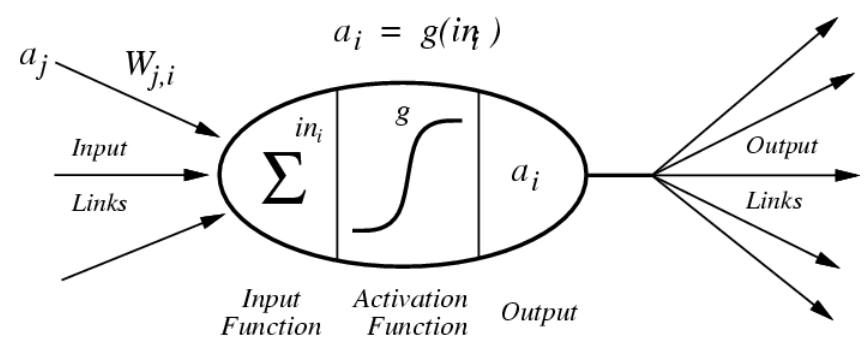

Backprop and Linear Seperability

Deep Learning

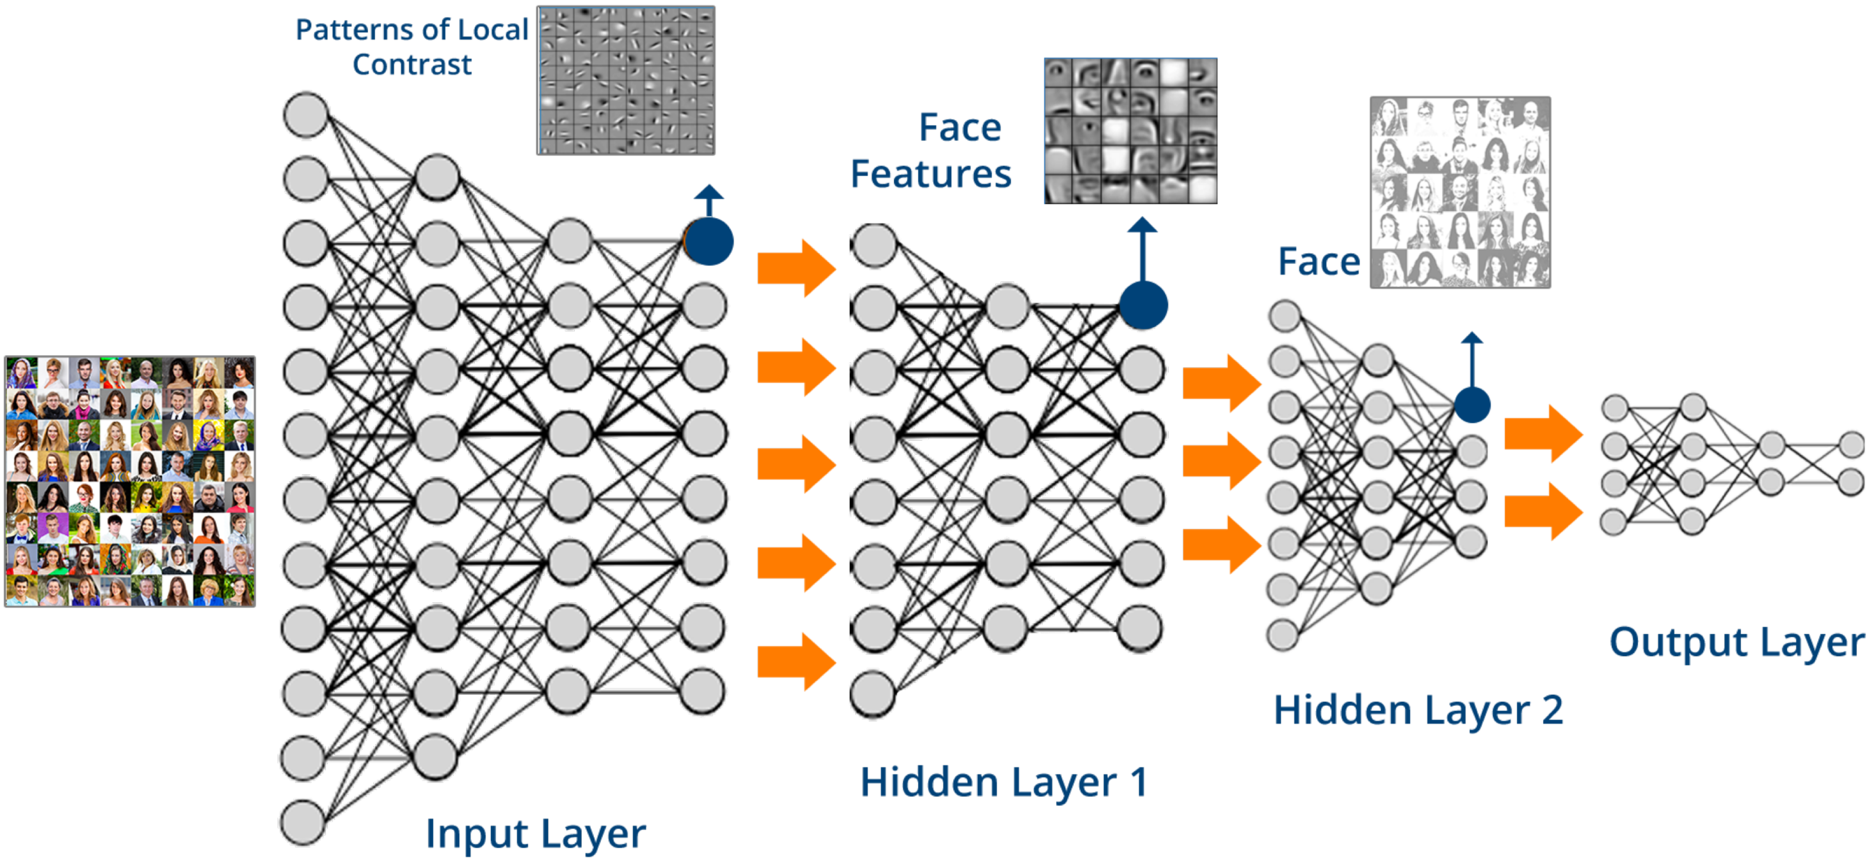

Reinforcement Learning

(Rewarded at the end of an action sequence)

Utility Learning (Temporal Difference)

- Learn a utility function that maps states to utilities and select an action by maximizing expected value
- Needs a model of the environment (needs to know the results of actions)
- Predictive

Action-Value Learning (Q-Learning)

- Learn an action-value function that gives the expected utility of taking a given action in a given state
- No need for an environment model
- Do not know where actions lead, so it cannot look ahead

$$Q(a,i) \leftarrow Q(a,i) + \alpha(R(i) + \max_{a'} Q(a',j) - Q(a,i))$$

3	−0.0380	0.0886	0.2152	+1
2	−0.1646		−0.4430	−1
1	−0.2911	−0.0380	−0.5443	−0.7722
	1	2	3	4

