


CS343: Artificial Intelligence

Introduction



Profs. Peter Stone and Yuke Zhu
The University of Texas at Austin

[Based on slides created by Dan Klein and Pieter Abbeel for CS188 Intro to AI at UC Berkeley, modified by Scott Niekum at UT Austin.]


All materials available at <http://ai.berkeley.edu>.]

Good Morning Colleagues

- **Welcome to a fun, but challenging course**
- **Goal: Learn about Artificial Intelligence**
 - Increase AI literacy (it's not magic!)
 - Prepare you for topics courses
 - Broad coverage of topics
 - NOT a current events class

A (Short) History of AI

- **1940-1950: Early days**
 - 1943: McCulloch & Pitts: Boolean circuit model of brain
 - 1950: Turing's "Computing Machinery and Intelligence"
- **1950—70: Excitement: Look, Ma, no hands!**
 - 1950s: Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
 - 1956: Dartmouth meeting: "Artificial Intelligence" adopted
 - 1965: Robinson's complete algorithm for logical reasoning
- **1970—90: Knowledge-based approaches**
 - 1969—79: Early development of knowledge-based systems
 - 1980—88: Expert systems industry booms
 - 1988—93: Expert systems industry busts: "AI Winter"
- **1990—: Statistical approaches**
 - Resurgence of probability, focus on uncertainty
 - General increase in technical depth
 - Agents and learning systems... "AI Spring"?
- **2010-: Neural networks (deep learning)**
 - Great progress in vision, NLP.. "AI Summer"?



The Big Scientific Questions of our Time

- How did the **universe** originate?
- How did **life** on Earth originate?
- What is the nature of **intelligence**?

The Nature of Intelligence

How can we Study it?

- Study human (or animal) **behavior** – Psychology
- Study human (or animal) **brains** – Neuroscience
- **Think** about it – Philosophy
- **Build and analyze** intelligent artifacts – Computer Science

