343H: Honors AI

Lecture 27:

Course wrap-up

5/1/2014

Kristen Grauman

UT Austin

Final exam

- Monday May 12, 2-5 pm in CPE 2.218
- Comprehensive
  - Some more emphasis on material since midterm
- Bring two 8.5 x 11” sheets of notes
- No phones, tablets, iPads, etc.
- Practice exam with solutions

Classification mini-contest!

- 3rd Place: Andy Wickham and Tony Wickham
  - 888 correct out of 1000 (88.8%)
- 2nd Place: Adam Faulkner
  - 904 correct out of 1000 (90.4%)
- 1st Place: David Yu and Qi Guo
  - 915 correct out of 1000 (91.5%)

Big picture recap

- 1. Making decisions
  - Search, planning
  - Adversarial and uncertain search
- 2. Reasoning under uncertainty
  - Bayes' nets
  - Decision theory
  - Machine learning

Search

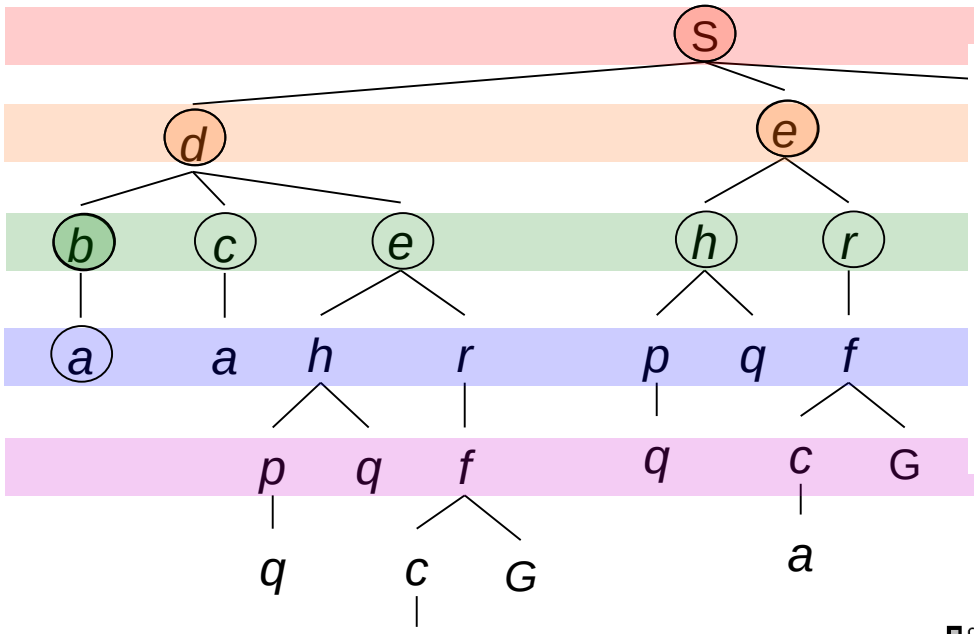

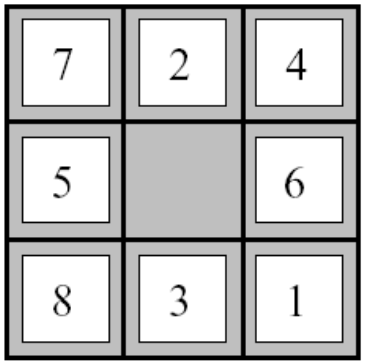

Start State

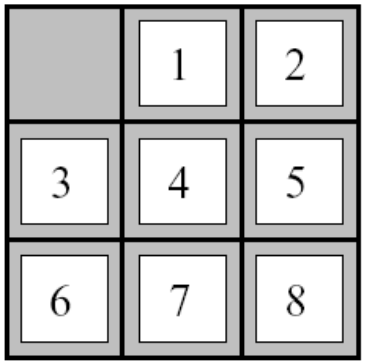

Goal State

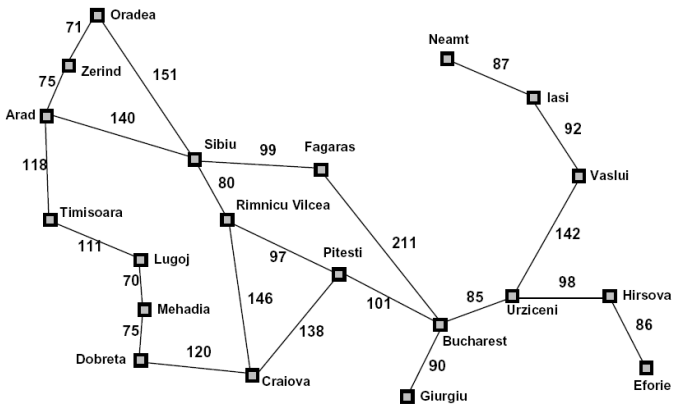

Straight−line distance to Bucharest	
Arad	366
Bucharest	0
Craiova	160
Dobreta	242
Eforie	161
Fagaras	178
Giurgiu	77
Hirsova	151
Iasi	226
Lugoj	244
Mehadia	241
Neamt	234
Oradea	380
Pitesti	98
Rimnicu Vilcea	193
Sibiu	253
Timisoara	329
Urziceni	80
Vaslui	199
Zerind	374