Estimated Utility Values

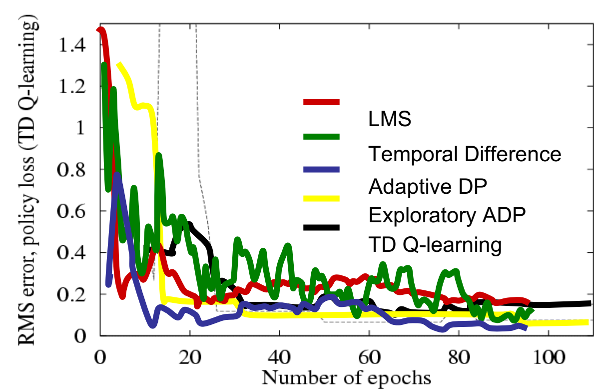

Communication: Grammars, Syntax, and Semantics

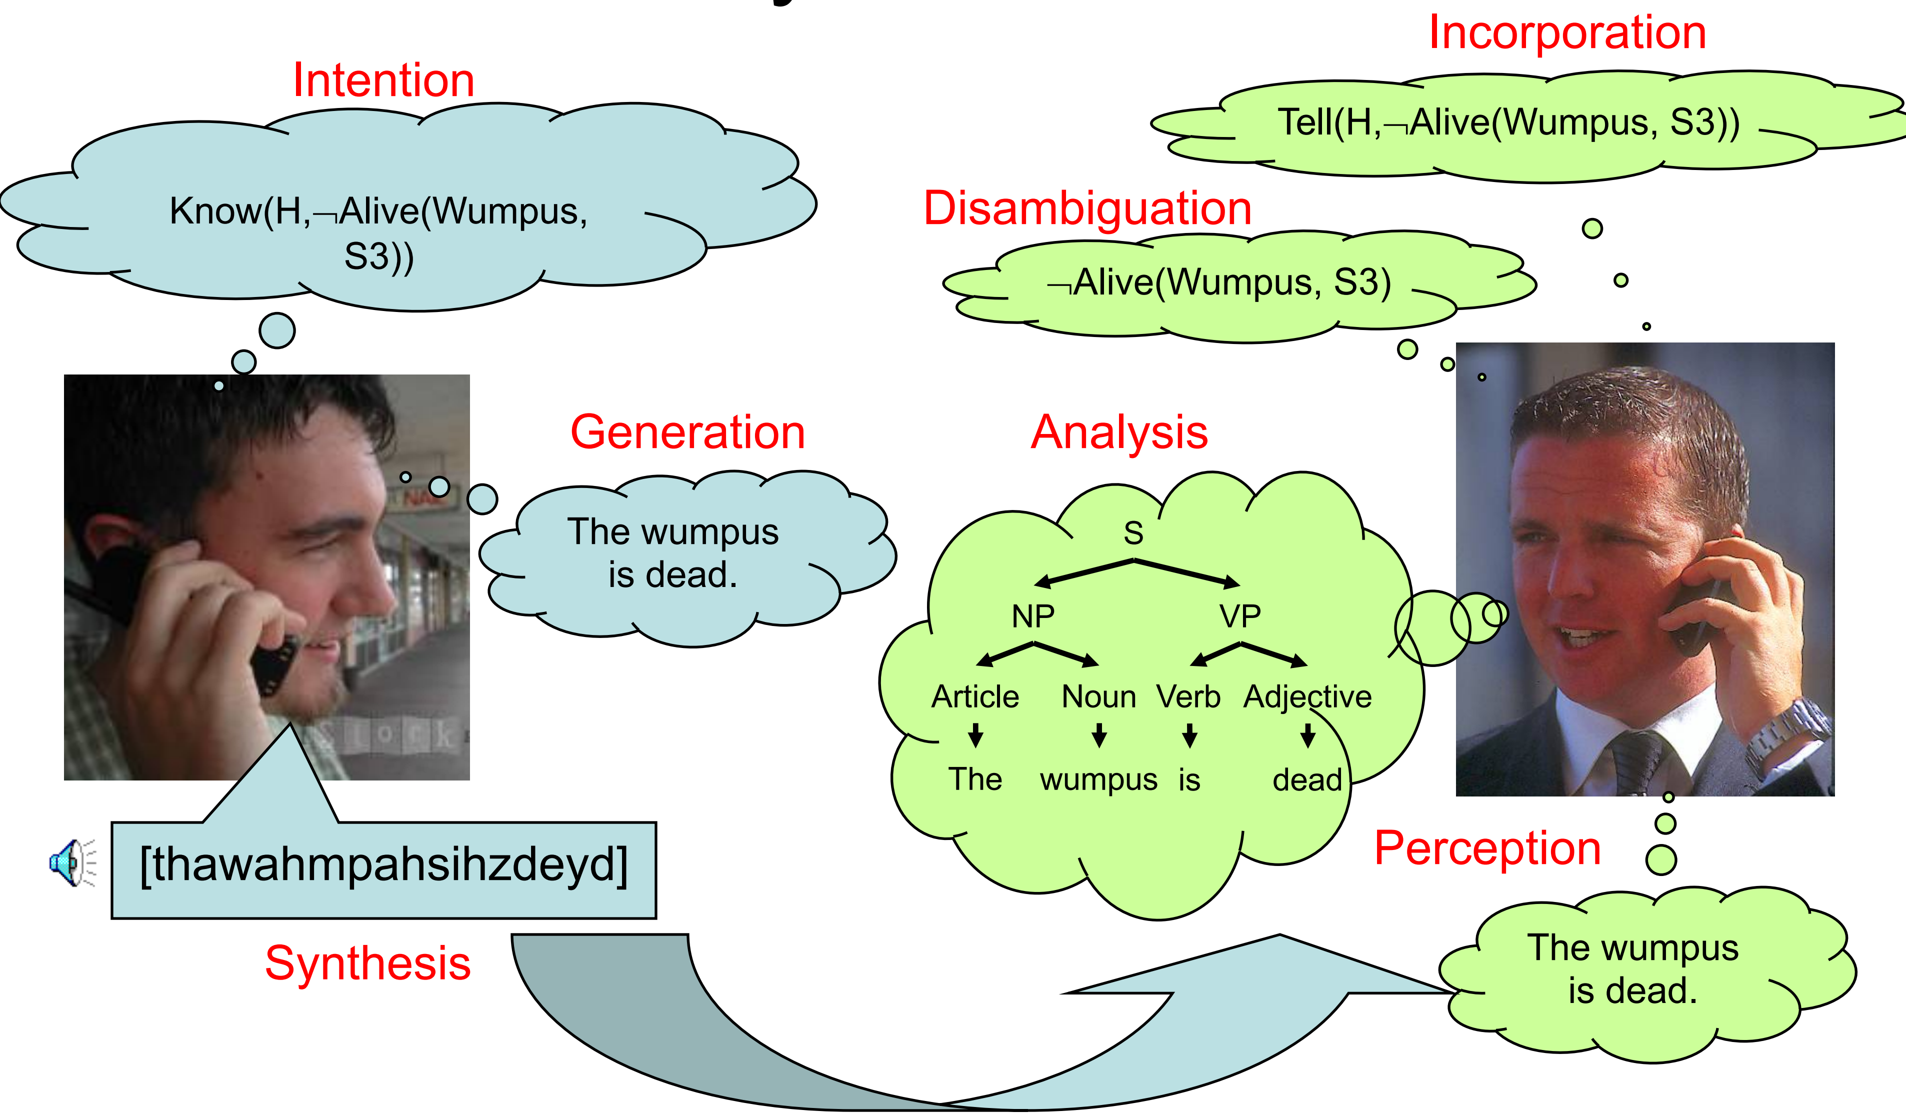

Communication

Speech Generation