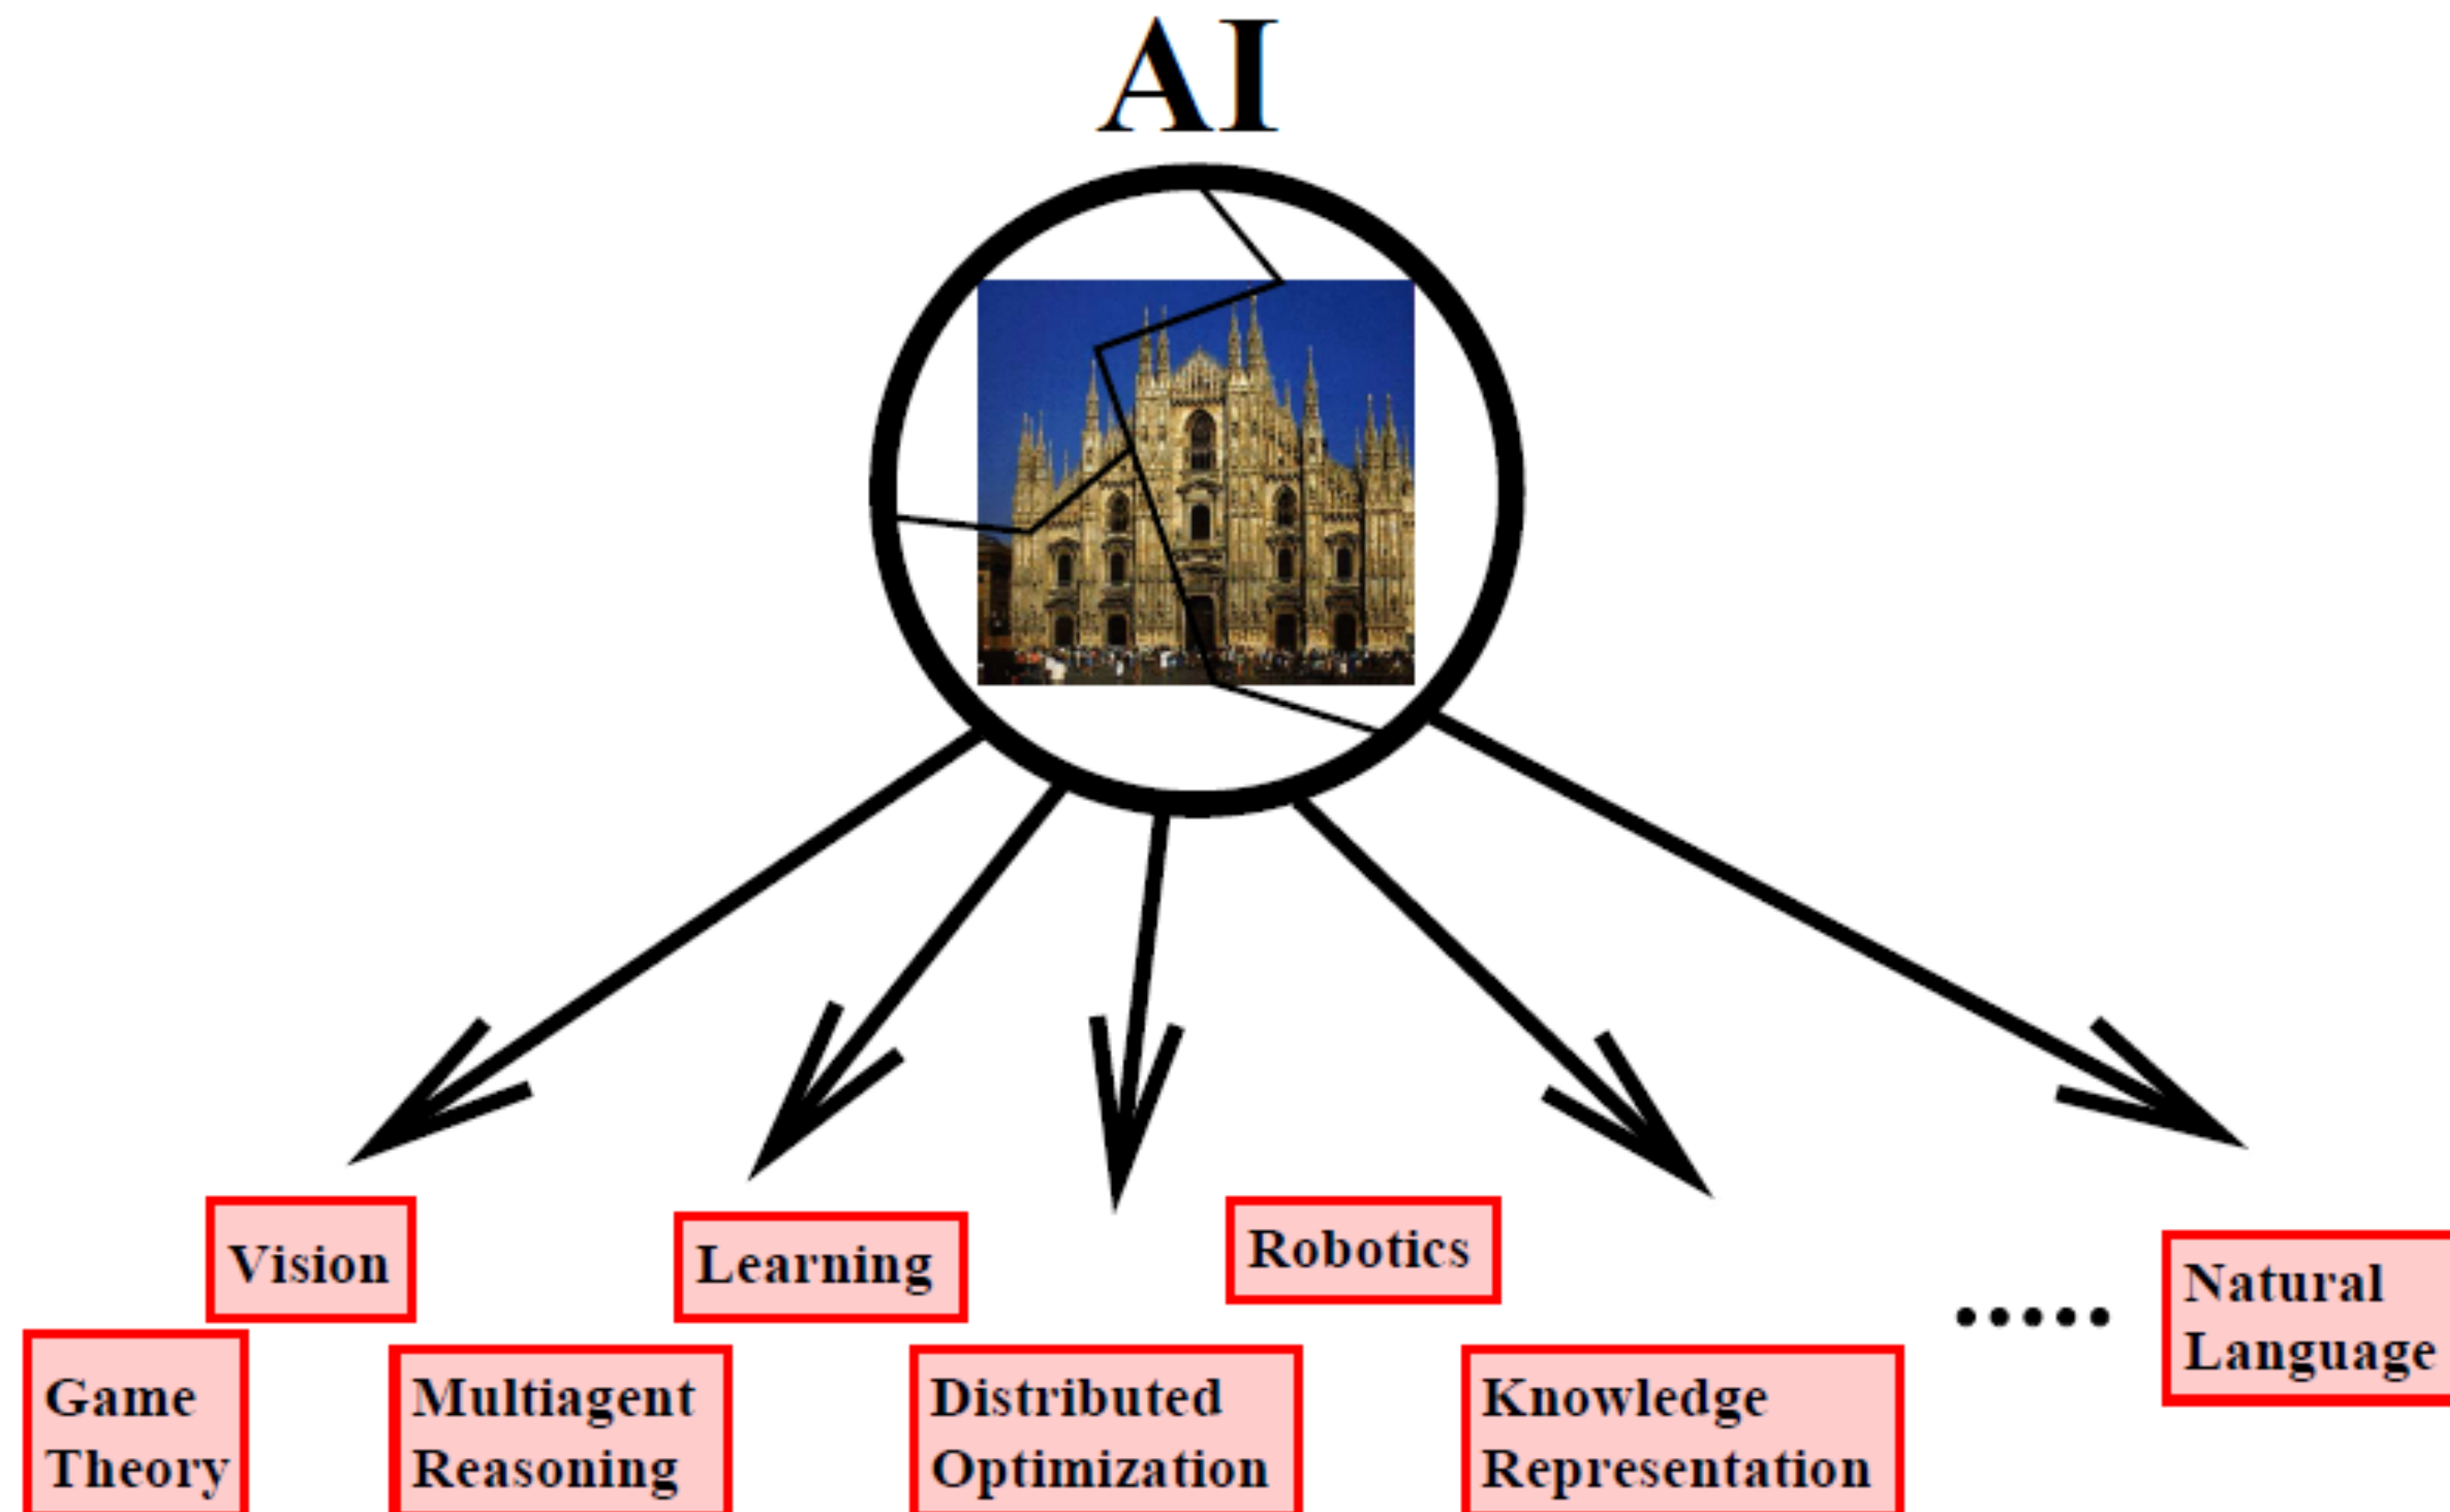
Building Intelligent Artifacts

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

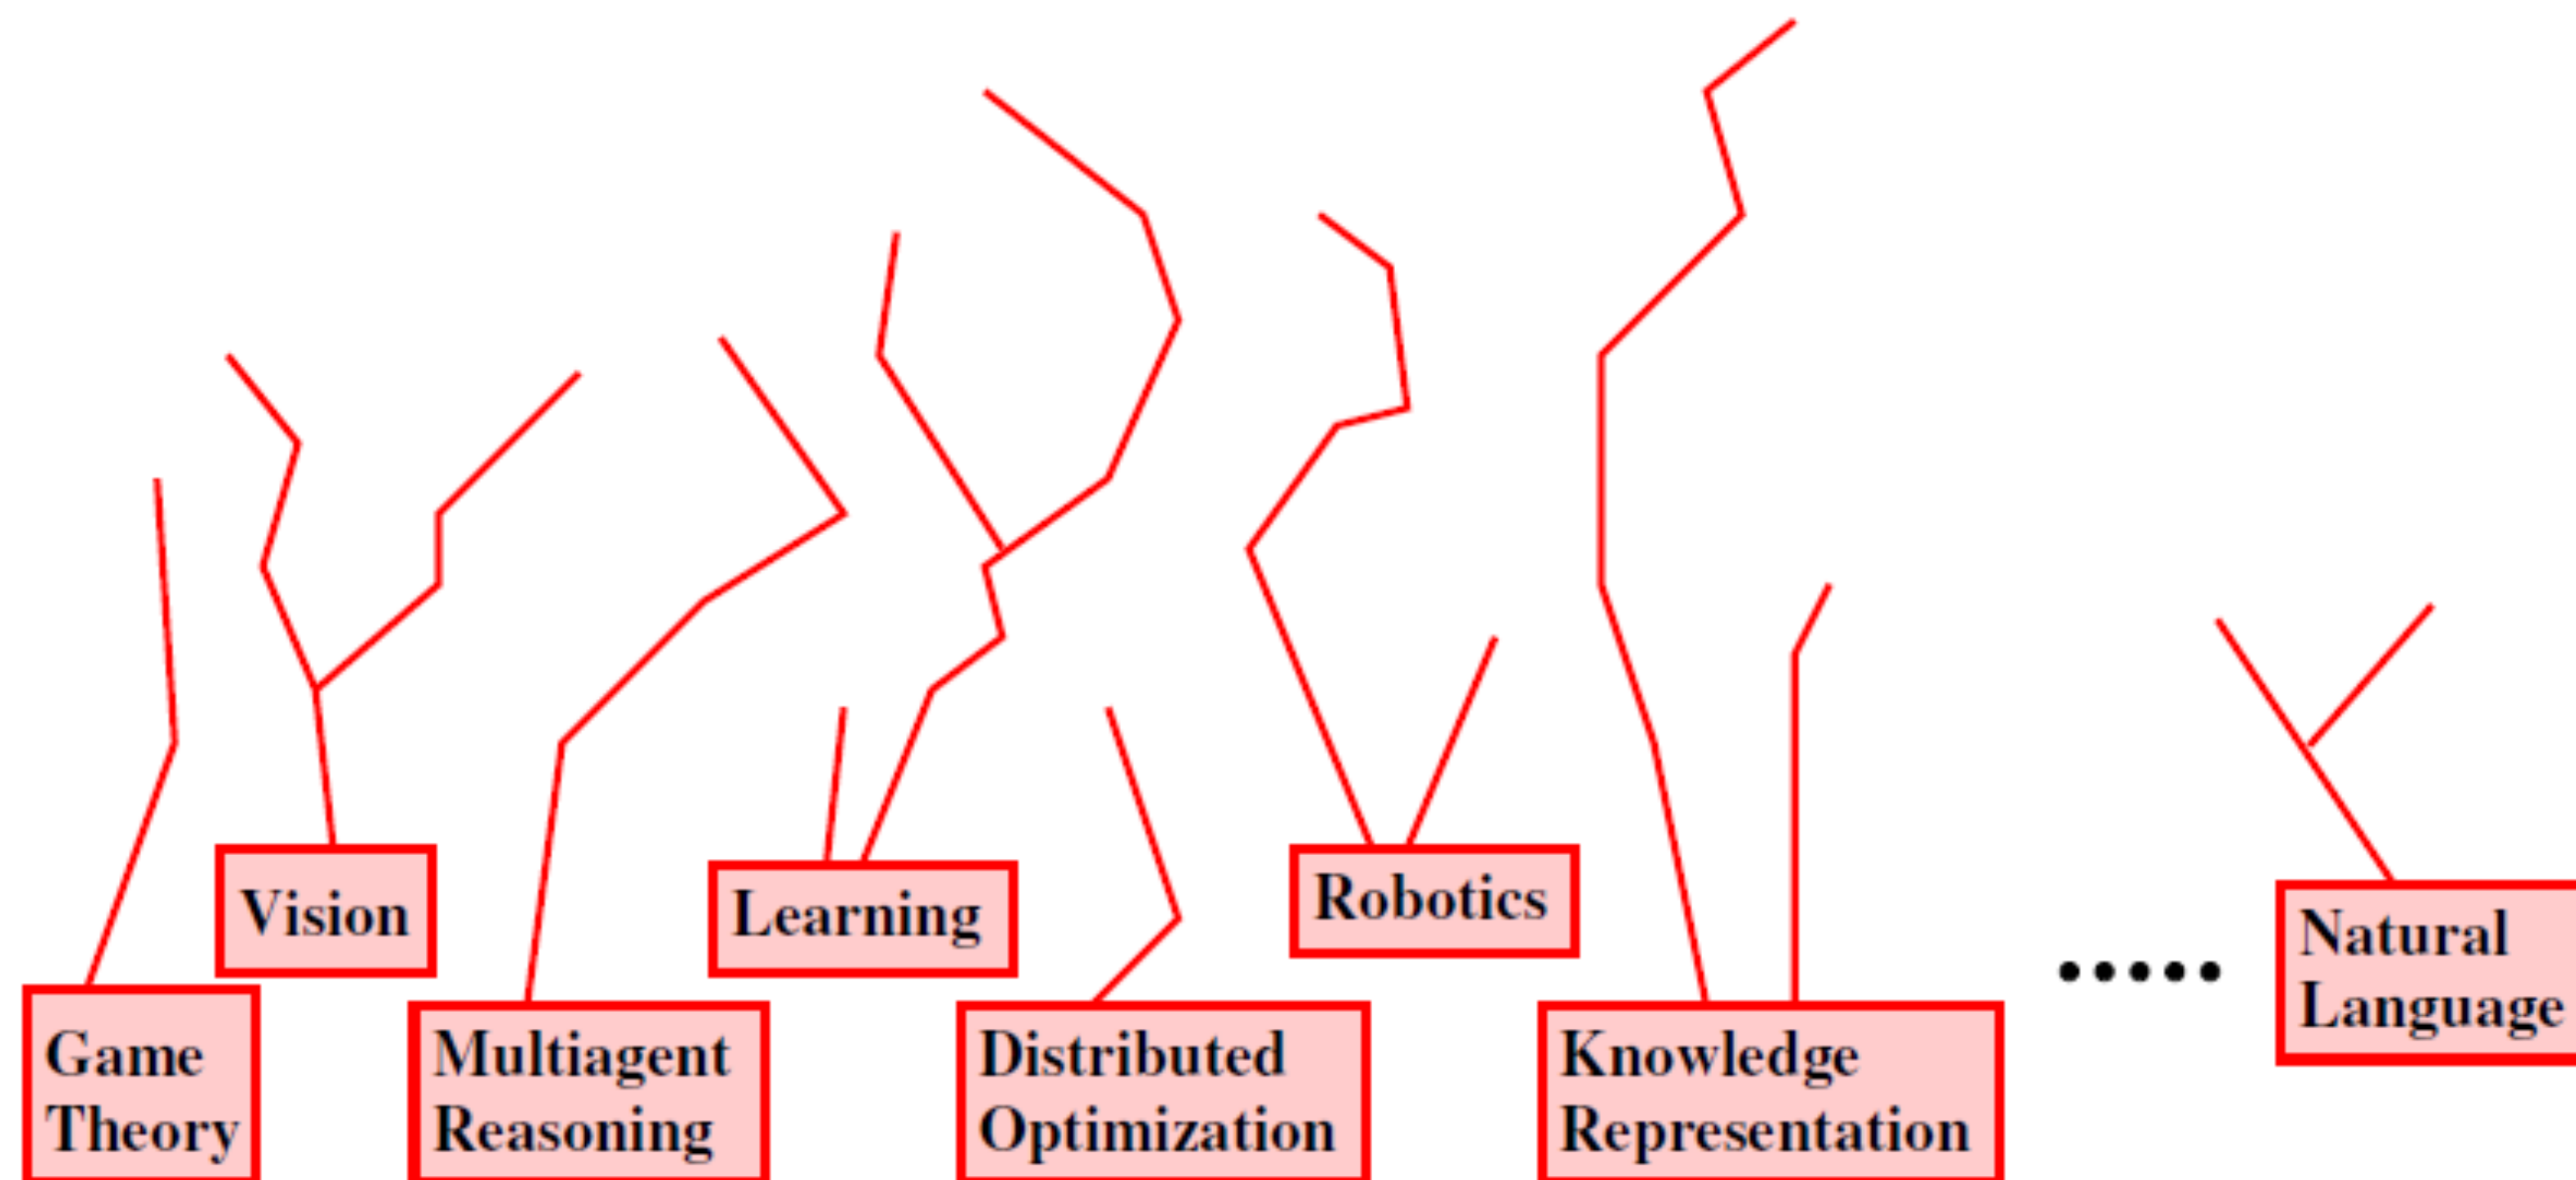
- Bottom-up metaphor:

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

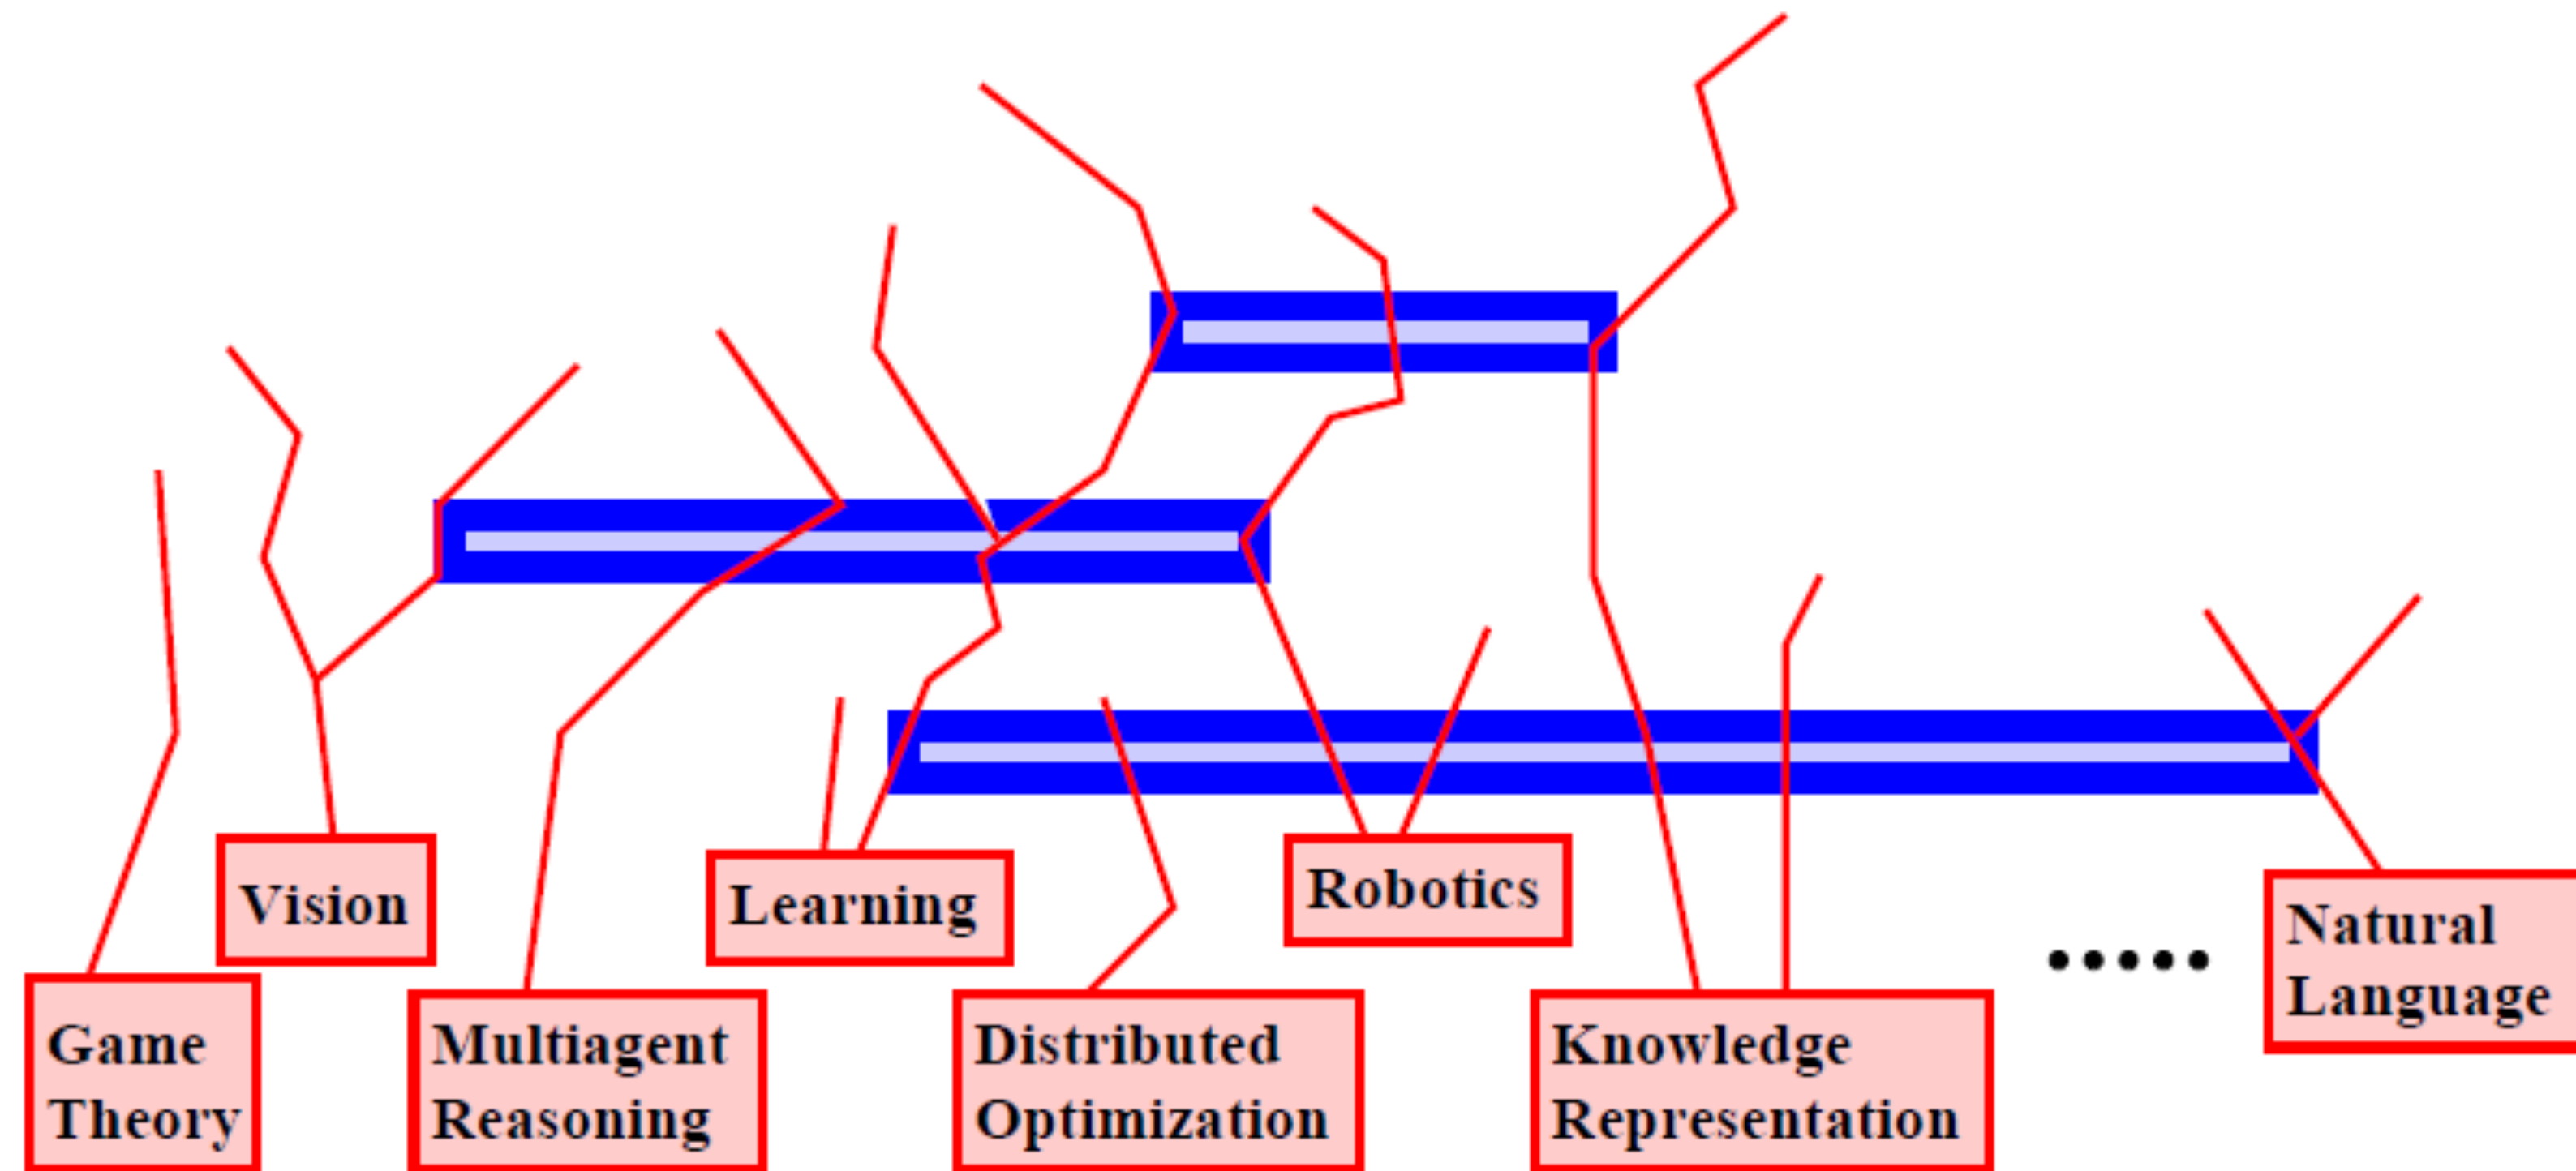
Bottom-up approach