Adversarial search

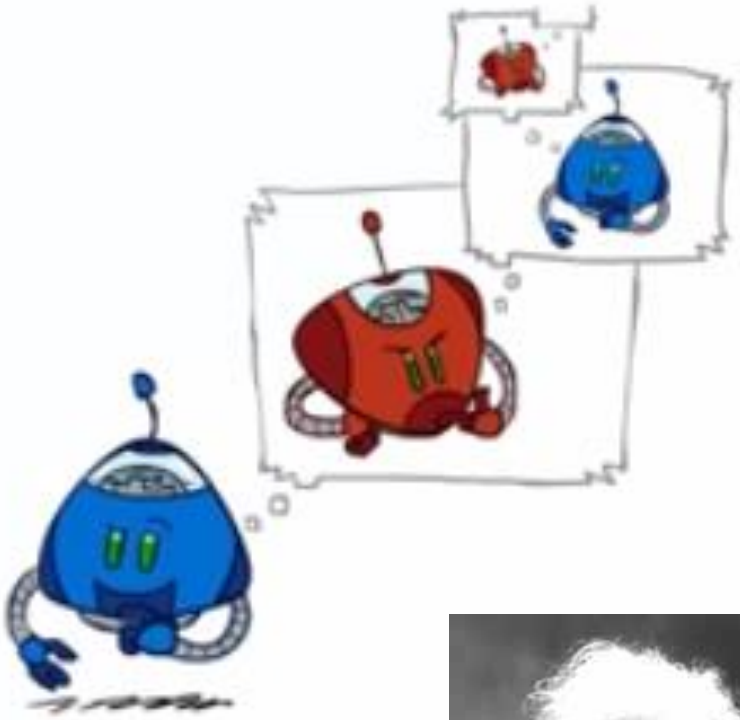

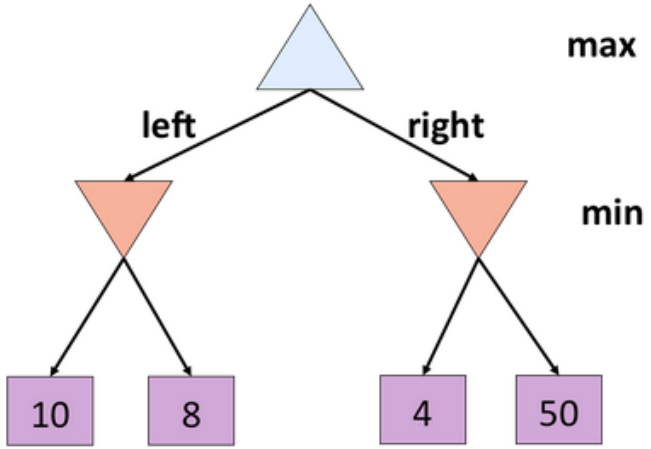

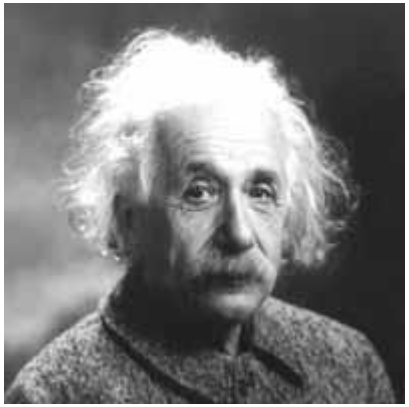

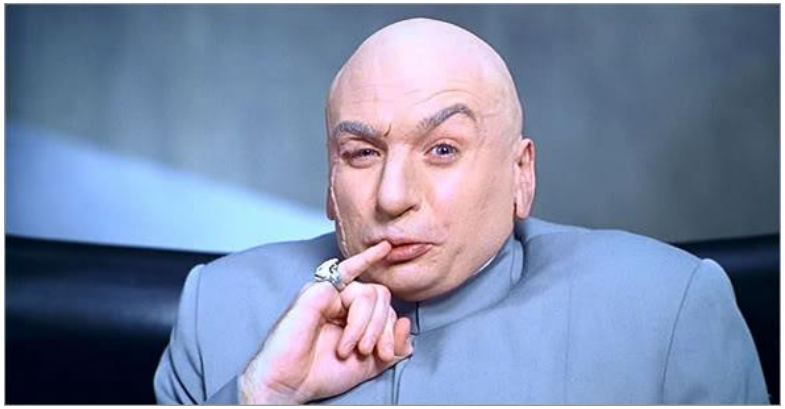

Utilities

20 points

10 points

5 points

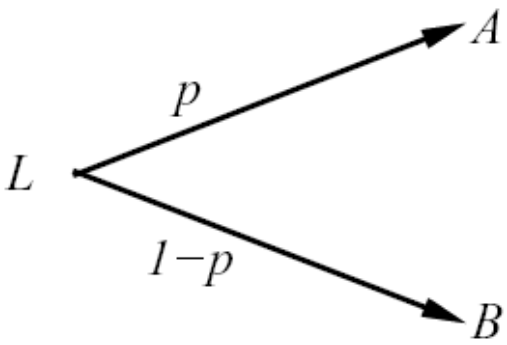

Markov Decision Processes

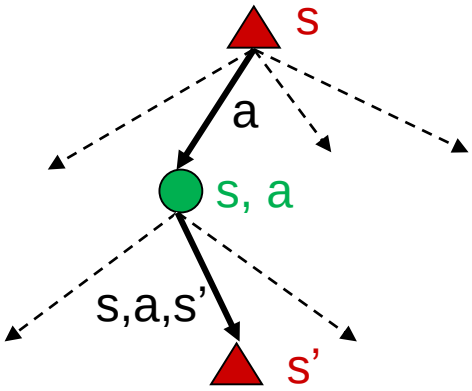

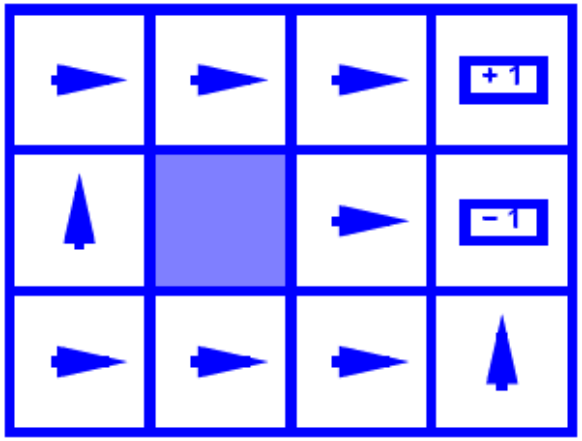

Reinforcement learning

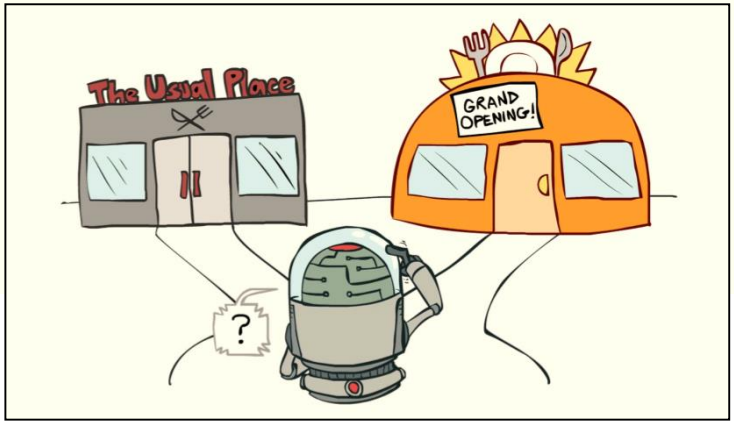

Bayesian networks

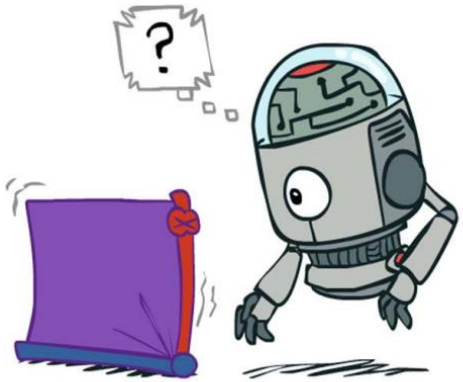

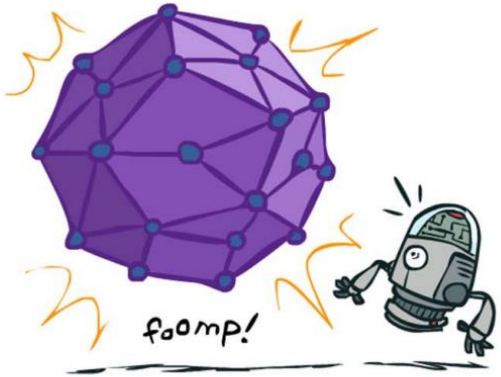

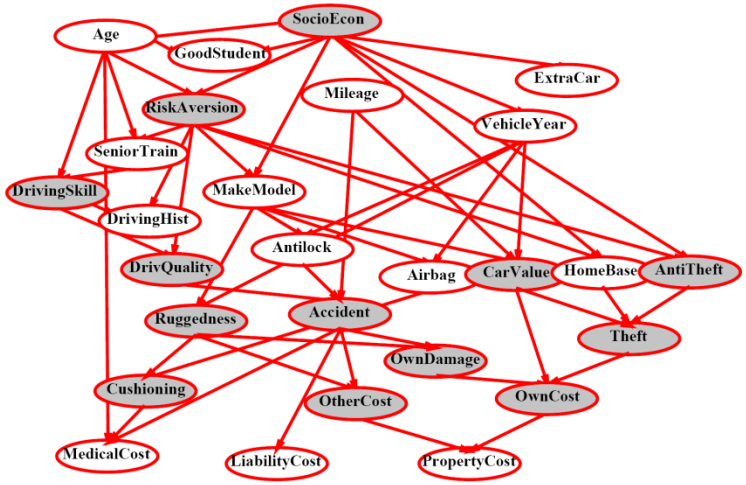