Klatt Synthesizer

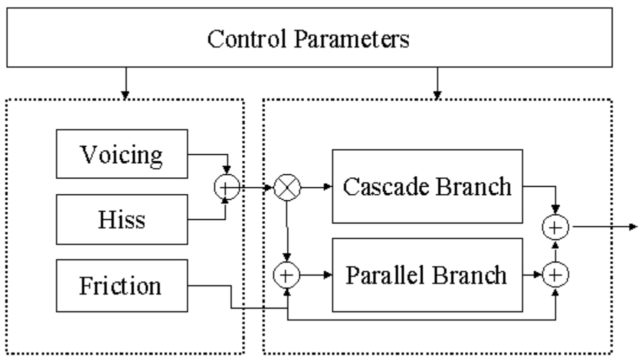

Source Model Filter Model

DECTalk Demo

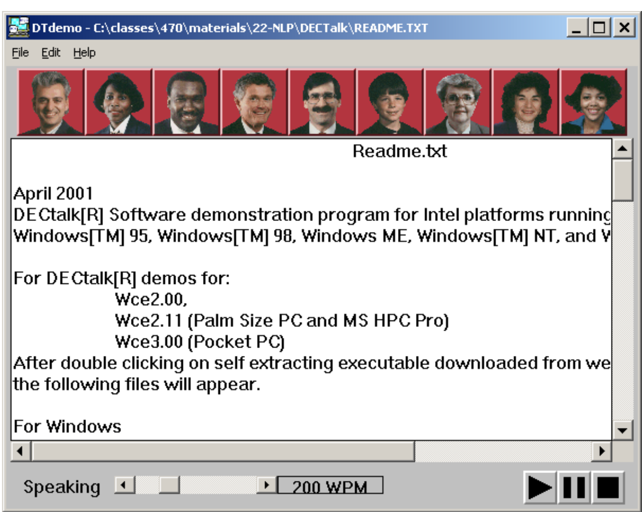

Speech Recognition

Hidden Markov Models

Phone HMM for [m]:

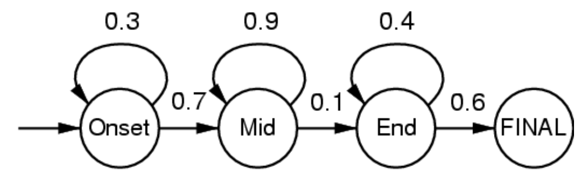

Output probabilities for the phone HMM:

Onset:	Mid:	End:
C1: 0.5	C3: 0.2	C4: 0.1
C2: 0.2	C4: 0.7	C6: 0.5
C3: 0.3	C5: 0.1	C7: 0.4

Demo of Dragon NaturallySpeaking

Perception

Mathematical Tools: Convolution

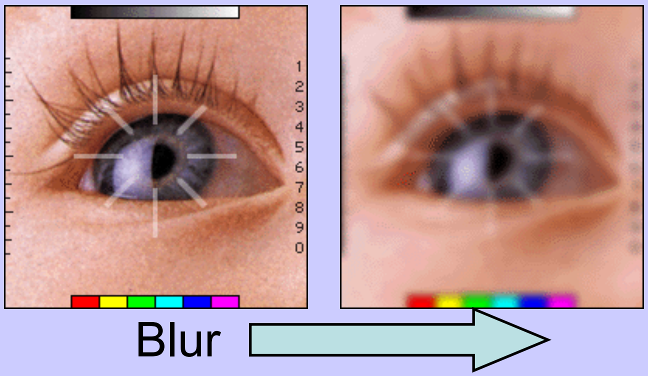

$$h(x) = \int_{-\infty}^{+\infty} f(u)g(x-u)du$$

$$h(x) = \sum_{u=-\infty}^{+\infty} f(u)g(x-u)$$

Applications: Edge Detection

Pre-attentive and Post-attentive

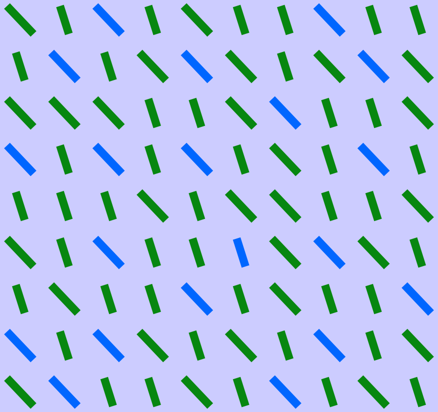

Q	A	S	D	F	G	E	R	O	P
U	K	J	E	R	T	H	C	F	M
A	Z	E	S	F	G	Q	W	R	T
F	G	H	U	Y	Y	B	X	L	W
V	N	R	H	J	B	D	K	W	L
R	T	G	F	M	X	V	P	O	S
Q	P	F	S	H	F	R	T	Y	U