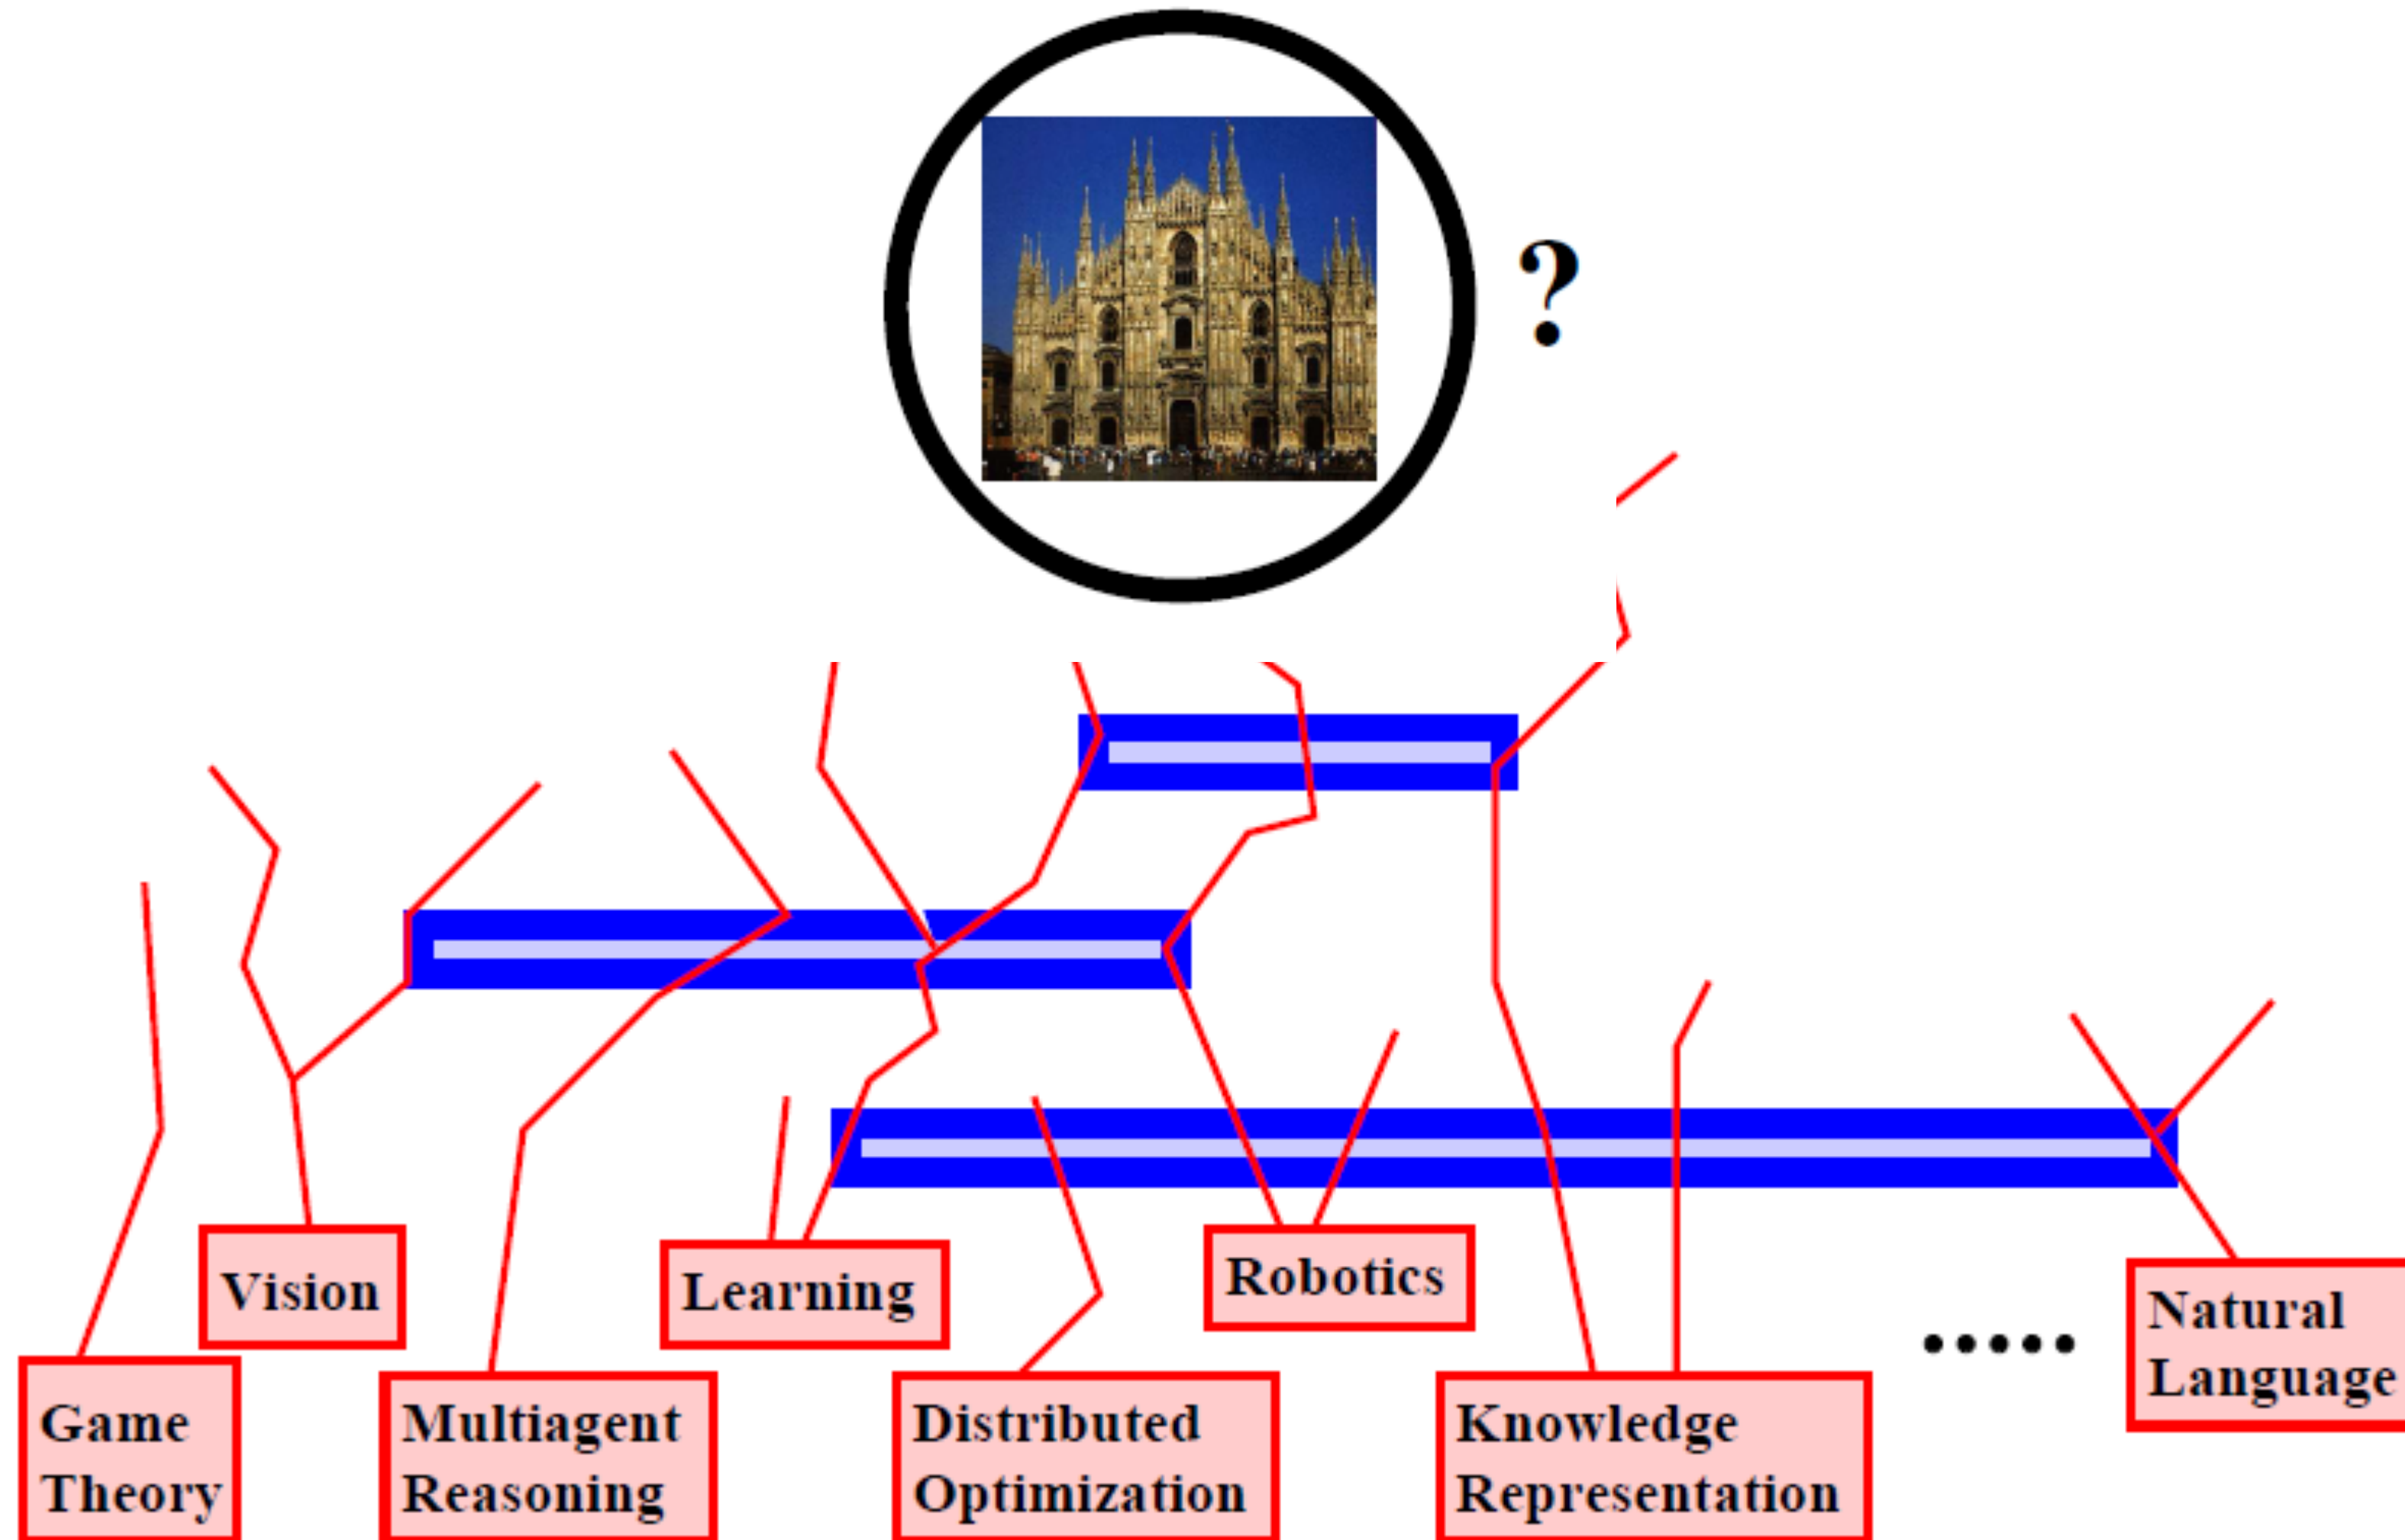
The bricks



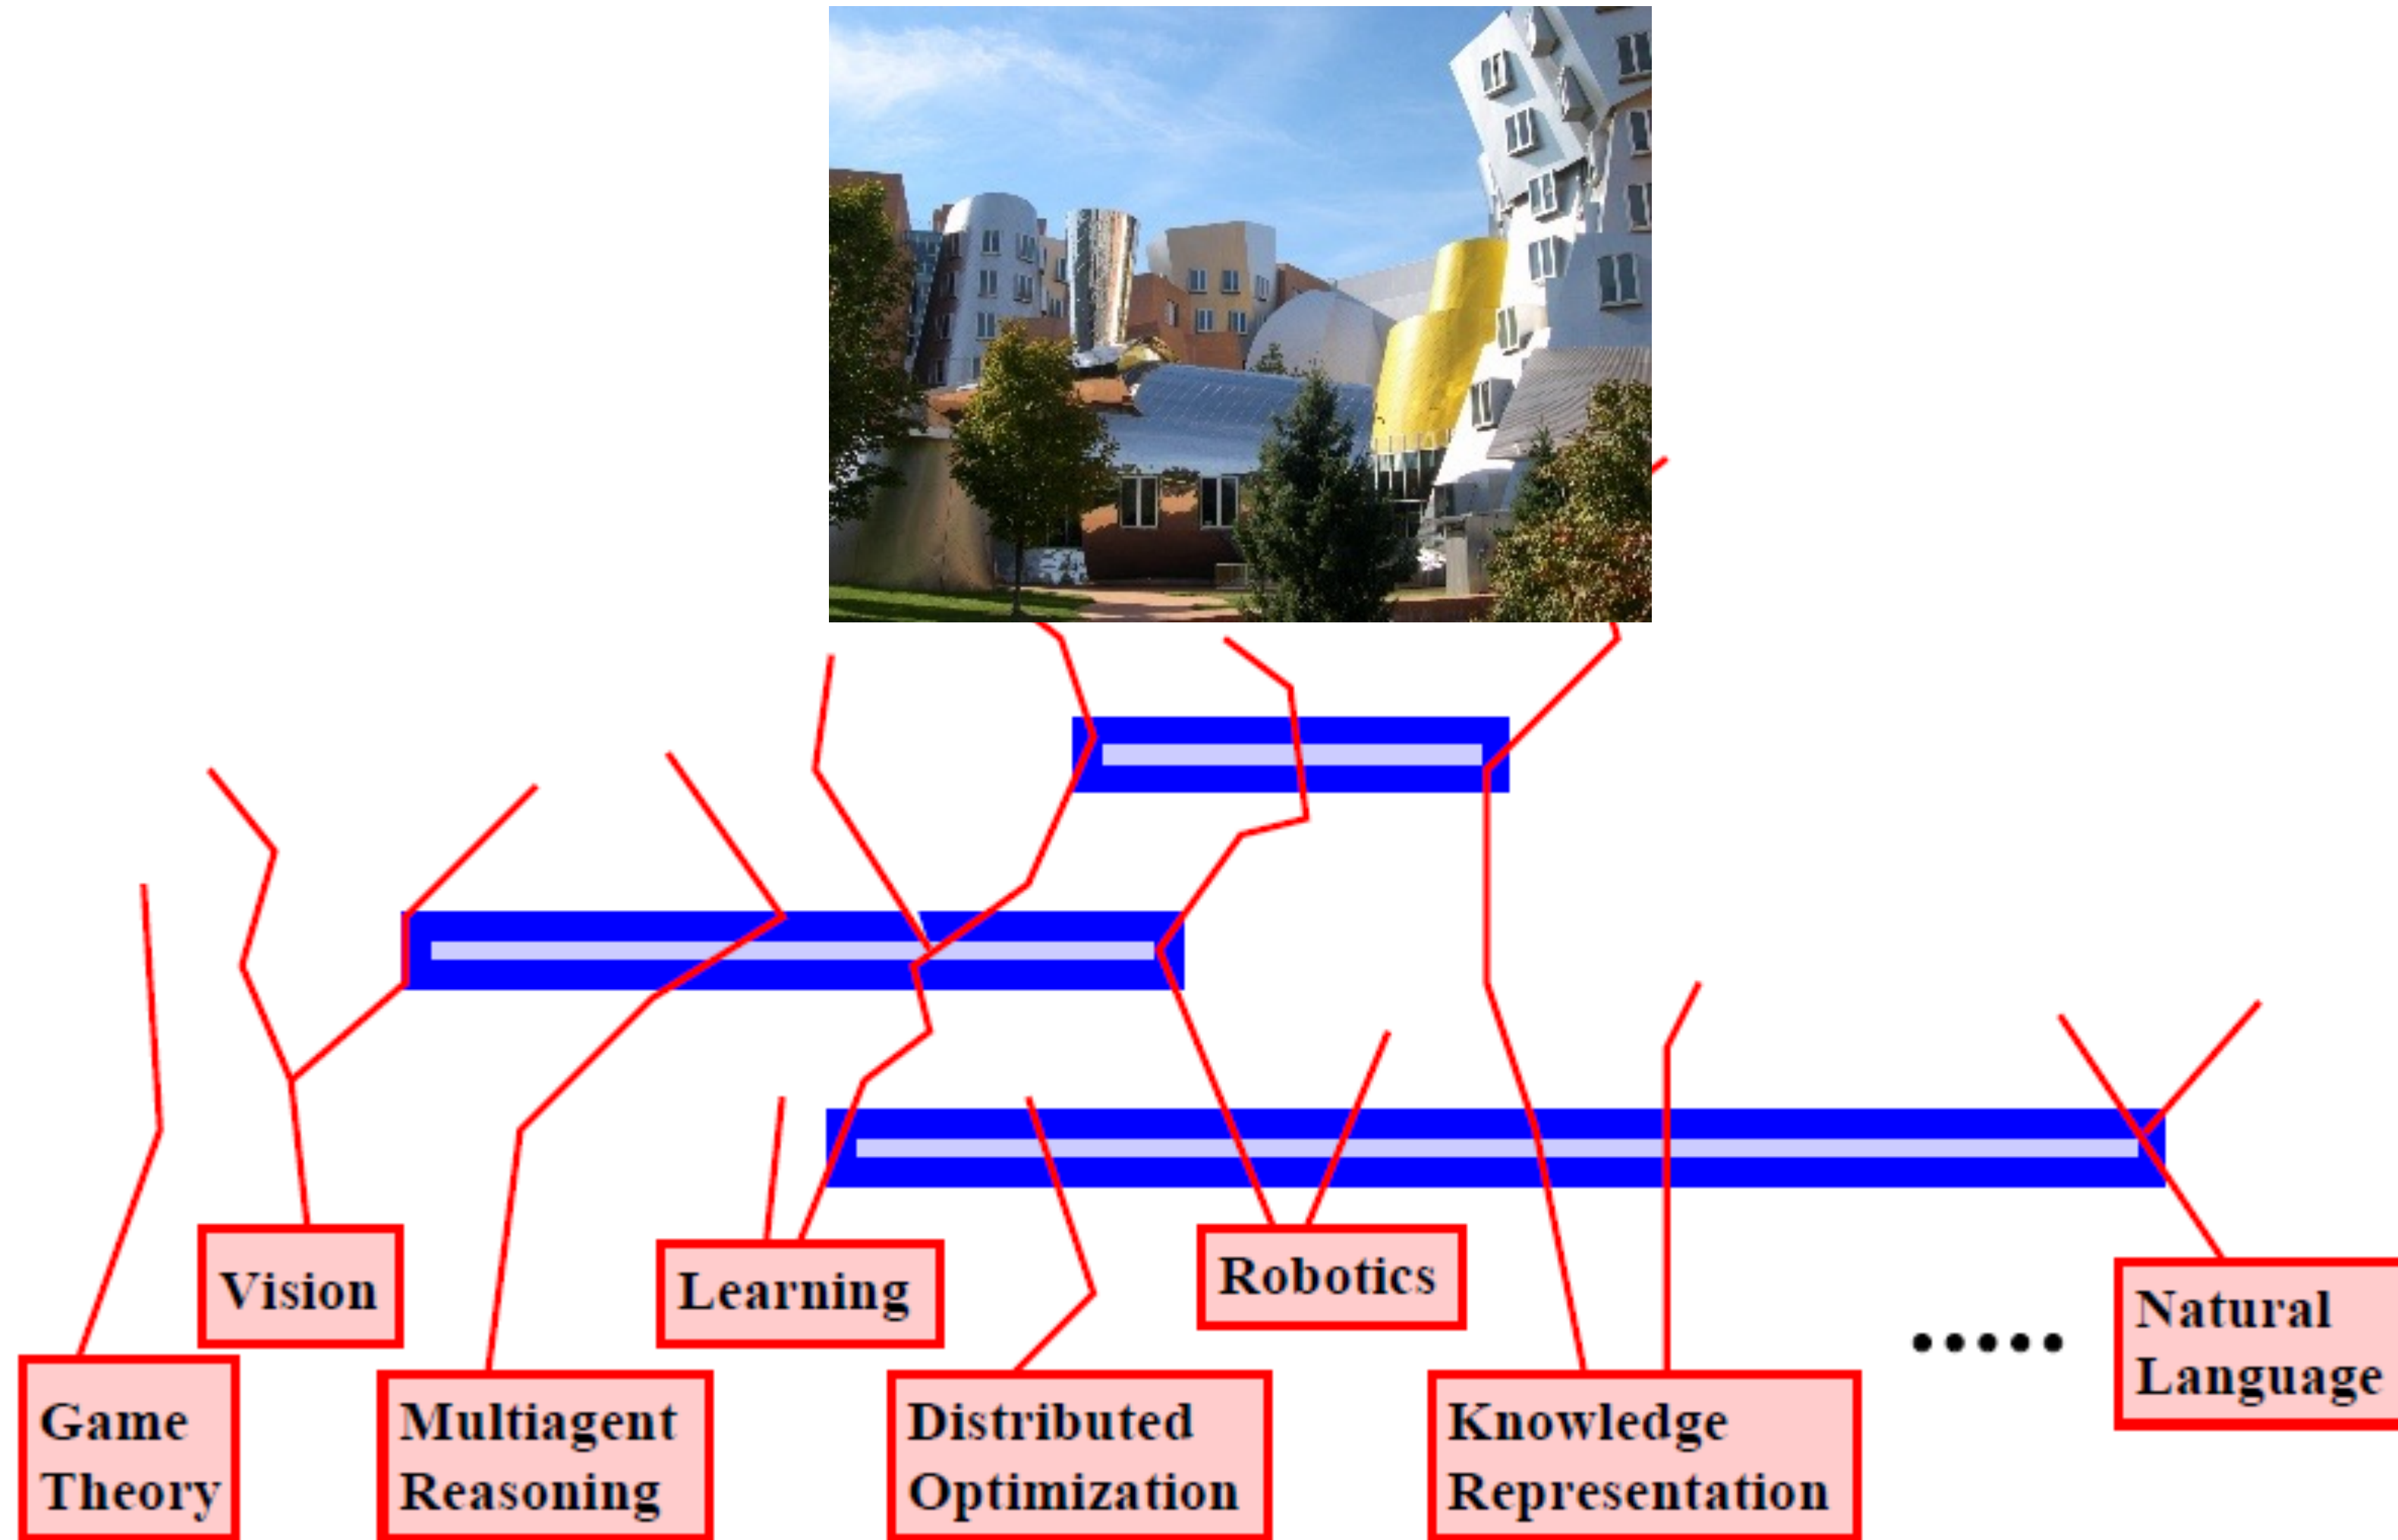