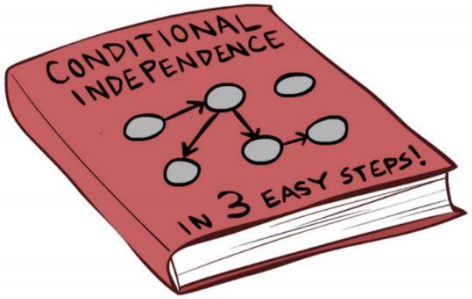

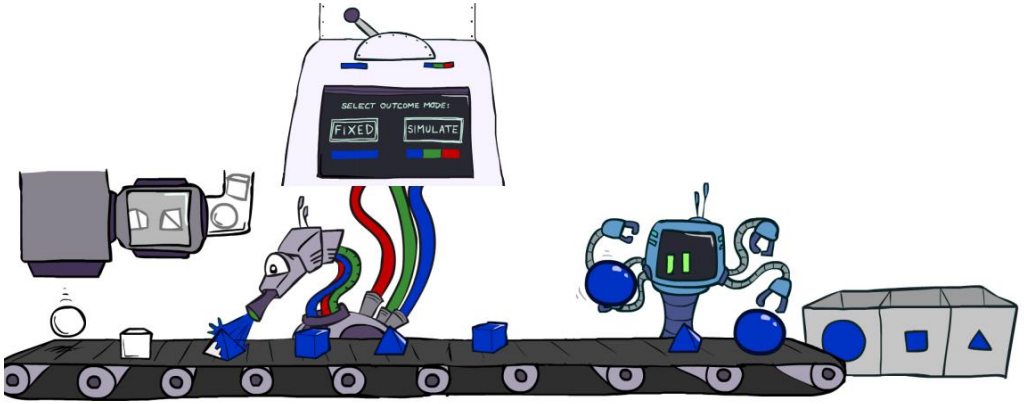

Decision networks & VPI

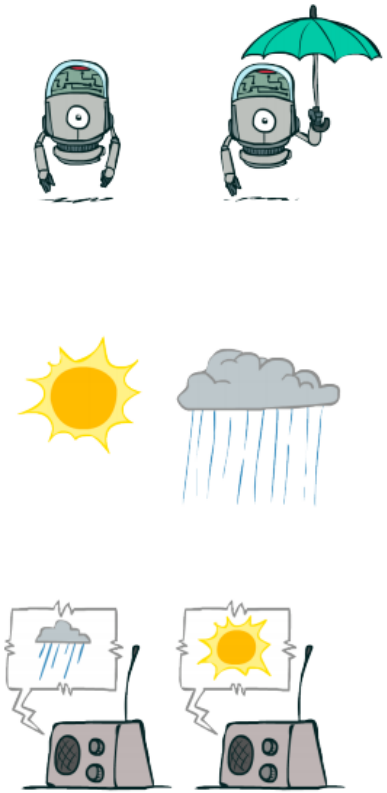

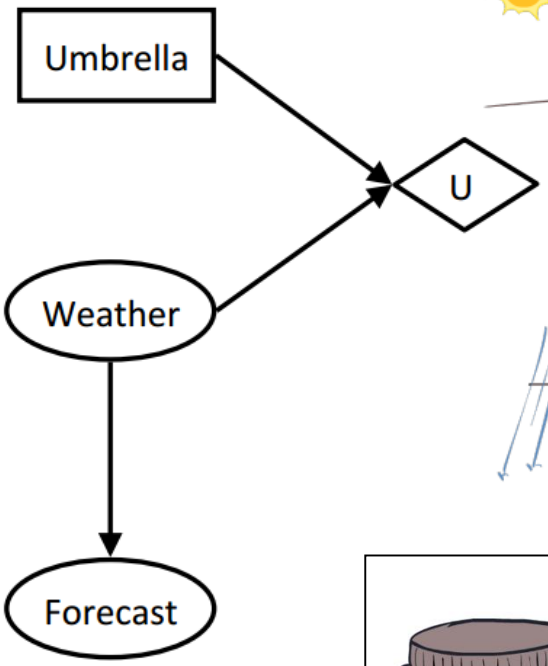

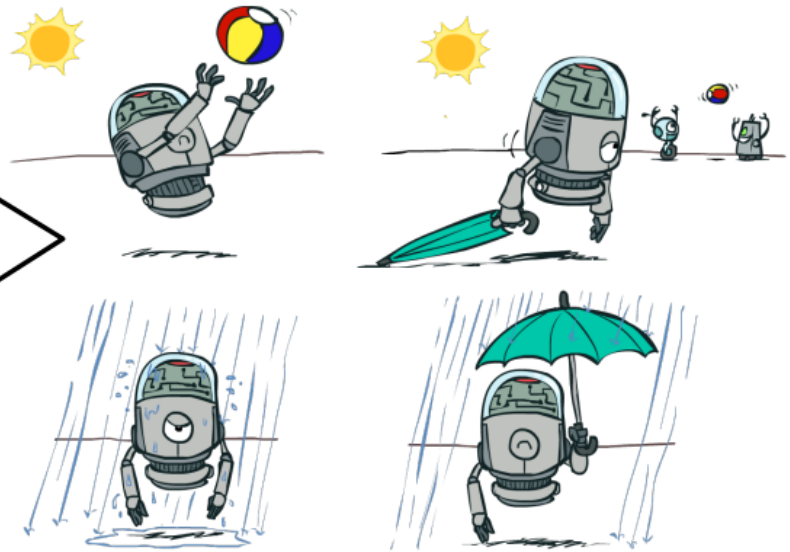

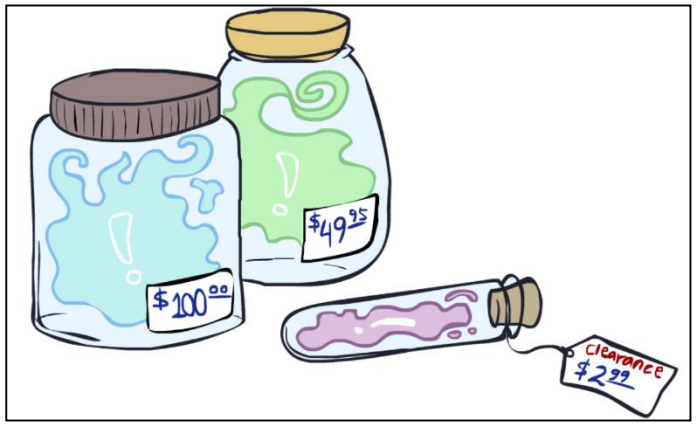

Probabilistic reasoning over time

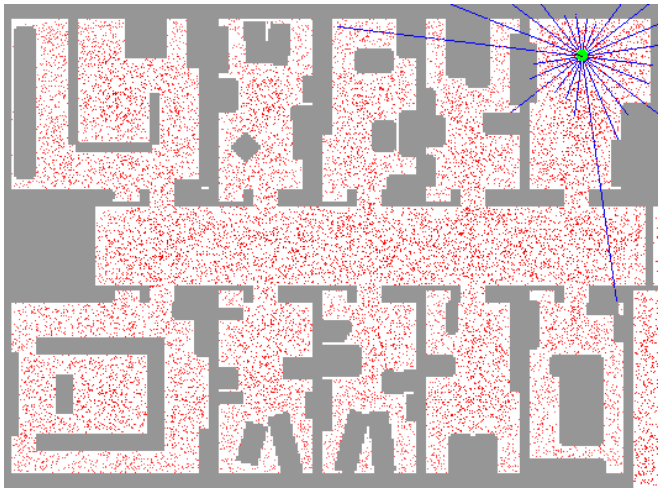

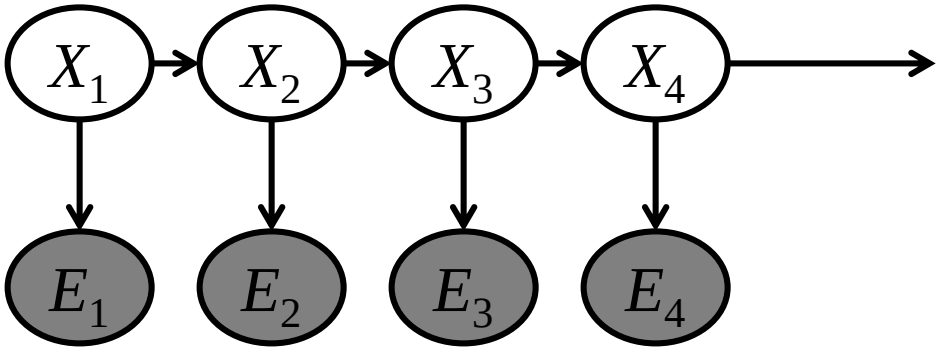

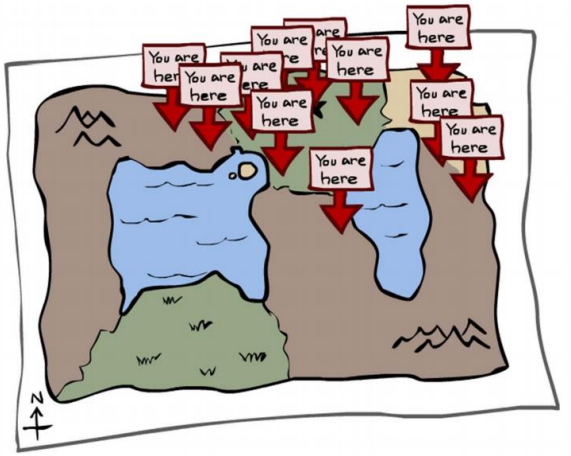

Learning (Classifiers, Clustering)

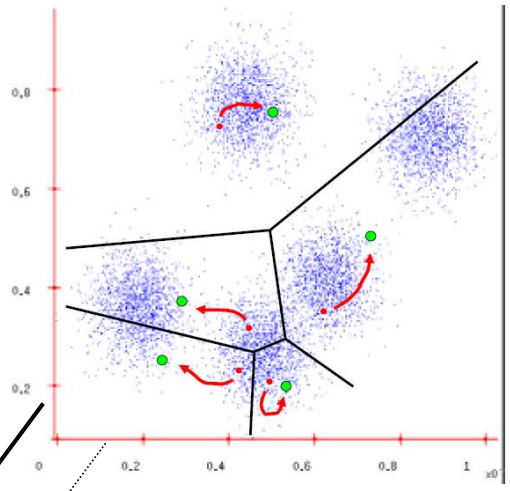

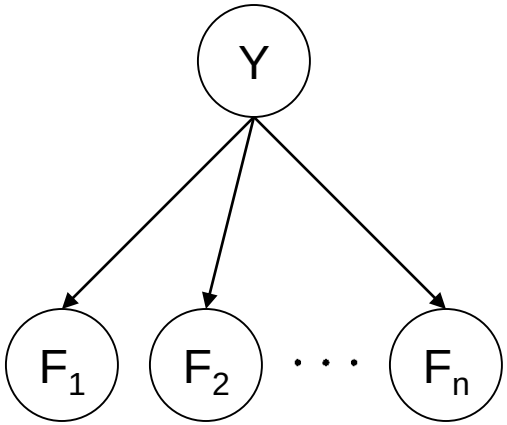

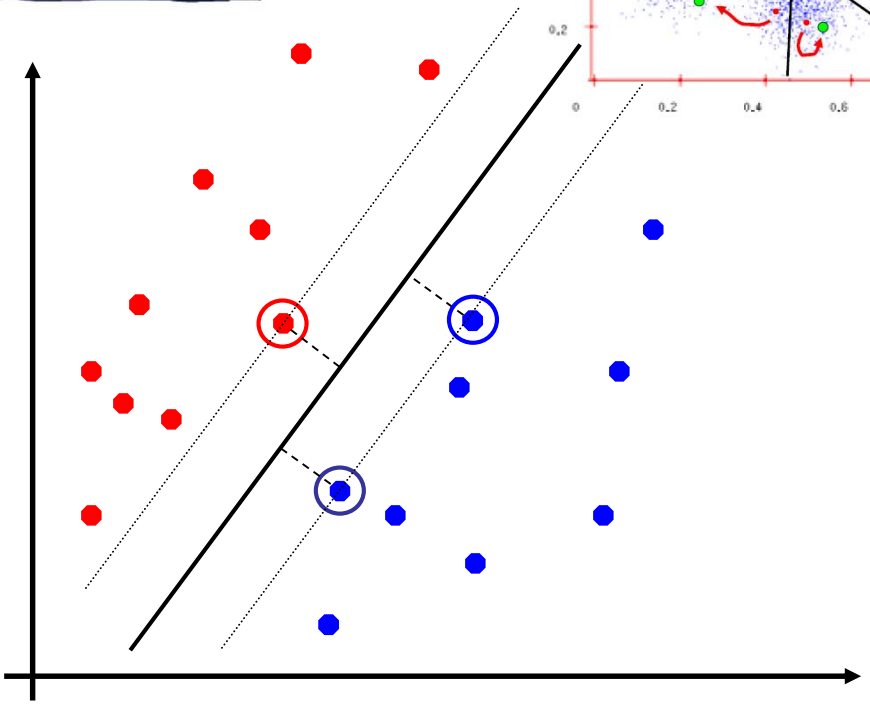

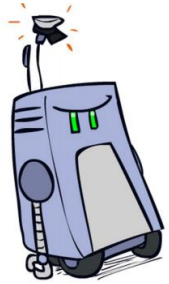

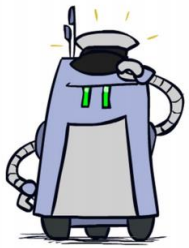

Intelligent agents

- Sense, decide, act
- Maximize expected utility

Things we didn’t cover

- Constraint satisfaction
- Knowledge representation and reasoning
- Game theory and auctions
- Aspects of learning
- Natural language
- Vision
- Robotics
- …

Strategy in AI

- A goal of AI: Robust, fully autonomous agents in the real world

- Bottom-up metaphor:

  Russell, '95: *"Theoreticians can produce the AI equivalent of bricks, beams, and mortar with which AI architects can build the equivalent of cathedrals."*