Higher-Level Perception

Finding Similar Images

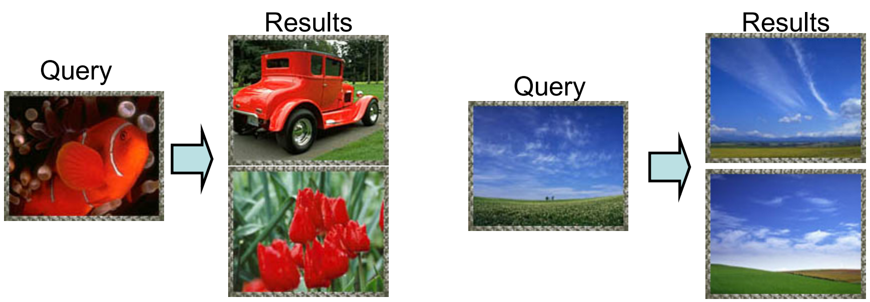

Motion Identification

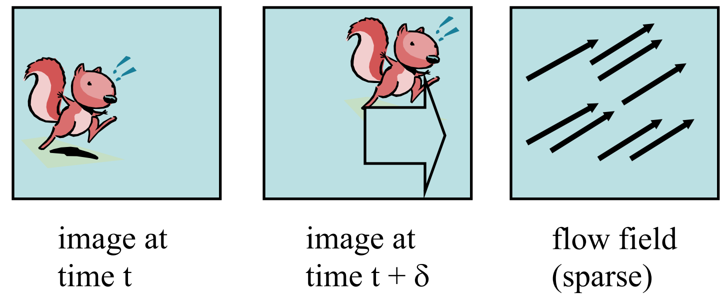

Region Segmentation

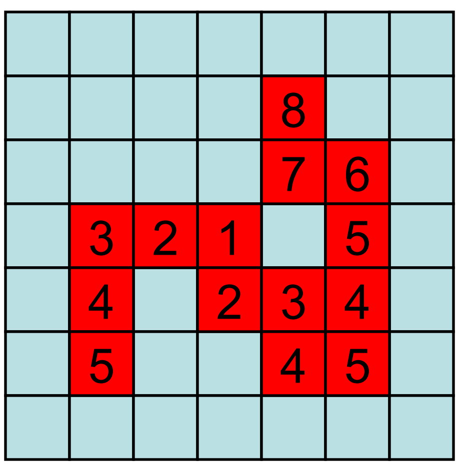

Object Detection and Recognition

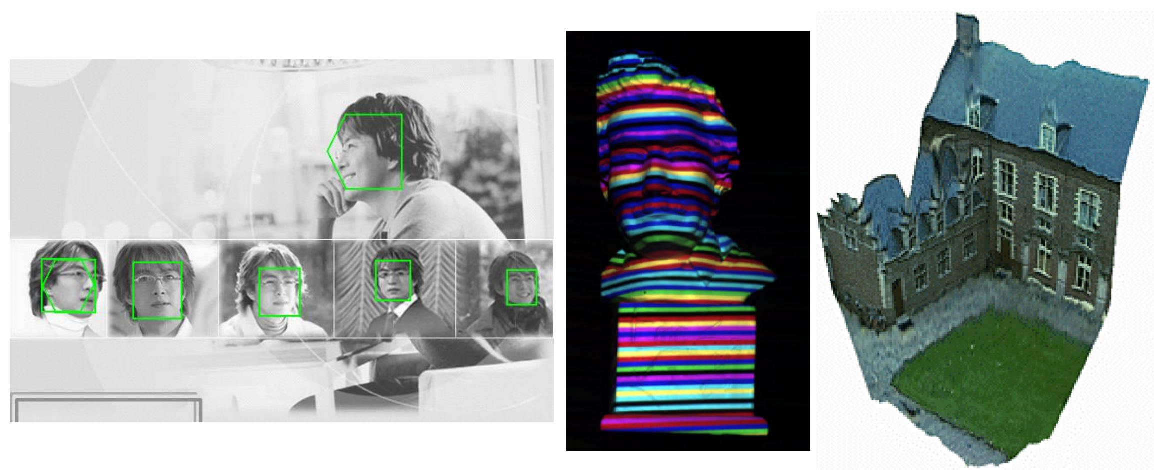

Robotics: Kinematics

Basic Joint Types

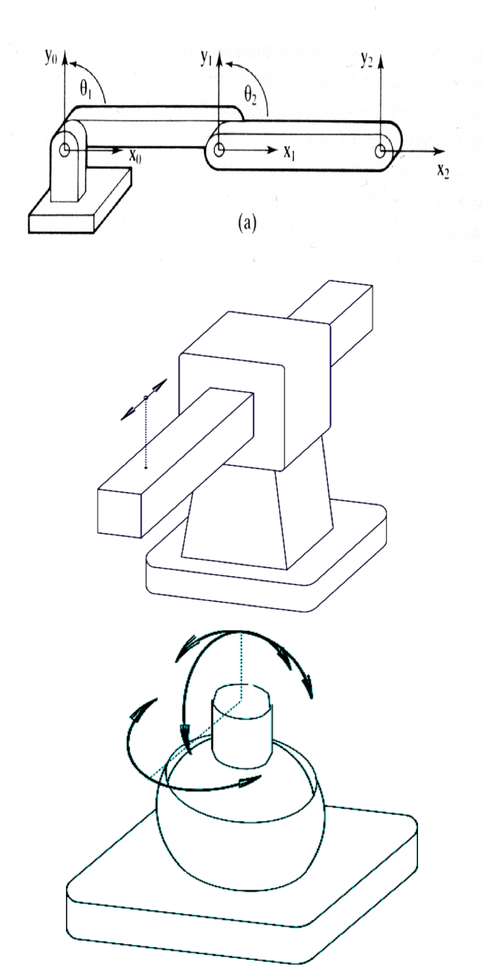

Forward Kinematics

(from joints to positions)

Given Y1, Y2 find x,y

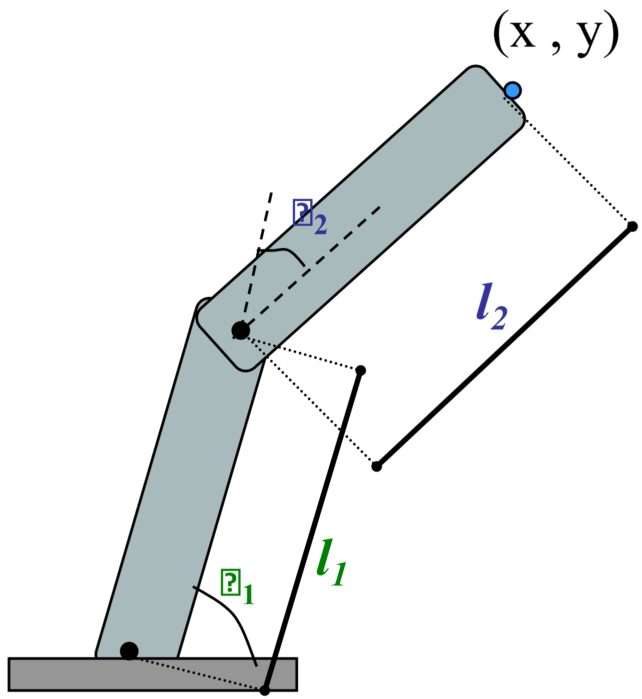

Inverse Kinematics

(from positions to joints)