The beams and mortar



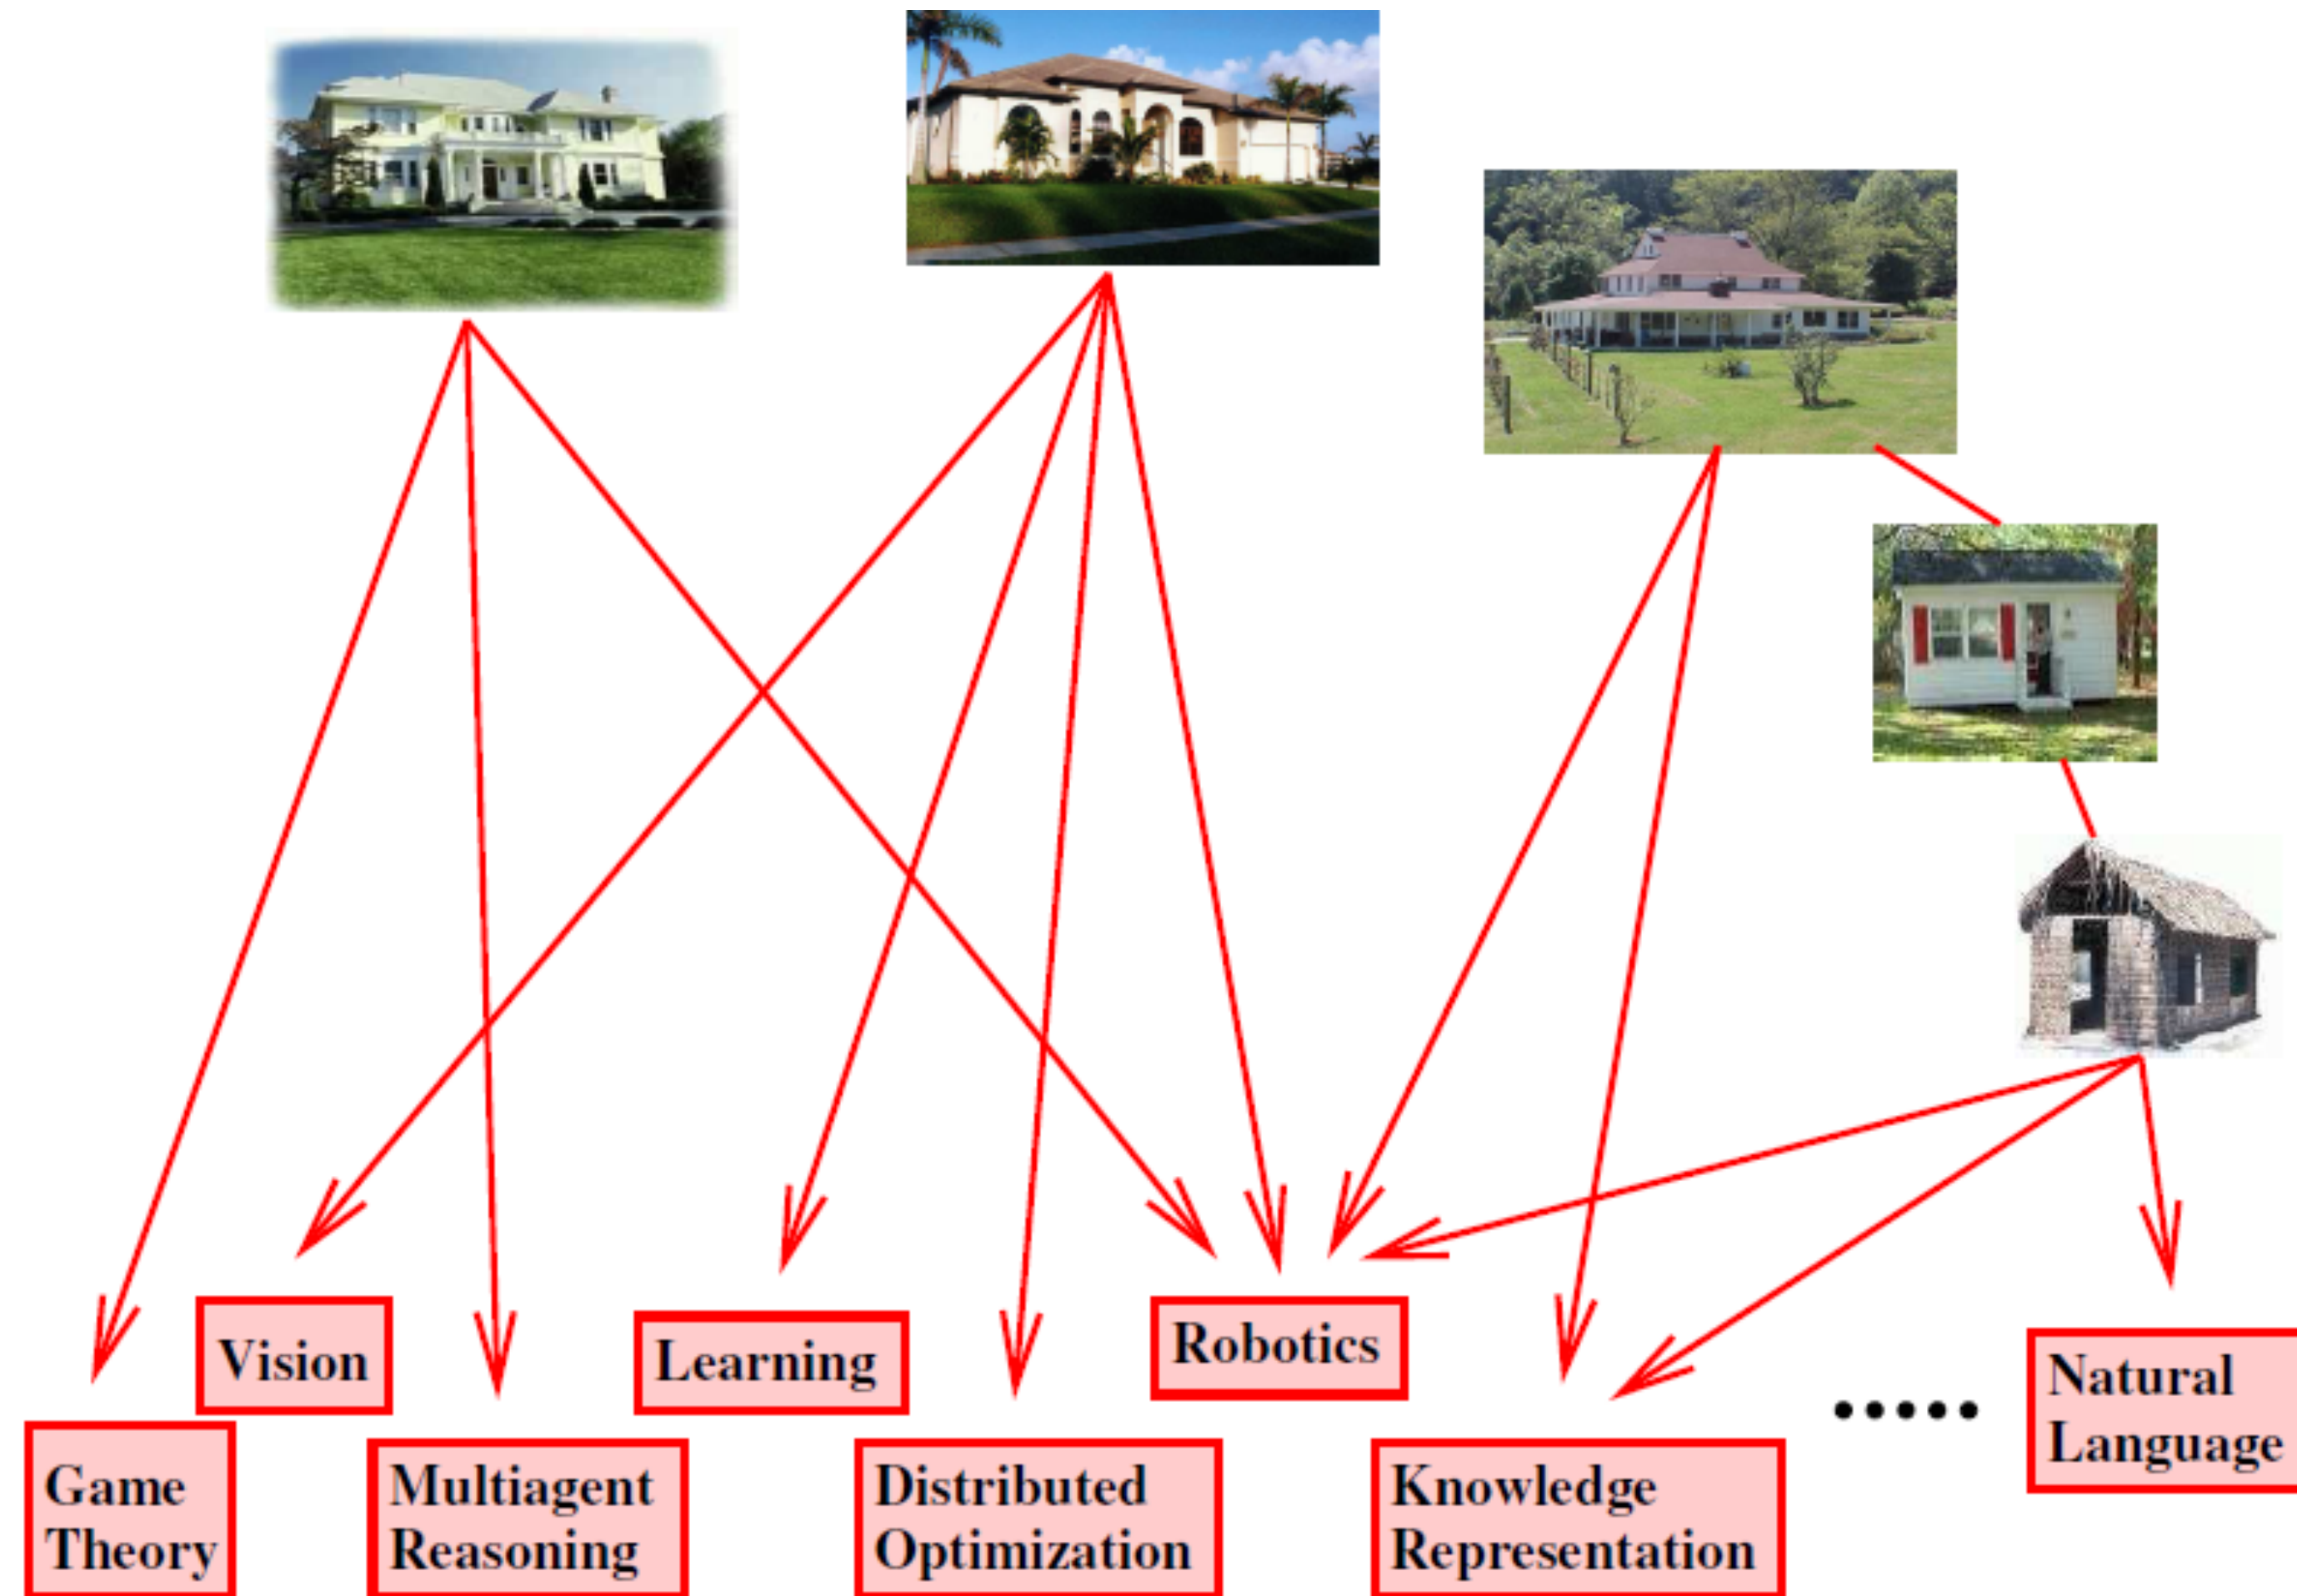
Towards a cathedral?



Or something else?