Slide credit: Peter Stone

Bottom-up approach

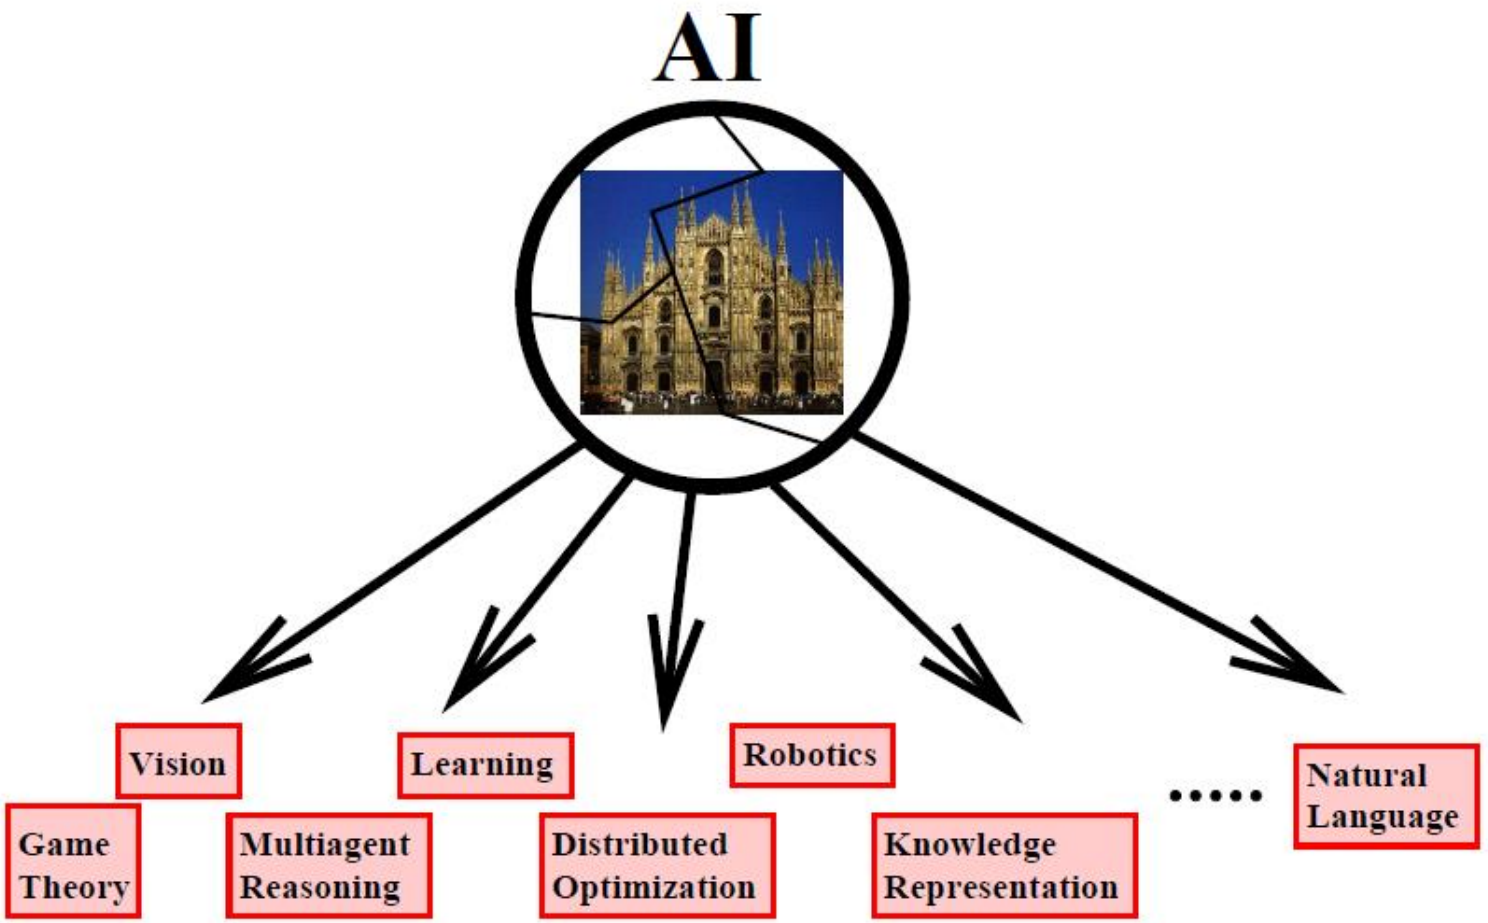

Slide credit: Peter Stone

The bricks

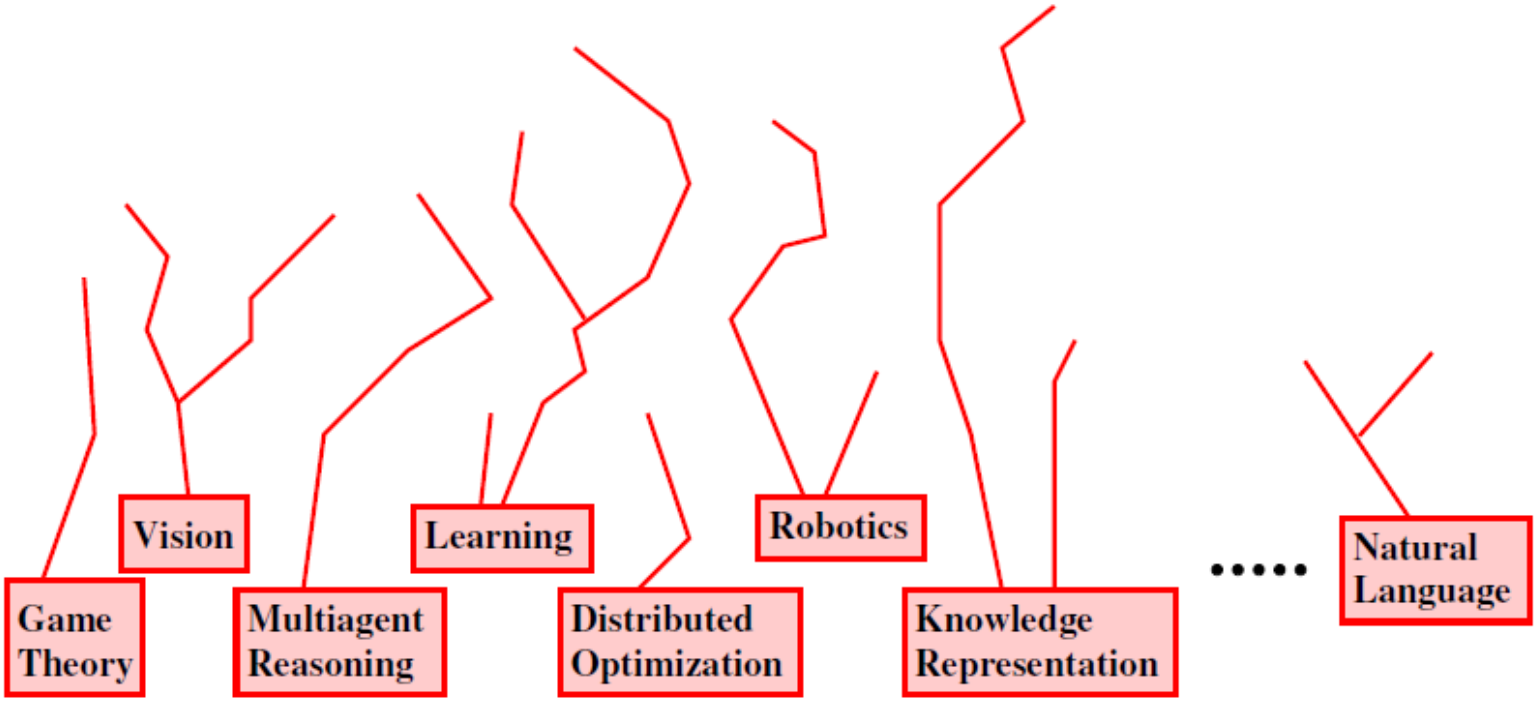
Game Theory
Vision
Multiagent Reasoning
Learning
Distributed Optimization
Robotics
Knowledge Representation
.....
Natural Language