Given x,y find Y1, Y2

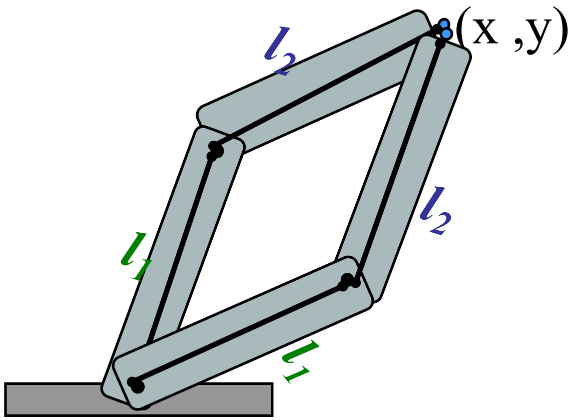

Problems with ambiguous solutions (or no solutions)

Robot Control Architectures

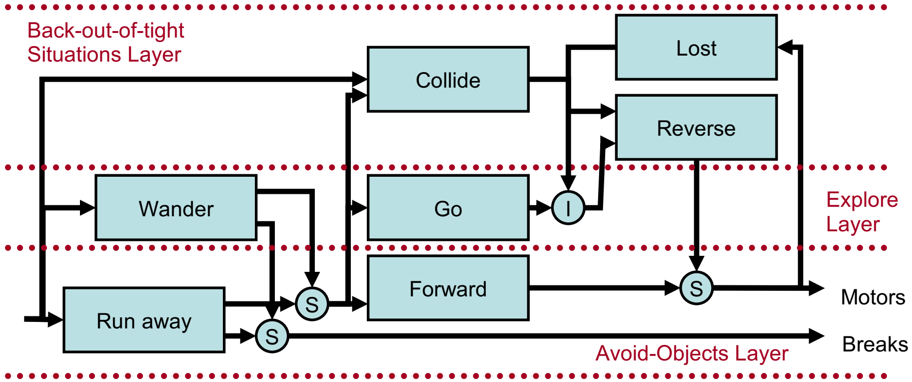

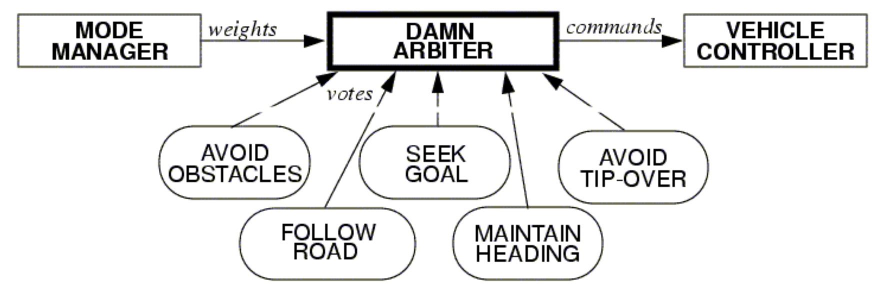

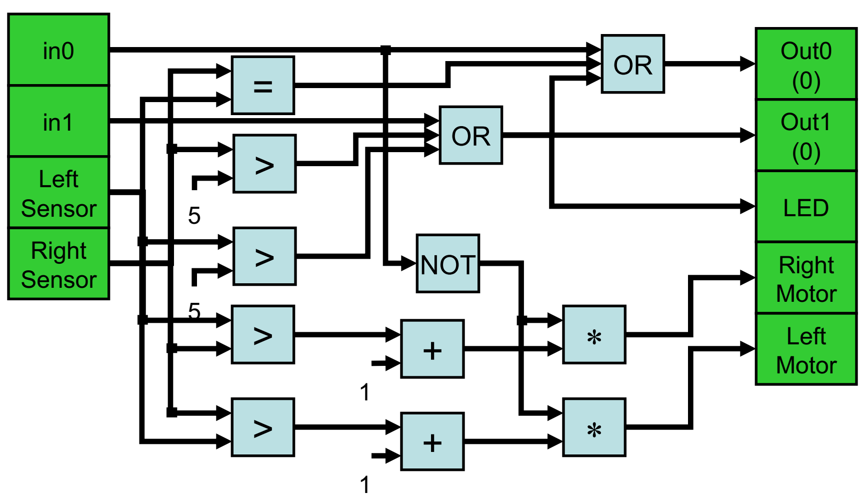

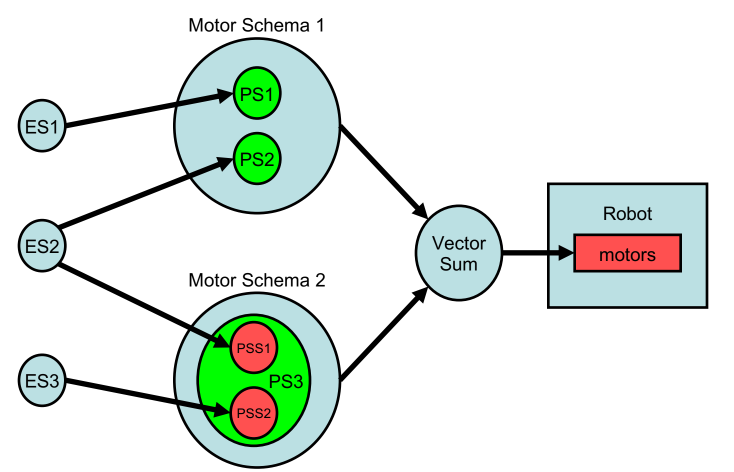

Humanoid Robots

Service Robots

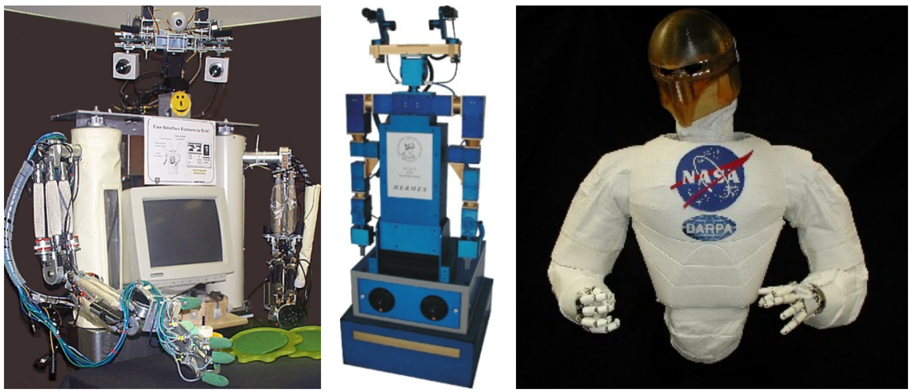

Adult-sized Research Robots

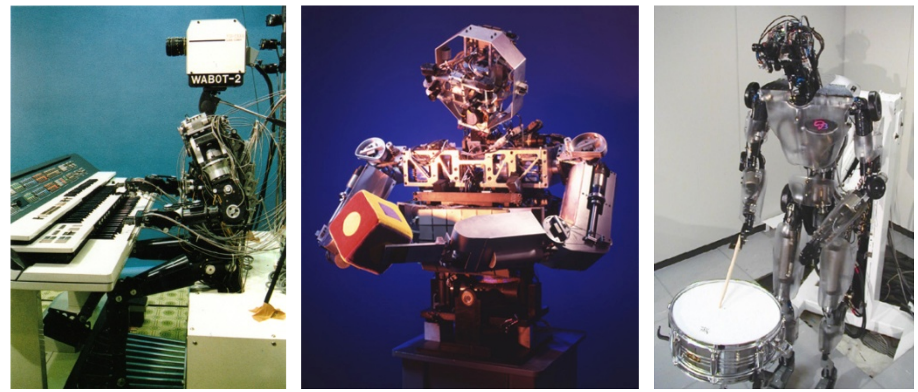

Commercial Robots

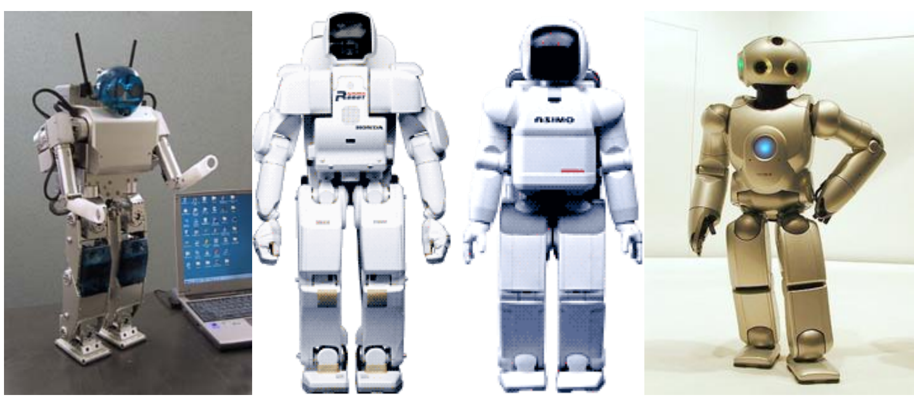

Child-sized Research Robots

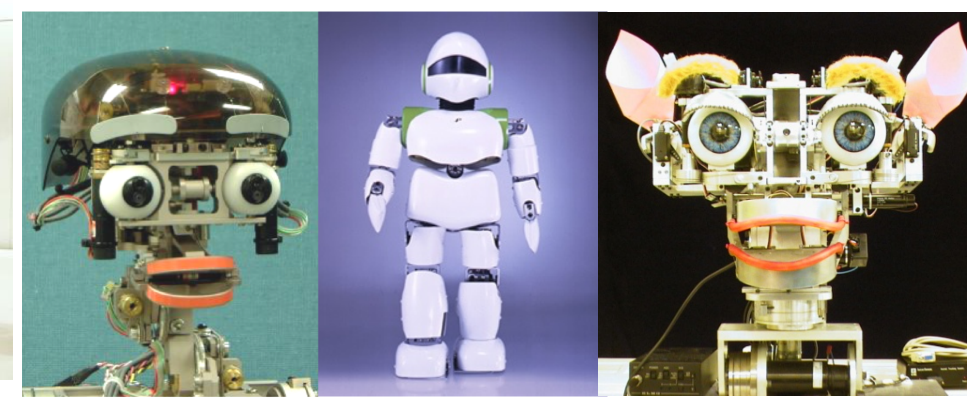

Social Robotics and HRI

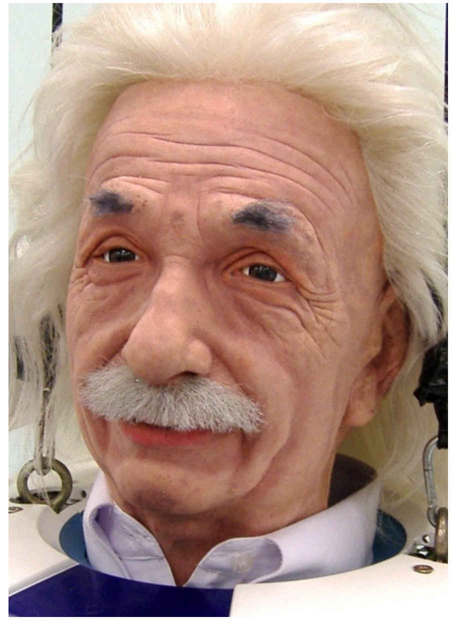

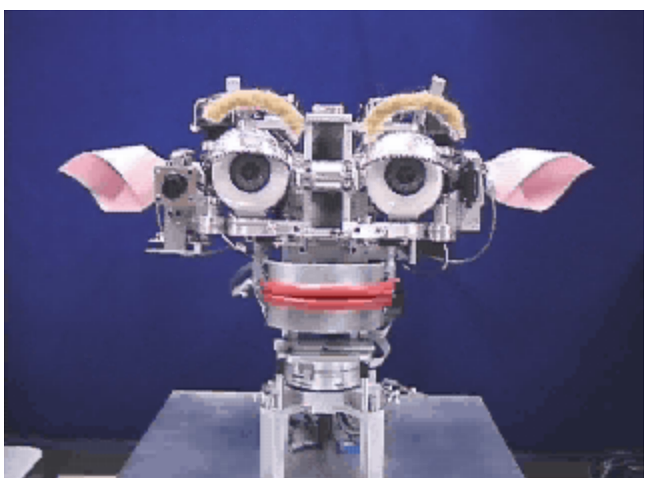

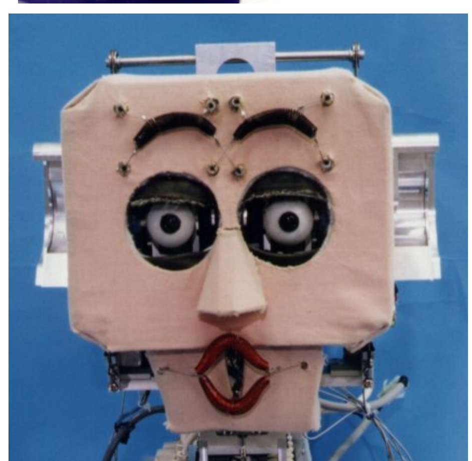

The Future of AI

Course Information

- Official prerequisites:
  - After CPSC 201 and 202 (or by permission of the instructor)
- Description:
  - Introduction to artificial intelligence research, focusing on reasoning and perception. Topics include knowledge representation, predicate calculus, temporal reasoning, vision, robotics, planning, and learning.
- Skills
  - Quantitative Reasoning

Grading

- Grading will be determined as follows:
  - Final Exam : 30%
  - Midterm Exam : 20%
  - Problem Sets : 50%
- These weights are subject to minor variations.
- Each problem on the problem sets and exams will be worth a specified number of points, which will be shown with the problem.