Top-down approach




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


Good problems produce good science



Manned flight



Apollo mission




Manhattan project



Autonomous vehicles

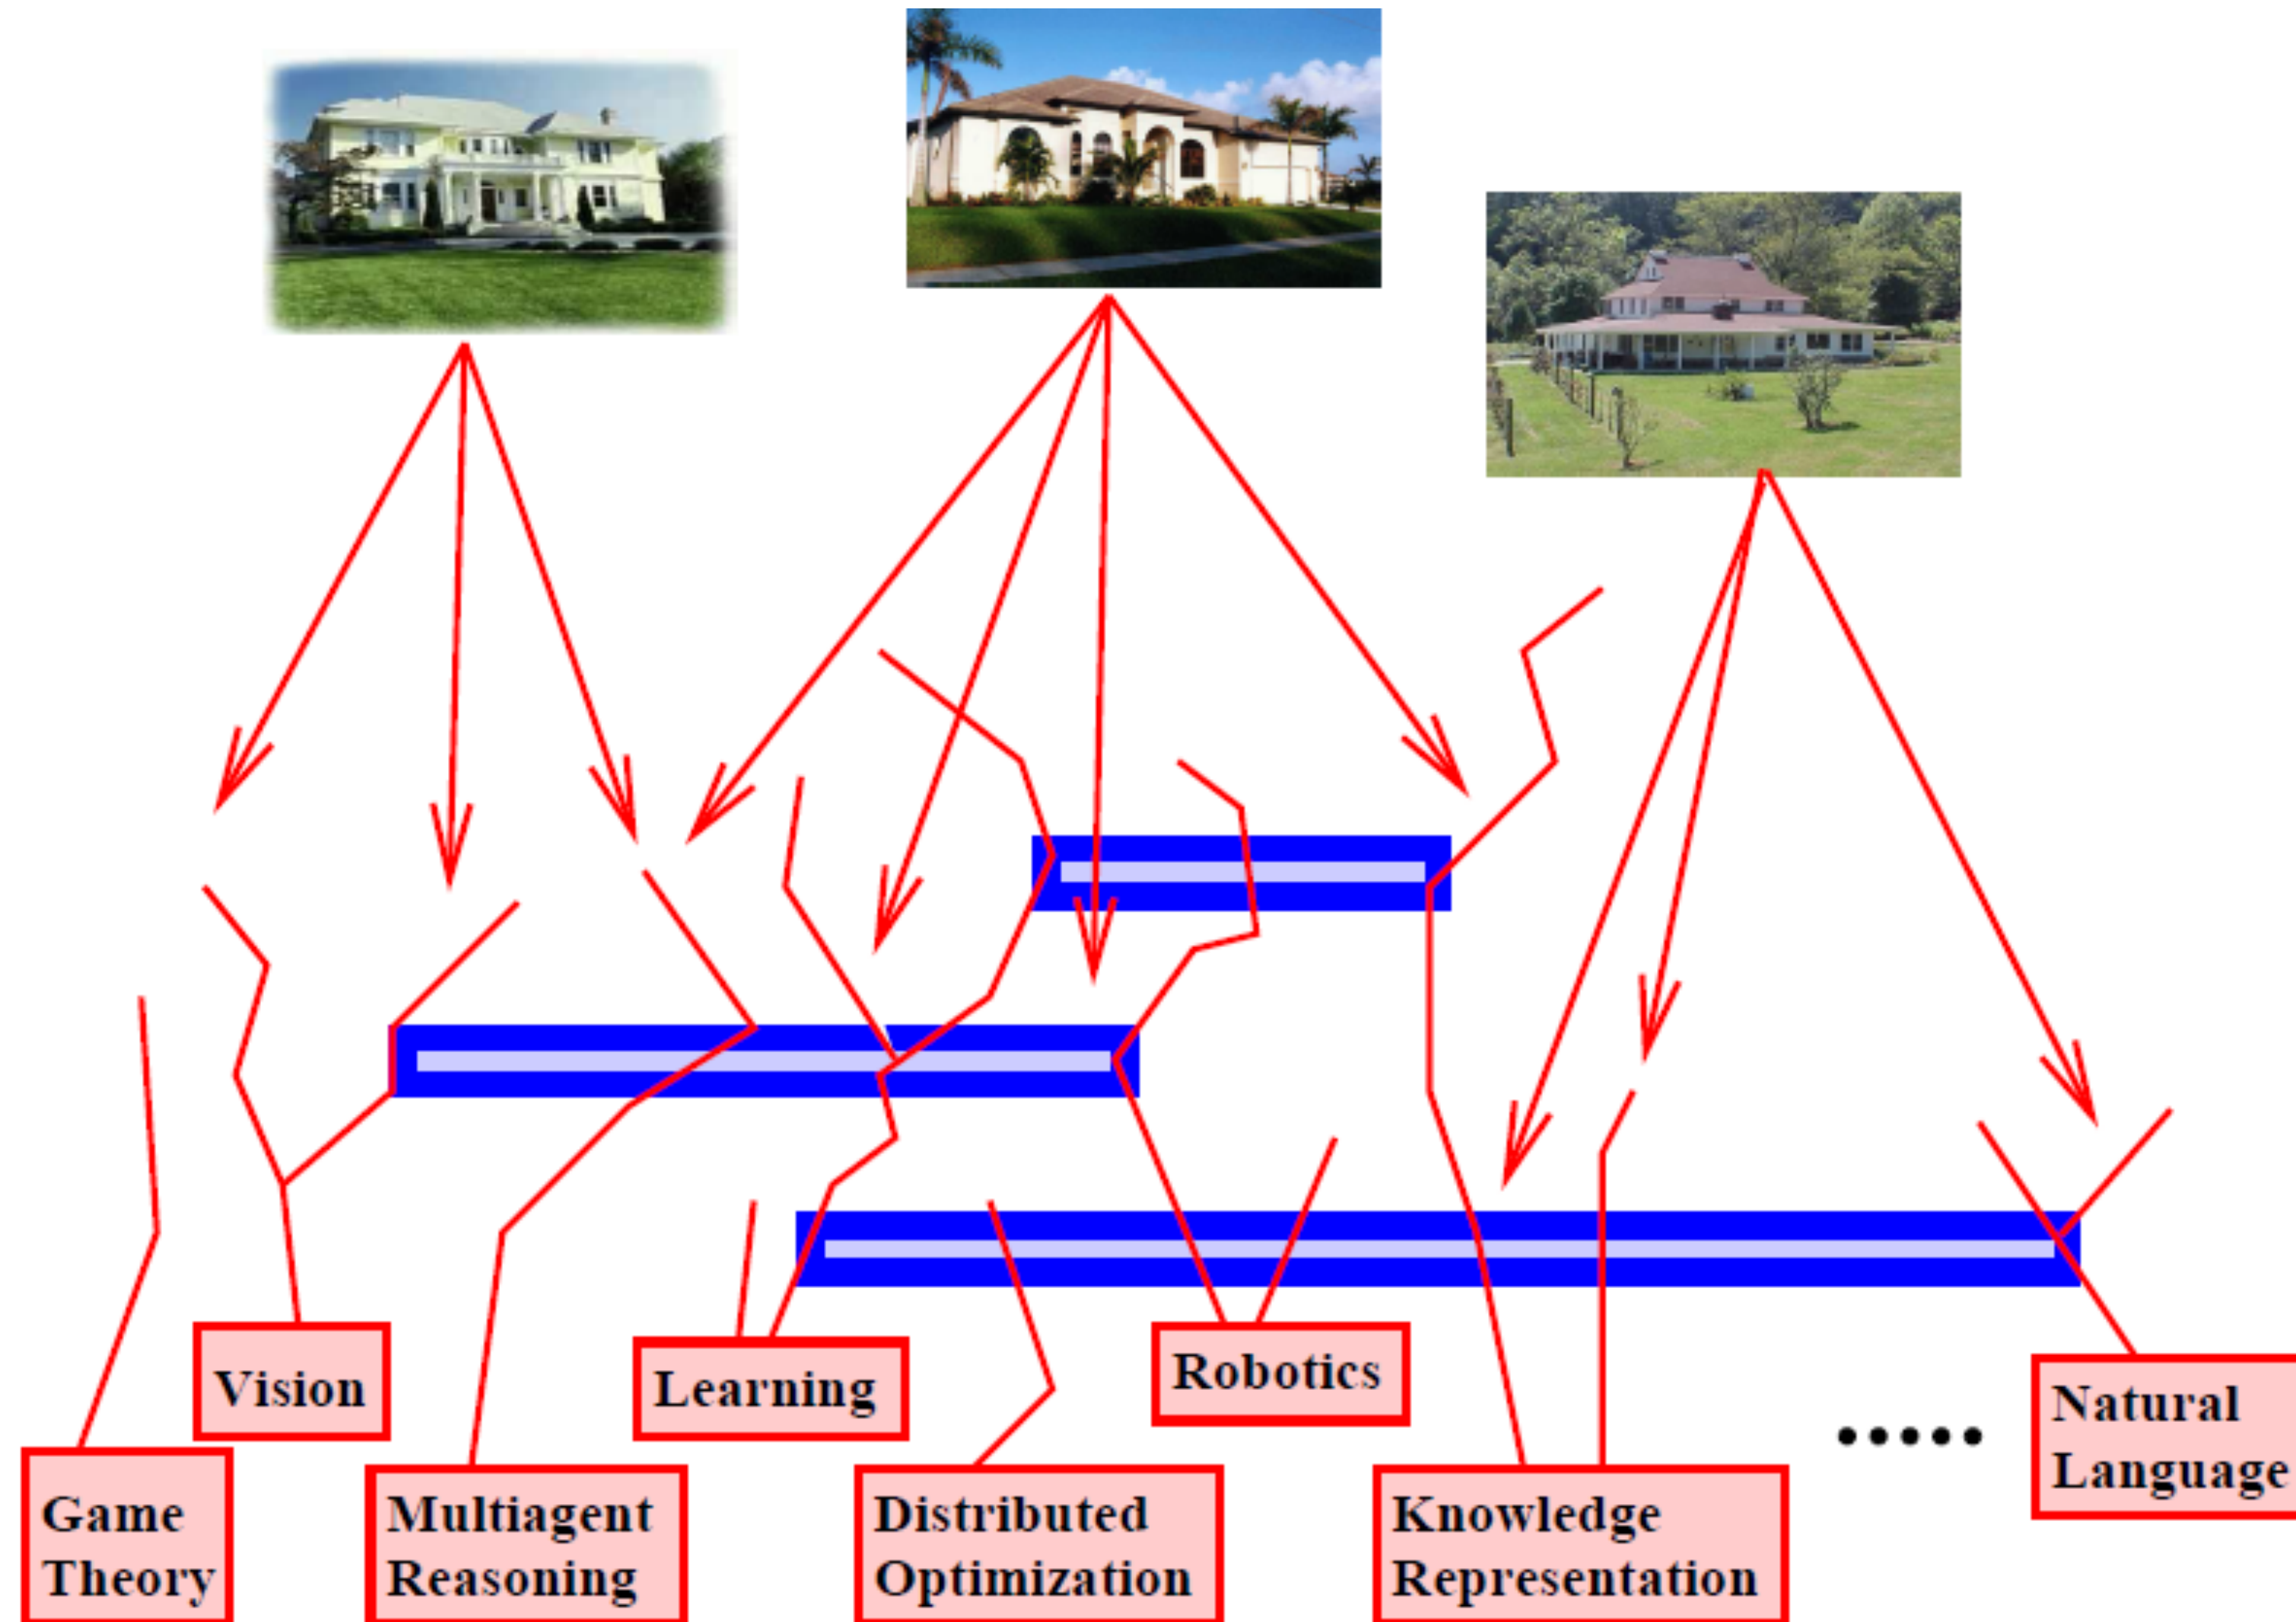


RoboCup soccer




Assistive robots

Meeting in the middle



Today


- What is this course?
- What is artificial intelligence?
- (What can AI do?)



Course Staff



Peter Stone
pstone@cs.utexas.edu



Yuke Zhu
yukez@cs.utexas.edu



Xiang Gu
xiangu@utexas.edu
Interests: Software
Engineering, Distributed
Computing



Soroush Nasiriany
soroush@cs.utexas.edu
Lab: RPL
Research Interest: Robot
Learning



Zizhao Wang
zizhao.wang@utexas.edu
Lab: LARG
Research Interest:
Robot Learning



Yifeng Zhu
yifeng.zhu@utexas.edu
Lab: LARG / RPL
Research Interest: Robot
Learning

Course Information

■ Communication:

- Announcements on webpage
- Grades on Canvas / edX
- Piazza for discussion

Class website:

http://rpl.cs.utexas.edu/cs343_spring2021

(or Google “Peter Stone” and go to the Teaching tab)

■ Course technology:

- edX for interactive homework (unlimited submissions!)
- Autograded programming projects (submit via Canvas)
- Make sure you have a CS Unix account IMMEDIATELY!
- Create an edX account


The screenshot displays the edX course interface for 'BerkeleyX: CS188x Artificial Intelligence - Berkeley (Spring 2014)'. The top navigation bar includes links for Courseware, Course Info (selected), Discussion, Wiki, Progress, Syllabus, Course Policies, Course Staff, Office Hours, Exams, and Student view. The main content area is divided into two columns. The left column, titled 'Course Updates & News', features a calendar icon and the date 'JANUARY 3, 2014', followed by the text 'Welcome to CS188, Spring 2014!'. The right column, titled 'Course Schedule', lists 'Self Diagnostic (ungraded)' and 'Project 0 1/24, 5pm'. The footer contains links for About, Jobs, Press, FAQ, and Contact, along with the edX logo and a brief description of the organization's mission. Social media icons for Weibo, Facebook, Twitter, Google+, and YouTube are also present, along with copyright information and links to Terms of Service and Privacy Policy.

Course Information

- Prerequisites:
 - Upper division standing
 - No formal class pre-reqs
 - There will be a lot of math (and programming)
- Coursework
 - Reading assignments with written responses
 - ~9 homework assignments:
 - ~2 weeks for each, but overlapping
 - Online, autograded, solve together, submit alone
 - No late submissions accepted
 - 6 programming projects
 - Python, groups of 1 or 2 (except Project 0)
 - ~2 weeks for each, non-overlapping
 - 5 late days for semester, (maximum 2 per project ?)
 - One midterm, one final
 - (Final Contest)

Textbook

Russell & Norvig, AI: A Modern Approach, 3rd Ed.



- After classes we'll post slides
- There will also be "step by step" videos posted for some topics

Warning: Not everything covered in the book will be covered in class (and to a small extent, vice versa). You are responsible for both.

Course Topics

- **Part I: Making Decisions**
 - Fast search / planning
 - Constraint satisfaction
 - Adversarial and uncertain search
 - MDPs and Reinforcement learning
- **Part II: Reasoning under Uncertainty**
 - Bayes nets
 - Decision theory and value of information
 - Statistical machine learning
- **Part III: Additional Topics**
 - Neural networks
 - Planning
 - Ethics
- **Throughout: Applications**
 - Natural language, vision, robotics, games, ...



Homework Exercises

- Online on edX
- Autograded text boxes / multiple choice
- Try as many times as you want!
 - (Might need to come to office hours
 - to unlock after a number of attempts)
- Goal: self-assess and prepare for tests
- Can discuss at high-level, but work alone
- No spoilers on Piazza discussions!


hw1_search_q4_a*_graph_search

[VIEW UNIT IN STUDIO](#)[Bookmark this page](#)

Q4: A* Graph Search

8.0 points possible (graded)

Consider A* *graph* search on the graph below. Arcs are labeled with action costs and states are labeled with heuristic values. Assume that ties are broken alphabetically (so a partial plan $S \rightarrow X \rightarrow A$ would be expanded before $S \rightarrow X \rightarrow B$ and $S \rightarrow A \rightarrow Z$ would be expanded before $S \rightarrow B \rightarrow A$).



In what order are states expanded by A* graph search? You may find it helpful to execute the search on scratch paper.

☐ Start, A, B, C, D, Goal

☐ Start, A, C, Goal

☐ Start, B, A, D, C, Goal


☐ Start, A, D, Goal

☐ Start, A, B, Goal

☐ Start, B, A, D, B, C, Goal

Programming Assignments

Pacman domain



Projects include:

- path planning and search
- multi-agent game trees
- reinforcement learning
- state estimation
- classification
- final contest

Highly suggested: Pair programming
(switch “driver” and “observer” roles often)

Midterm and Final

- Midterm will cover roughly half the class material
- Final will be comprehensive
- (When it was in person: One page of notes, but not open book)

Syllabus

Official syllabus is online

Grades will be weighted as follows:

- Written responses to readings (5%)
- Class attendance and participation (10%)
- Homework Exercises (20%)
- Programming Assignments (25%)
- Midterm (15%)
- Final (25%)

Academic Honesty

READ THE STATEMENT IN THE SYLLABUS

- Discuss concepts, but don't share solutions or written work with other students
- Don't look for answers / code online or elsewhere
- Automated tools will be used to discover cheating
- If unsure, check departmental guidelines or ask — ignorance is not an excuse
- We will pursue the harshest penalties available, so please don't cheat!
- To be clear: you will fail the class automatically and be reported to the university

Important This Week


- Important this week:
 - **Read and respond** to AI 100 report (past due, but no late penalty)
 - **Read** the syllabus
 - **Reading assignment and email response** by Wednesday at 9:30am
 - **Register** for the class on edX (Click on link for HW1 to register for class after making account)
 - **Enroll and post something** on Piazza
 - **Be sure** that you have a usable CS Unix account - <https://apps.cs.utexas.edu/udb/newaccount/>
 - **P0: Python tutorial** is out (due on Thursday 1/21 at 9:30am via Canvas)
- Also important:
 - **If you are wait-listed**, you might or might not get in depending on how many students drop. Be patient if possible — many students often drop early in the course.
 - **Office Hours** see website

How to Ace this Course

- Do the readings! – That's your primary source for the material
- Do all the assignments on time – That's a big chunk of the grade
- Prepare for the exams – Historically, have been difficult for some
 - Come to class for practice
 - Keep up with the material – don't cram
- Become known to the class staff – for class participation
 - Come to class (with your camera on!) and be ready to interact
 - Reading responses
 - Piazza posts

Today


- What is this course?
- What is artificial intelligence?
- (What can AI do?)