The beams and mortar

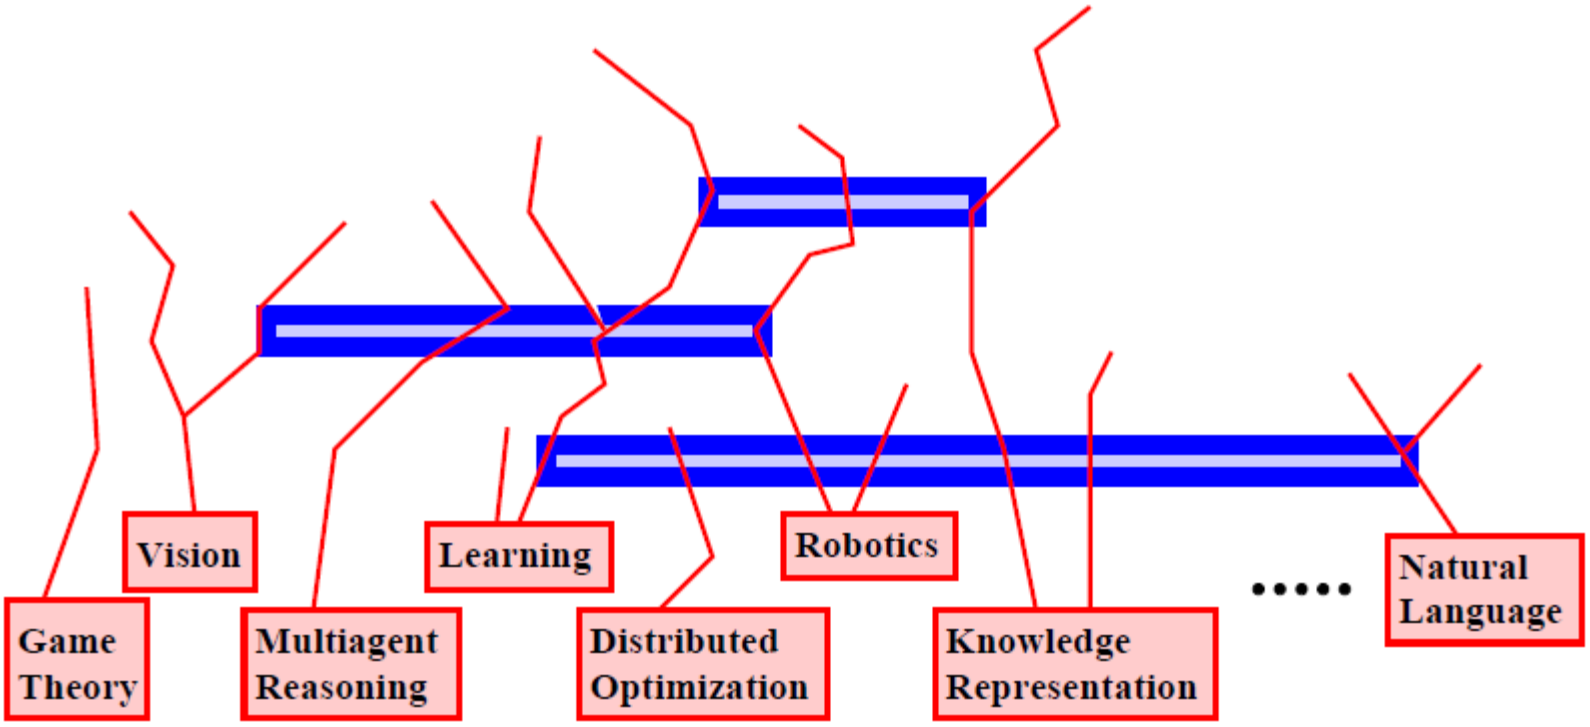

Towards and cathedral?

?

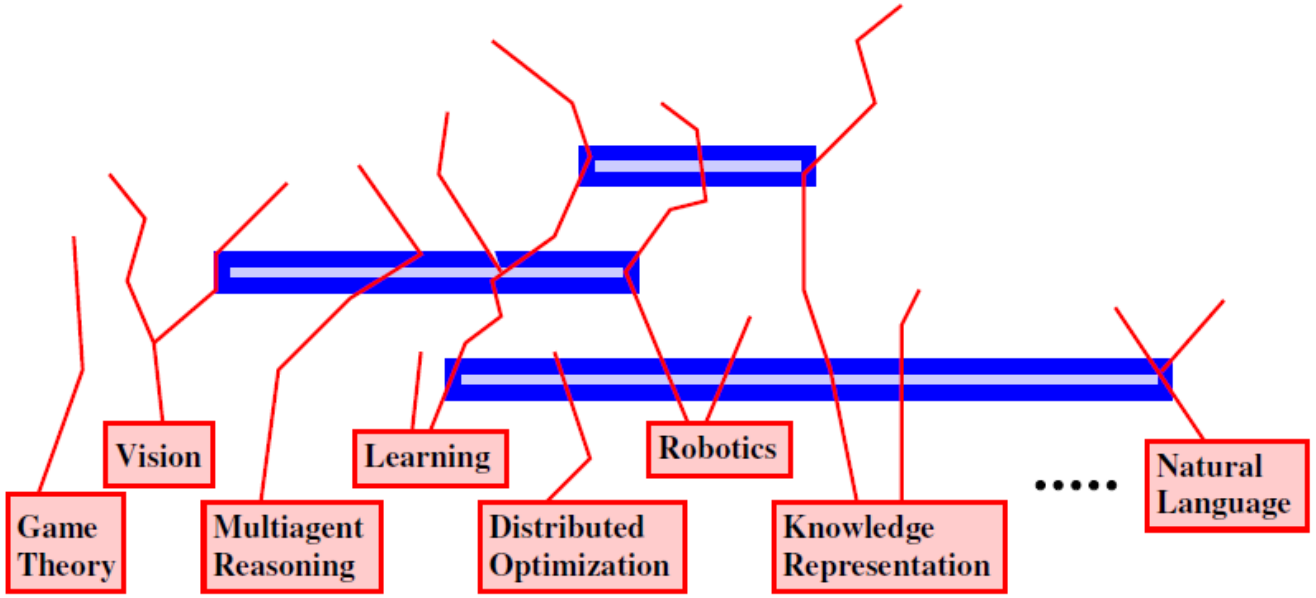

Or something else?

?