Syllabus

Date	Lecture Topic
01/14/19	Course Overview
01/16/19	Intelligent Agents
01/18/19	Python Intro
01/21/19	*No class - MLK*
01/23/19	Basic Search Algorithms
01/25/19	Informed Search Algorithms
01/28/19	Adversarial Search and Game Playing
01/30/19	TBA
02/01/19	*Guest– Dragomir Radev – NLP*
02/04/19	Constraint satisfaction problems
02/06/19	Propositional Logic
02/08/19	First Order Logic
02/11/19	Building a Knowledge Base
02/13/19	Inference
02/15/19	Planning
02/18/19	Motion Planning
02/20/19	Planning in the Real World
02/22/19	Reasoning Under Uncertainty
02/25/19	Learning from Observations
02/27/19	*Guest - Marynel Vasquez - robot navigation*
03/01/19	Supervised Learning
03/04/19	Midterm Exam
03/06/19	Genetic Algorithms
03/08/19	*Flex day*
	Spring break

Date	Lecture Topic
03/25/19	Neural Networks
03/27/19	Deep Learning
03/29/19	Reinforcement Learning I (utility functions)
04/01/19	Reinforcement Learning II (action-value learning)
04/03/19	Natural Language Processing
04/05/19	Communication
04/08/19	Introduction to Machine Perception
04/10/19	Higher-level Perception
04/12/19	Vision and Robotics
04/15/19	Robotics: Kinematics, Sensors and Actuators
04/17/19	Robotics: Control Architectures
04/19/19	Humanoid Robots
04/22/19	Emergence
04/24/19	Current Topics in AI
04/26/19	The Future of AI

Assignments (draft list)

- HW 0: Introduction to the Course Environment
- HW 1: Search (Pacman)
- HW 2: Game Playing (Othello)
- HW 3: Logic and Representations
- HW 4: Planning (Blocks world)
- HW 5: Supervised Learning (Muir Trail)
- HW 6: Deep learning (Autonomous vehicles)
- HW 7: Reinforcement learning (Pacman revisited)
- HW 8: Vision
- HW 9: Robotics Control

Collaboration Policy

- Homework assignments are your individual responsibility, and plagiarism will not be tolerated.
- You are encouraged to discuss assignments with the instructor, with the TAs, and with other students.
- However, each student is required to implement and write any assignment on their own.
- You will not copy, nor will you allow your work to be copied.

Specifics

- Coding and write up should be done independently
- Do not show your work to anyone
- Do not look at anyone’s work
- Do not use existing code (e.g., github)

Attendance Policy

- Attendance at lectures is critical to success in this course
- Lectures **will** contain material that is not covered by the text (and may not appear on the lecture slides).
- You are responsible for all material presented in lectures, material contained in the assigned reading, and material covered by the homework assignments.

How to Get Help

- Use the right channels for communication
  - Piazza (not canvas)
  - Email (always include **CPSC 470** in the subject line)
  - TAs and ULA staff listed on each assignment