But First.... Implications

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

What happens when we achieve this goal?



?



?

Is AI moving us in the right direction?

Discussion Question: Would you rather have been born:
– 50 years earlier? – 50 years later?

Other difficult questions...


- Who is liable if a robot driver has an accident?
- Will machines surpass human intelligence (in all ways)?
- Would such machines have conscious existence? Rights?
- What is a mind?
- How can a physical object have a mind?
- Can we build a mind?
-

**AI is one of the great intellectual adventures
of the 20th and 21st centuries!**


What is AI?

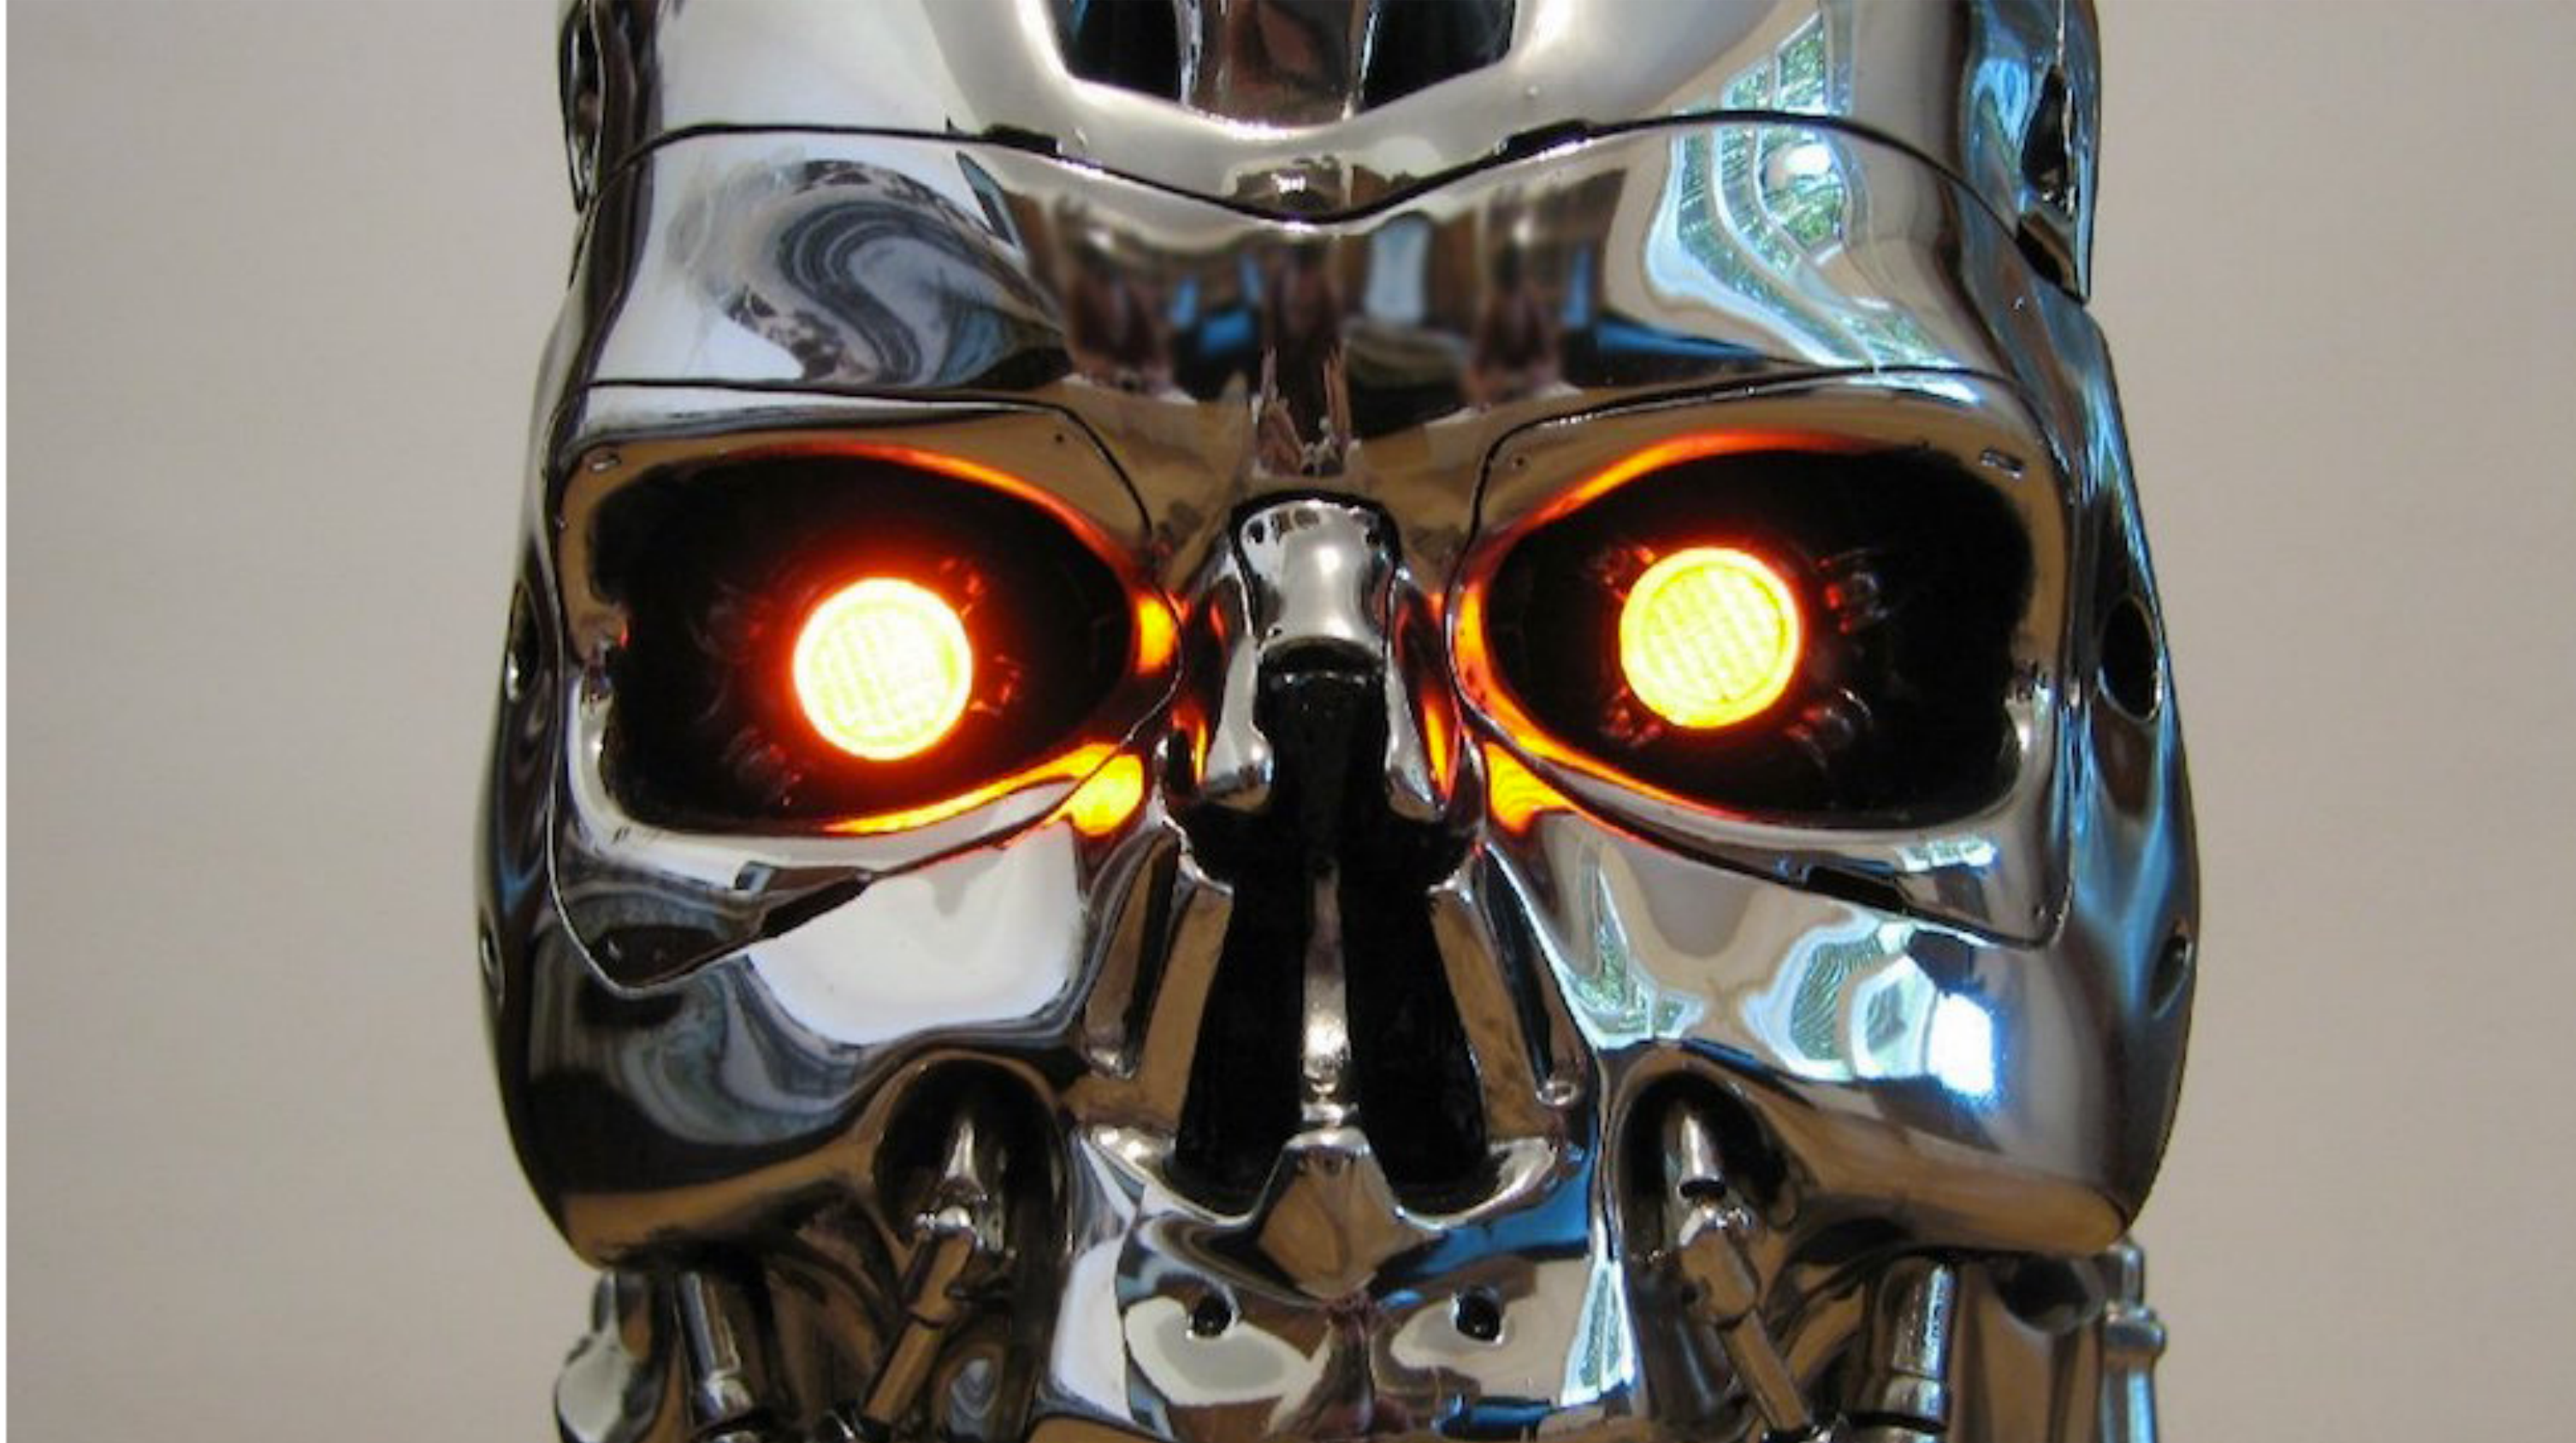
- Finding fast algorithms for NP-hard problems
- Getting computers to do the things they can't do yet
- Getting computers to do the things they do in the movies
-