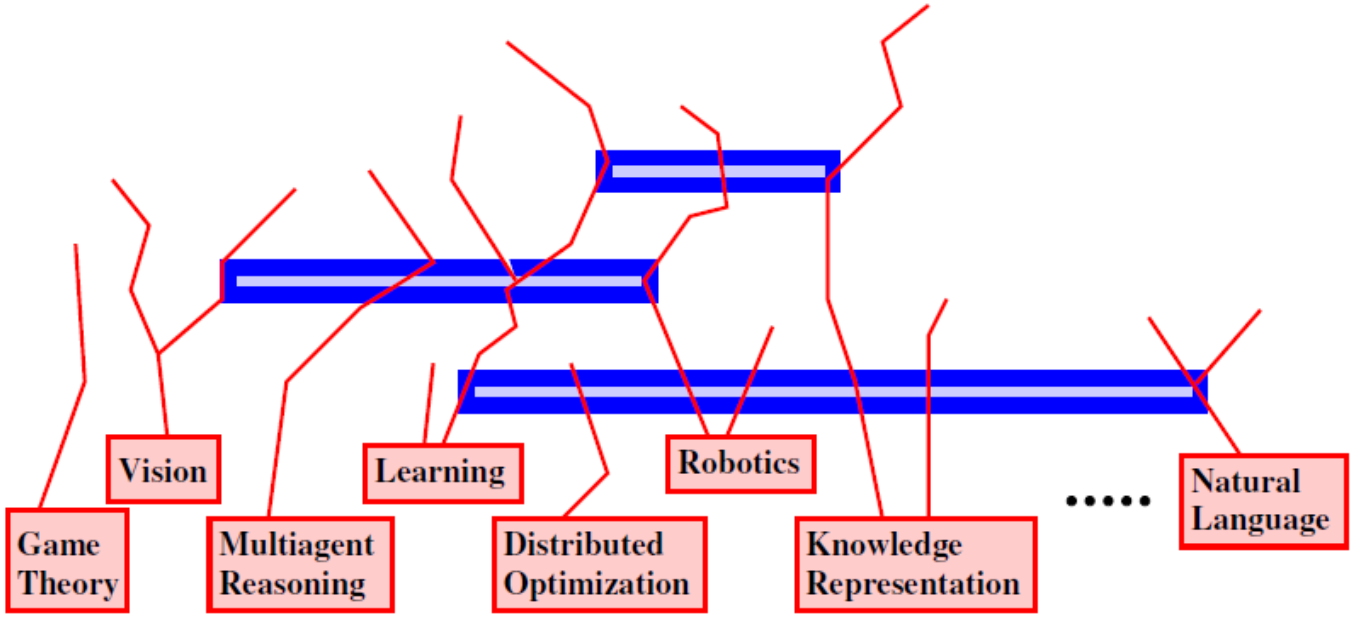

Top-down approach

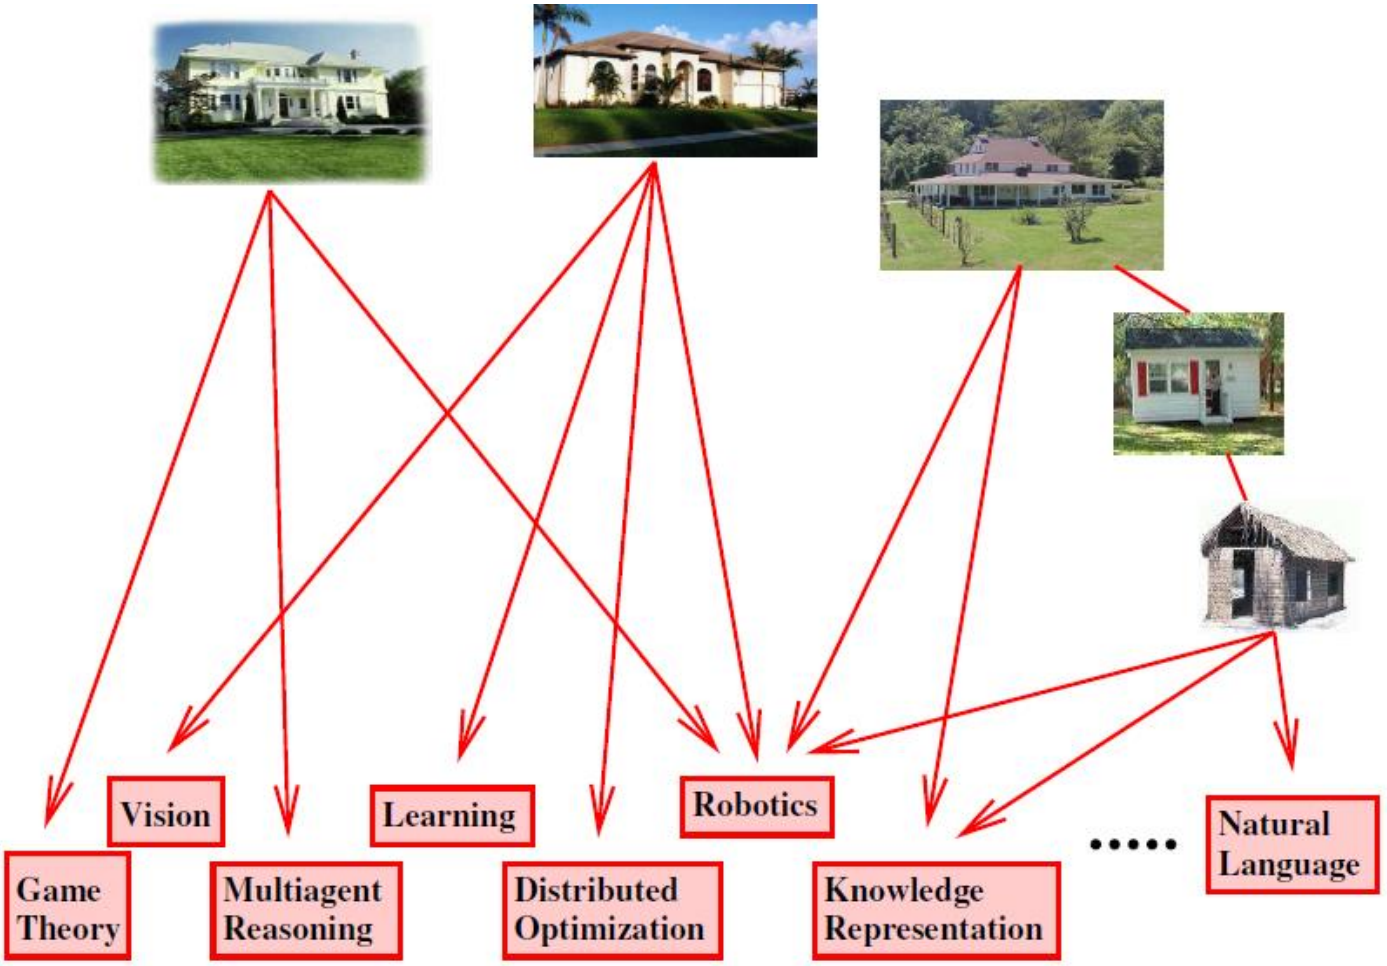

"Good problems . . . produce good science" [Cohen, '04]