Sci-Fi AI



AI in the news

**Nerdist** [PODCASTS](#) [VIDEOS](#) [ARTICLES](#) [EVENTS](#) [TWITCH](#) [STORE](#) [Search](#) [User](#)




ELON MUSK ANNOUNCES OPEN AI: A NON-PROFIT TO HELP PREVENT SKYNET

Advertisement


MOST POPULAR ARTICLES

AI Hype


**Realit
y**




Uncertainty



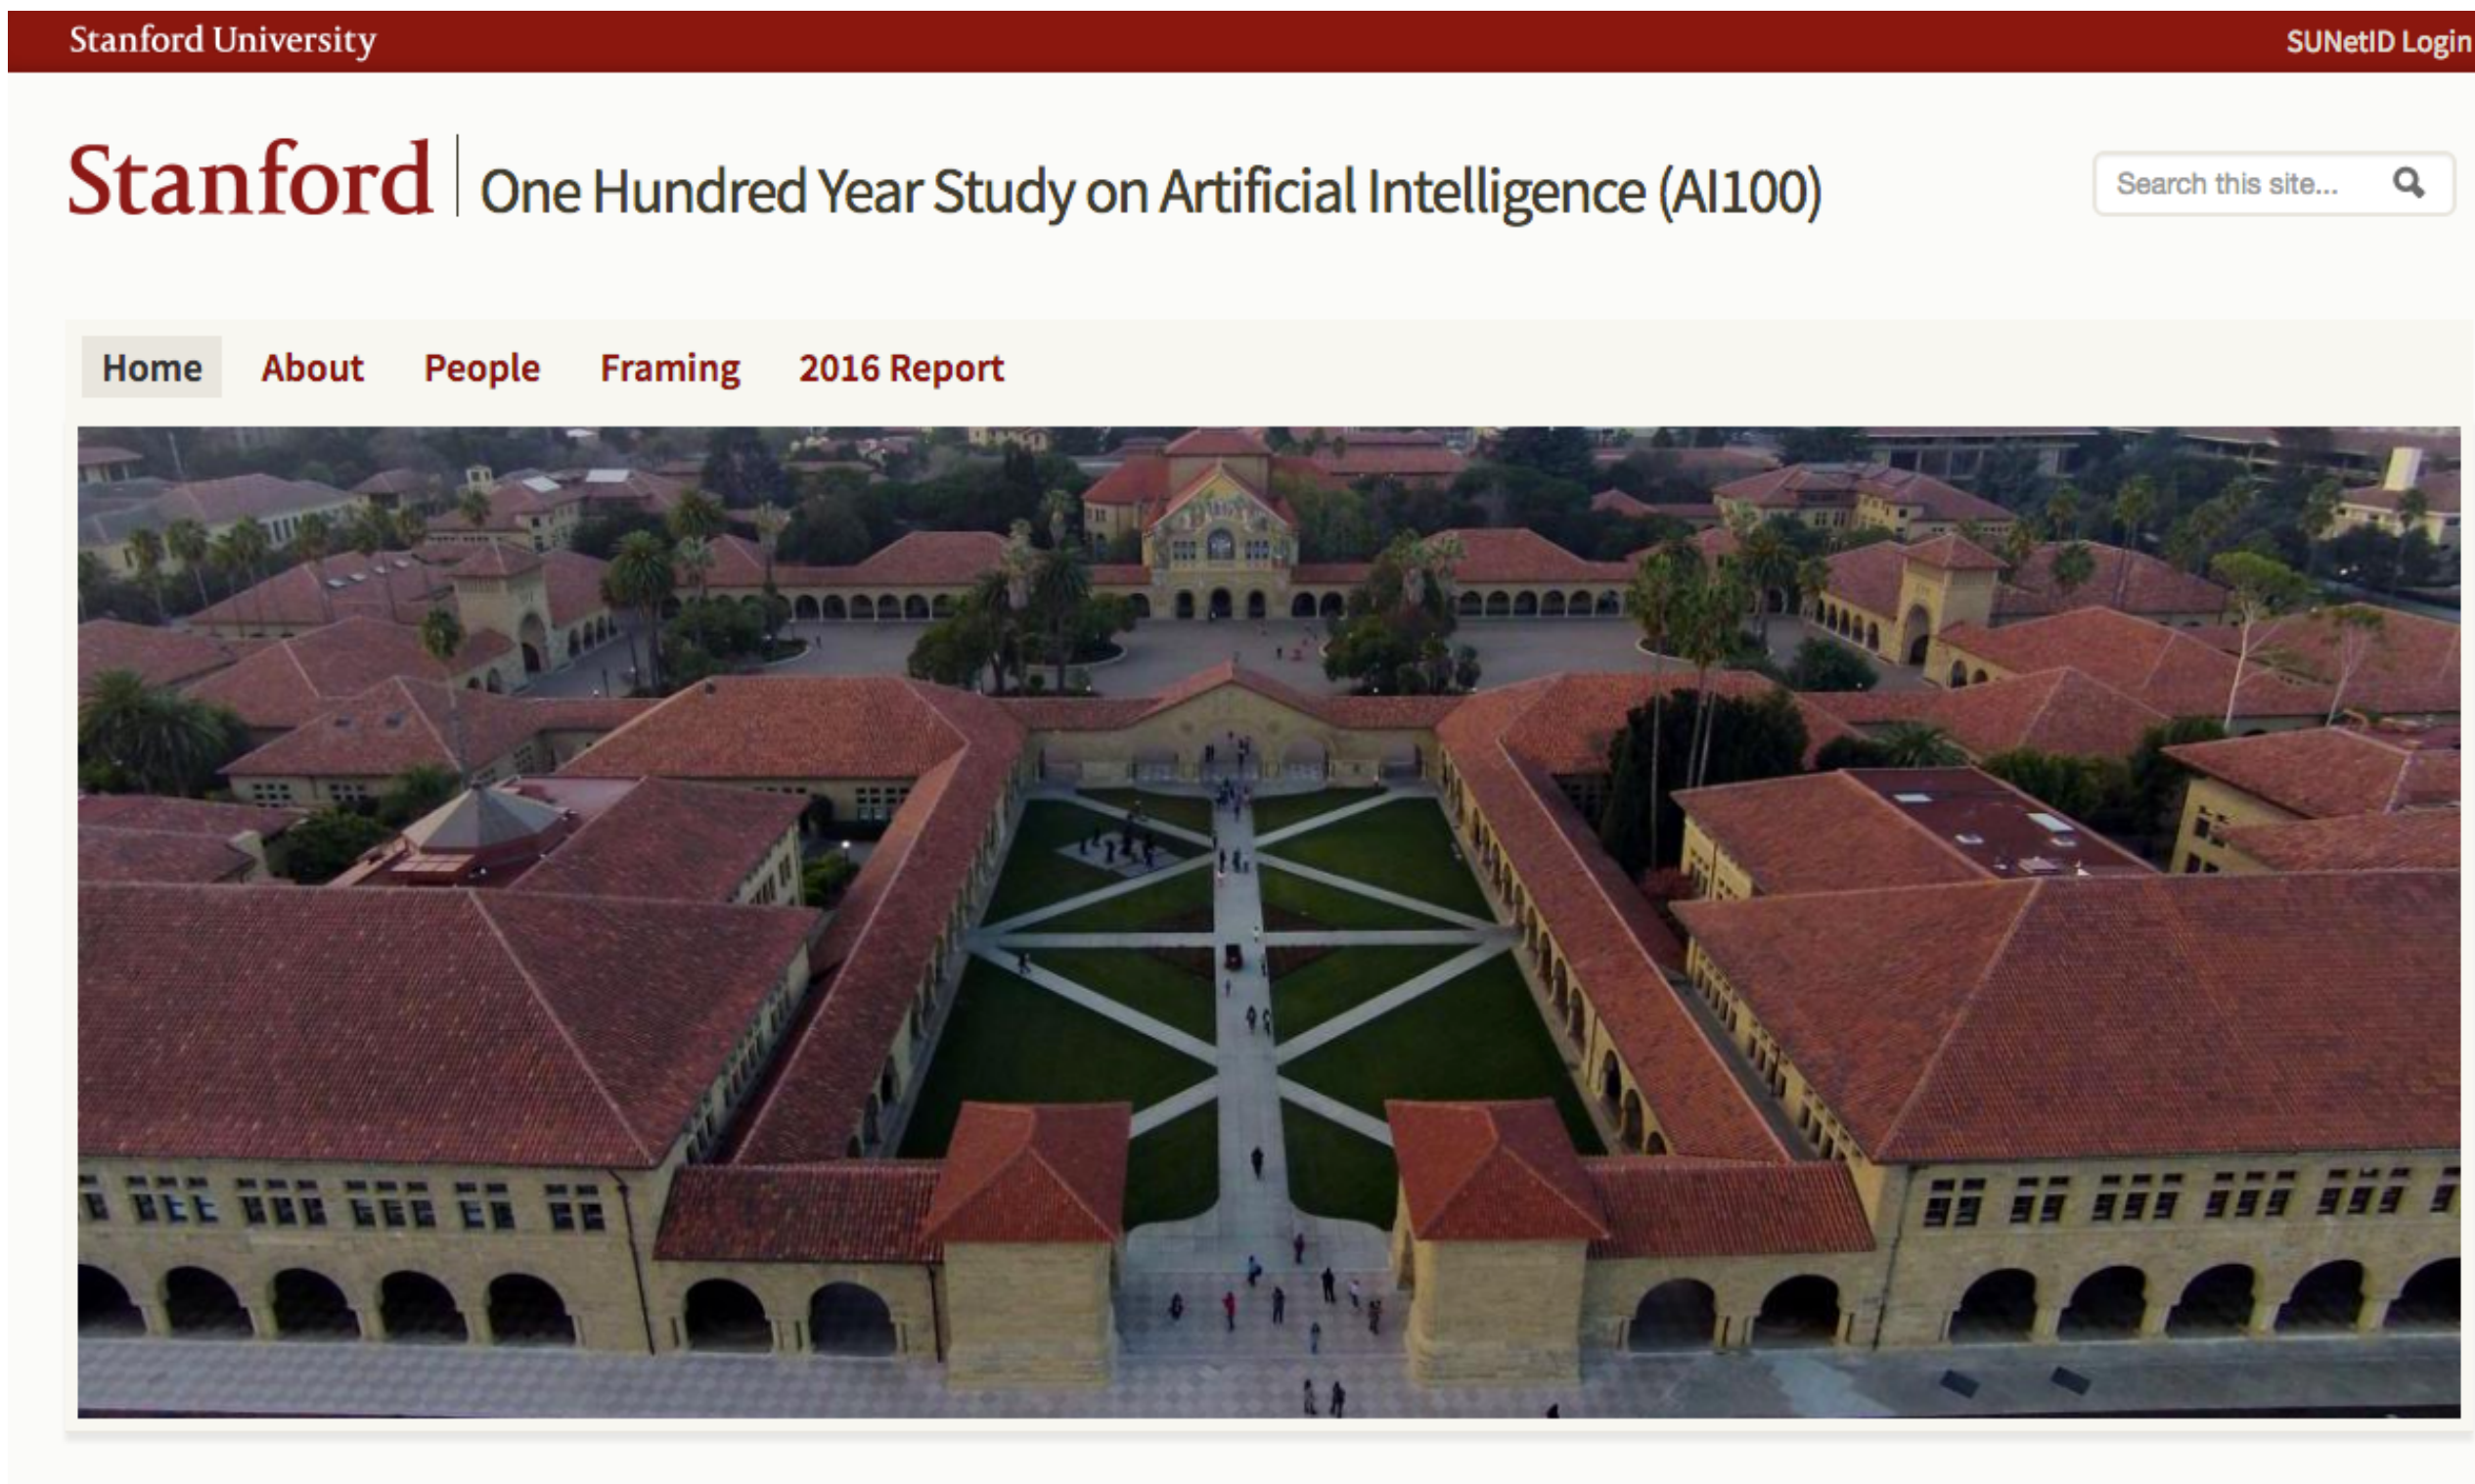
Perceptions



**Perception
Uncertainty**



A definition of AI



“Artificial Intelligence (AI) is a science and a set of computational technologies that are inspired by — but typically operate quite differently from — the ways people use their nervous systems and bodies to sense, learn, reason, and take action.”

Your Reactions

- **Self-driving cars not yet widely adopted (Madeline Huang)**
- **Bias may be unavoidable in data and usage (Winnie Chang)**
- **No mention of implicit bias leaking into systems (Esther Yeum)**
- **No oath for software engineers (Ramya Prasad)**
- **Policy recommendations not actionable (Shrivastava Shankar)**
- **Musk/Hawking predictions of AI getting out of control vs.
“...there is no race of superhuman robots on the horizon or
probably even possible.” (Daniel Deng)**


What is AI?

The science of making machines that:

What is AI?


The science of making machines that:

Think like people



Thinking Like Humans?

- The cognitive science approach:
 - 1960s ``cognitive revolution": information-processing psychology replaced prevailing orthodoxy of behaviorism (reflexive behaviors, classical conditioning, etc.)
- Scientific theories of internal activities of the brain
 - What level of abstraction? "Knowledge" or "circuits"?
 - **Cognitive science:** Predicting and testing behavior of human subjects (top-down)
 - **Cognitive neuroscience:** Direct identification from neurological data (bottom-up)
 - Both approaches now distinct from AI
 - Both share with AI the following characteristic:
The available theories do not explain (or engender) anything resembling human-level general intelligence




Images from Oxford fMRI center

What is AI?

The science of making machines that:


Think like people




What is AI?

The science of making machines that:

Think like people




Act like people



Acting Like Humans?

- Turing (1950) “Computing machinery and intelligence”
 - “Can machines think?” → “Can machines behave intelligently?”
 - Operational test for intelligent behavior: the *Imitation Game*




- Predicted by 2000, a 30% chance of fooling a lay person for 5 minutes
- Anticipated all major arguments against AI in following 50 years
- Suggested major components of AI: knowledge, reasoning, language understanding, learning
- Problem: Turing test is not reproducible or amenable to mathematical analysis


What is AI?

The science of making machines that:

Think like people




Act like people




What is AI?

The science of making machines that:


Think like people



Think rationally



Act like people



Thinking Rationally?


- The “Laws of Thought” approach
 - What does it mean to “think rationally”?
 - Normative / prescriptive rather than descriptive
- Logicist tradition:
 - Logic: notation and rules of derivation for thoughts
 - Aristotle: what are correct arguments/thought processes?
 - Direct line through mathematics, philosophy, to modern AI
- Problems:
 - Not all intelligent behavior is mediated by logical deliberation
 - What is the purpose of thinking? What thoughts should I (bother to) have?
 - **Logical systems tend to do the wrong thing in the presence of uncertainty**




What is AI?

The science of making machines that:


Think like people



Think rationally




Act like people




What is AI?

The science of making machines that:


Think like people




Think rationally



Act like people



Act rationally



Acting Rationally

- Rational behavior: doing the “right thing”
 - The right thing: that which is expected to maximize goal achievement, given the available information
 - Doesn't necessarily involve thinking, e.g., blinking
 - Thinking can be in the service of rational action
 - Entirely dependent on goals!
 - Irrational \neq insane, irrationality is sub-optimal action
 - Rational \neq successful
- Our focus here: rational agents
 - Systems which make the best possible decisions given goals, evidence, and constraints
 - In the real world, usually lots of uncertainty
 - ... and lots of complexity
 - Usually, we're just approximating rationality

Rational Decisions

We'll use the term **rational** in a very specific, technical way:

- Rational: maximally achieving pre-defined goals
- Rationality only concerns what decisions are made
(not the thought process behind them)
- Goals are expressed in terms of the **utility** of outcomes
- Being rational means **maximizing your expected utility**

Another title for this course could be:

Computational Rationality

Maximize Your Expected Utility