Meeting in the middle

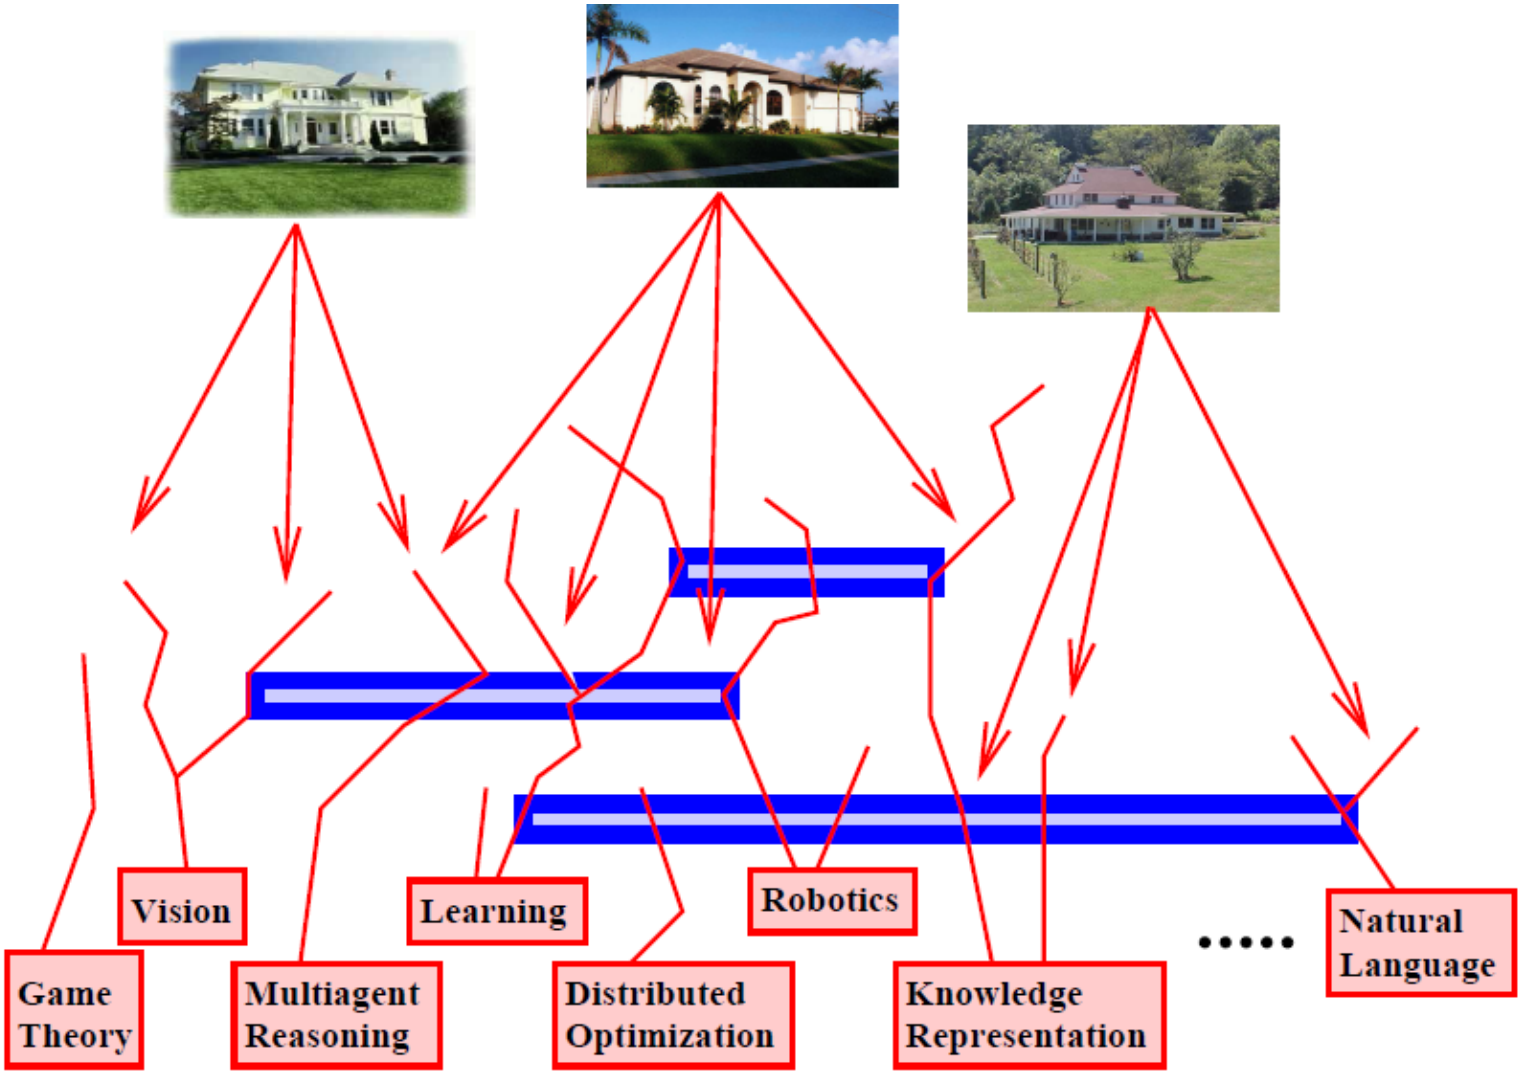

Ethics, implications

- Robust, fully autonomous agents in the real world

- What happens when we achieve this goal?

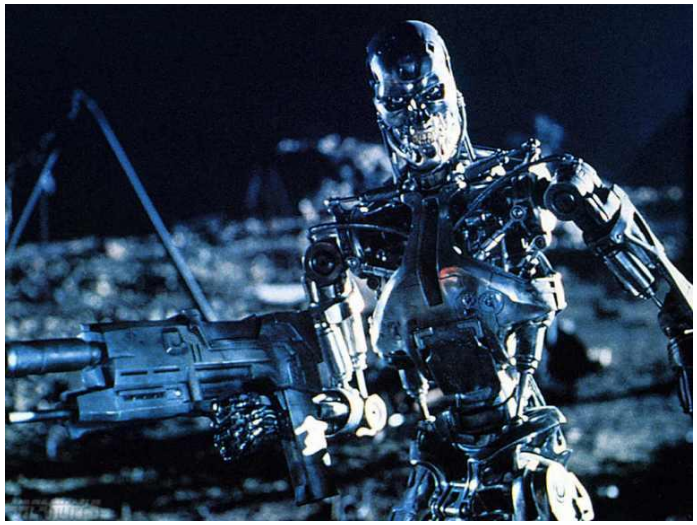

Some Hard Questions…

- Who is liable if a robot driver has an accident?
- Will machines surpass human intelligence?
- What will we do with superintelligent machines?
- Would such machines have conscious existence? Rights?
- Can human minds exist indefinitely within machines (in principle)?

Slide credit: Dan Klein

Tournament highlights

Tournament results

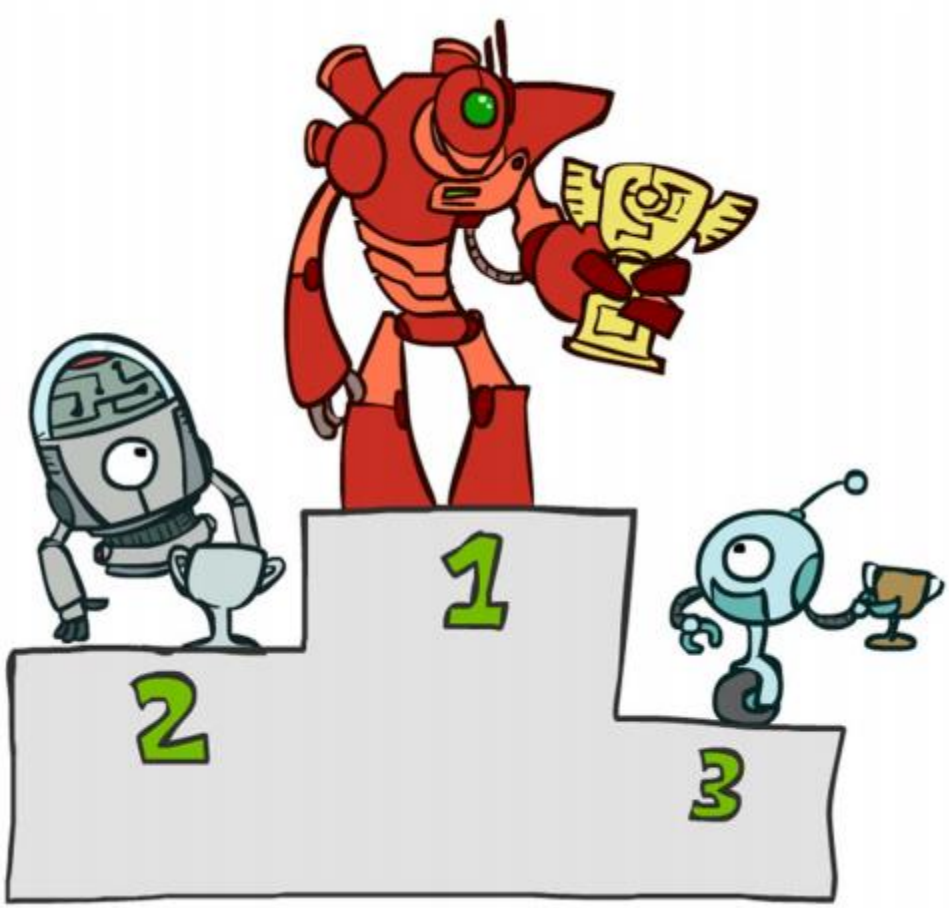

Winners

- 3rd place: Josh Kelle
- 2nd place: Sam Thompson
- 1st place: Jaime Rivera and K. K.

- Congrats to all!

Roomba pacman

Students at Colorado University: http://pacman.elstonj.com

[DEMO]

Bugman

- AI = Animal Intelligence?
  - Wim van Eck at Leiden University
  - Pacman controlled by a human
  - Ghosts controlled by crickets
  - Vibrations drive crickets toward or away from Pacman's location

[DEMO]

http://pong.hku.nl/~wim/bugman.htm