


CS343: Artificial Intelligence

Introduction



Prof. 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 Yuke Zhu at UT Austin.]

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

Welcome to CS 343!

- **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

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

Goals of AI

- Building robust, fully autonomous agents in the real world
- Understanding human intelligence from a computational perspective


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"?




AI Hype


Reality



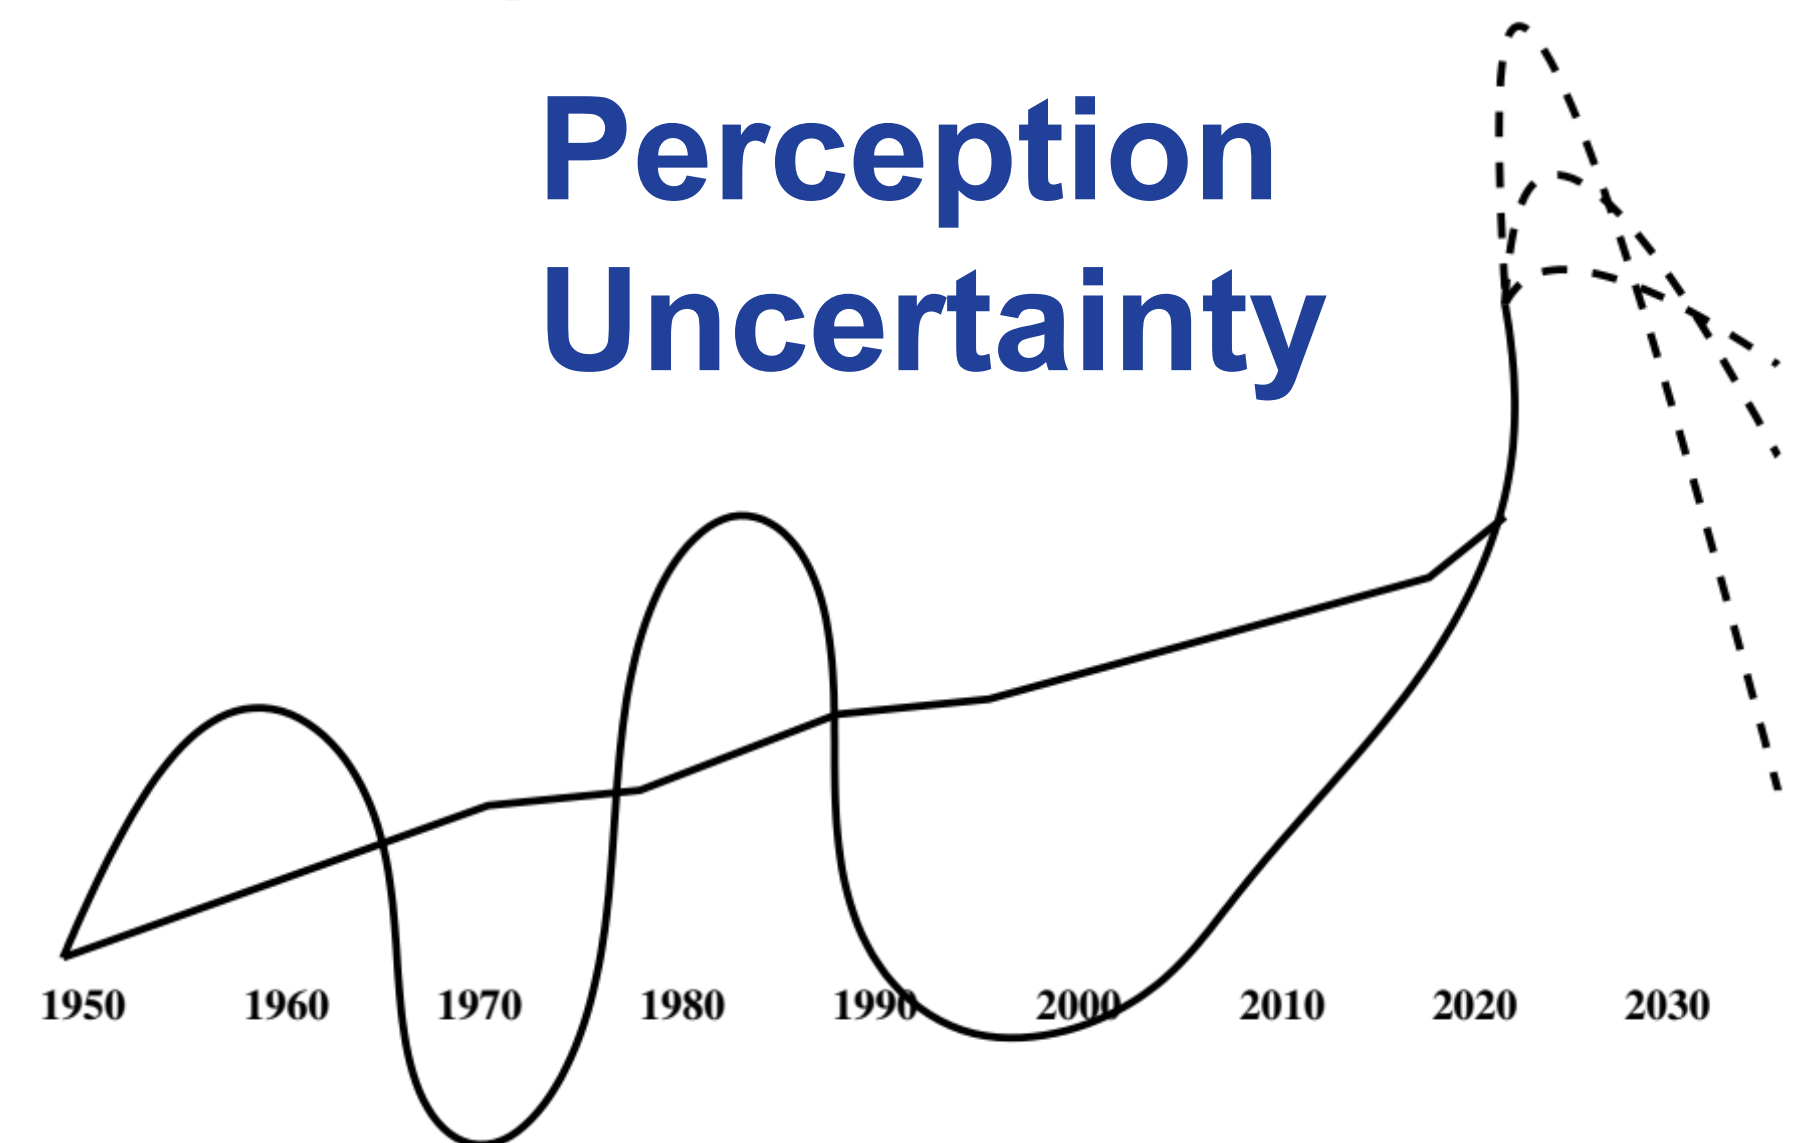
Perceptions



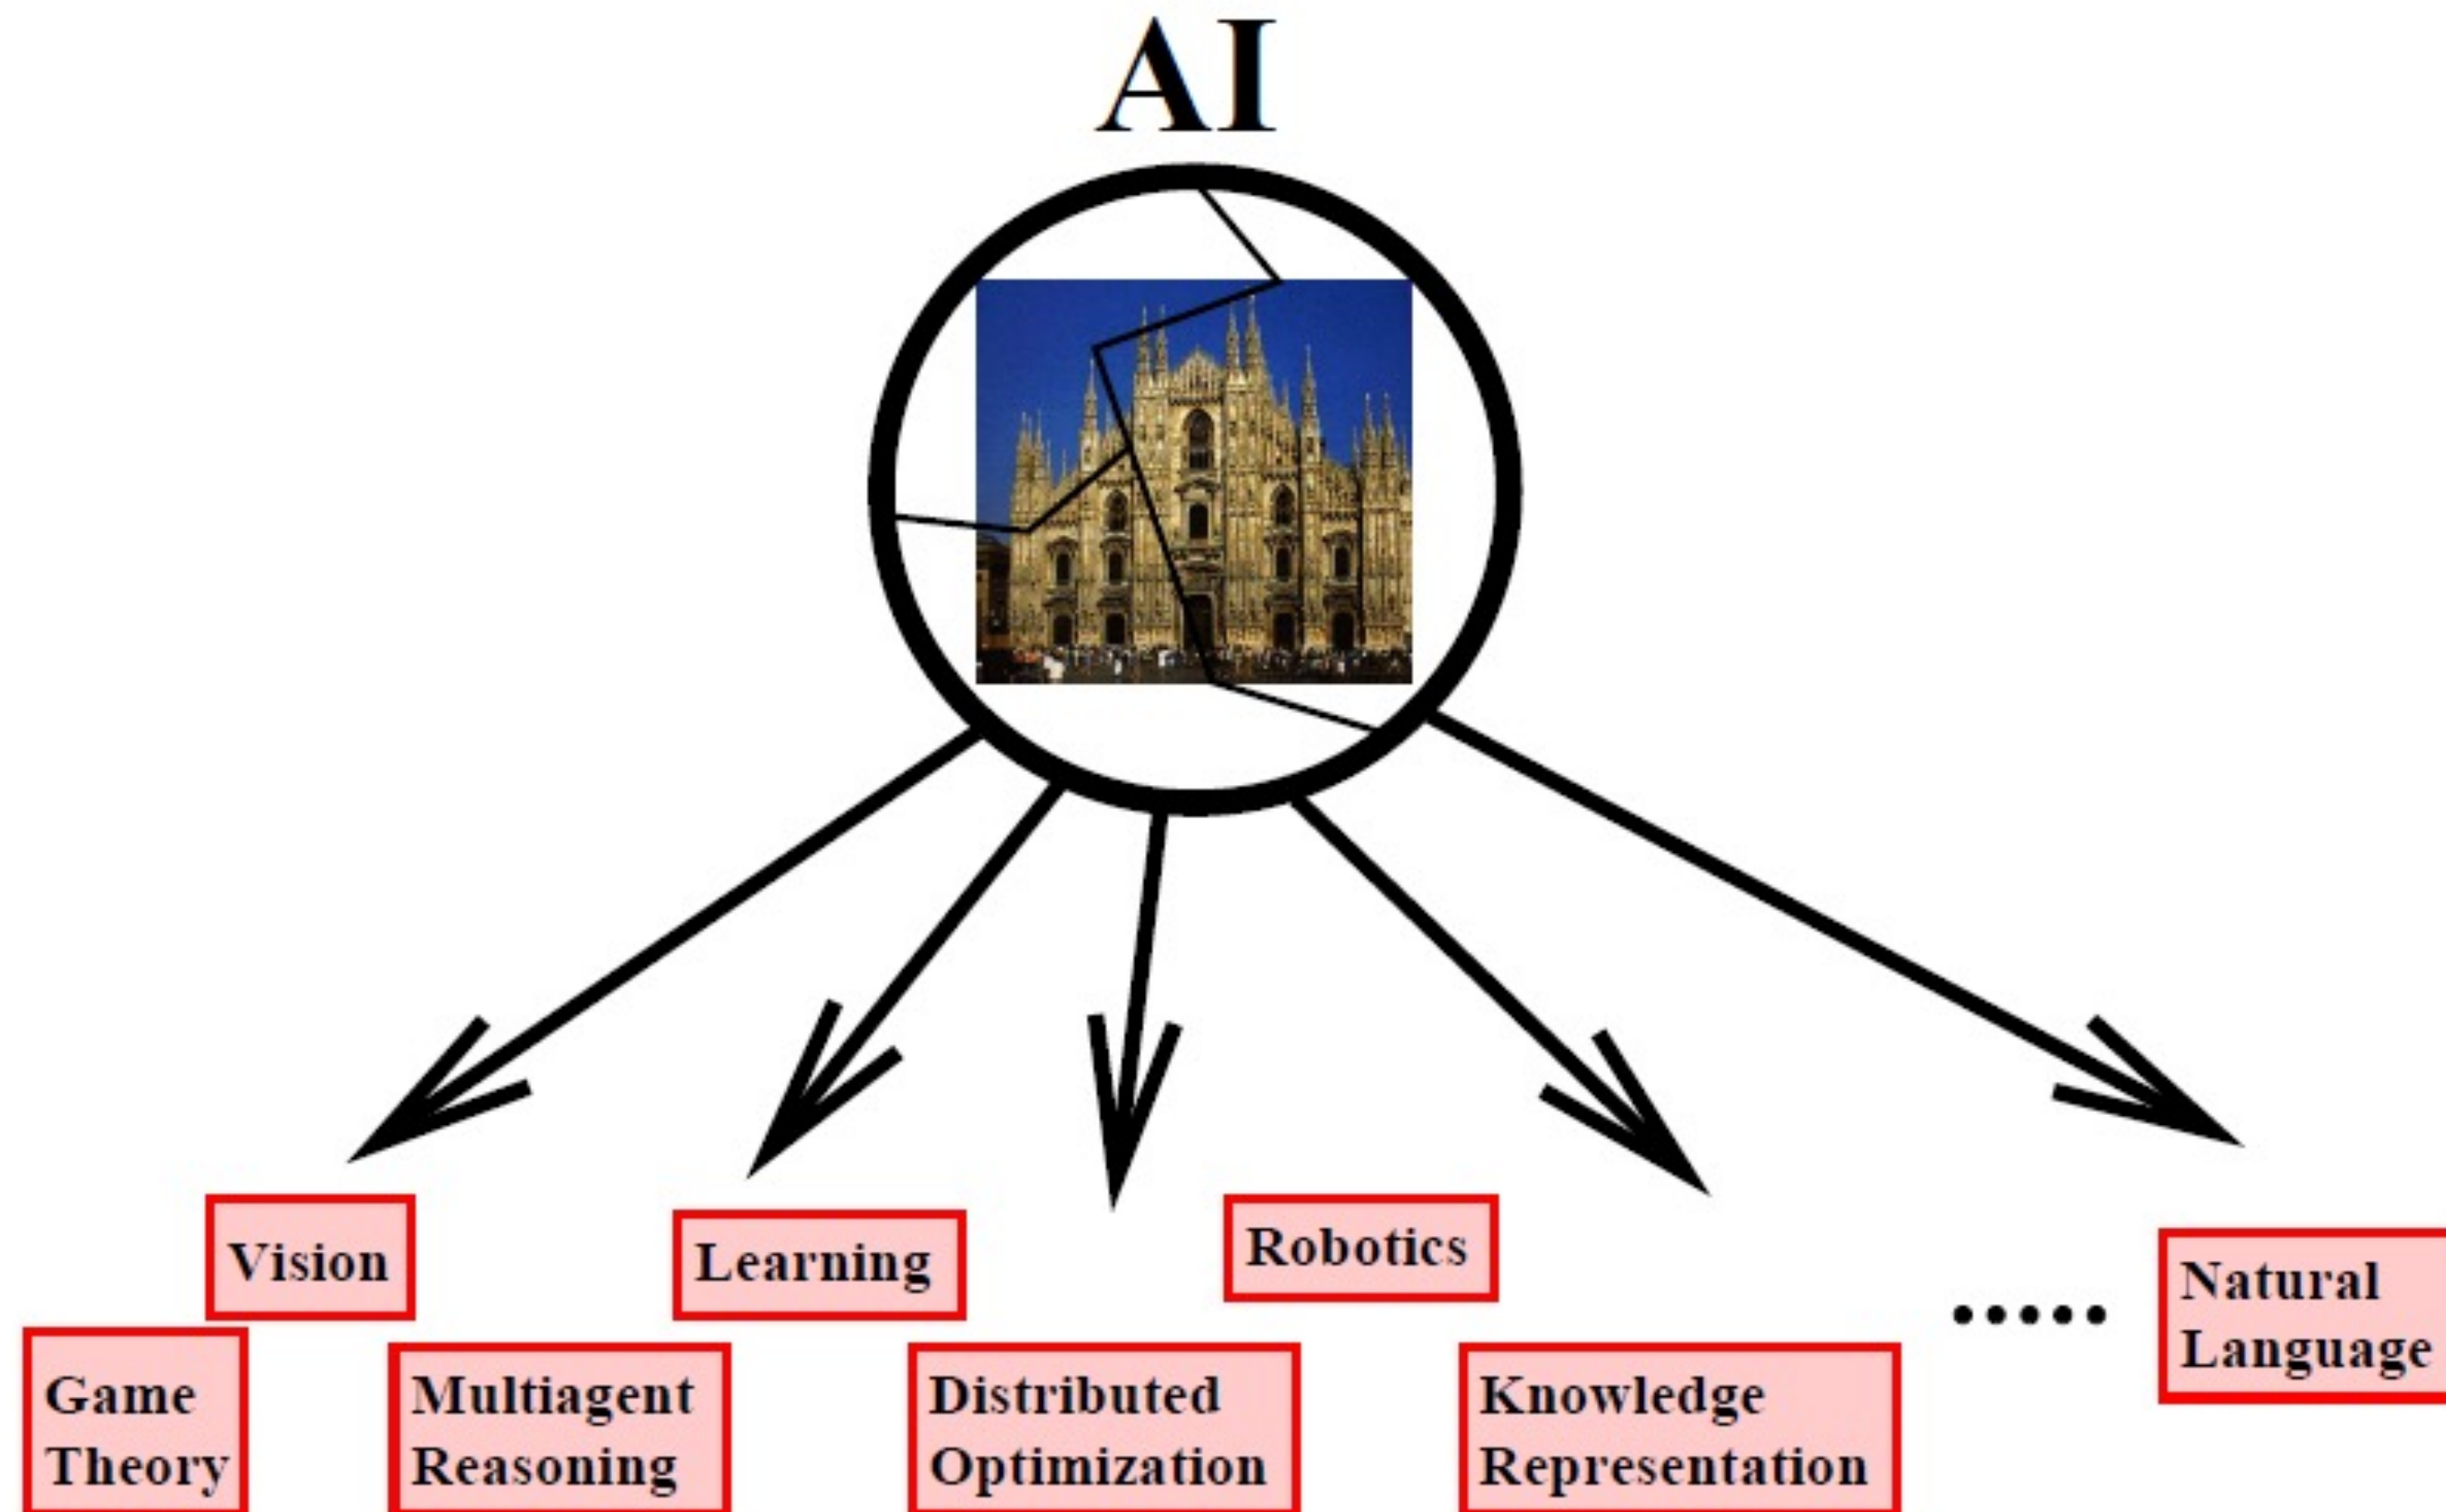
Uncertainty



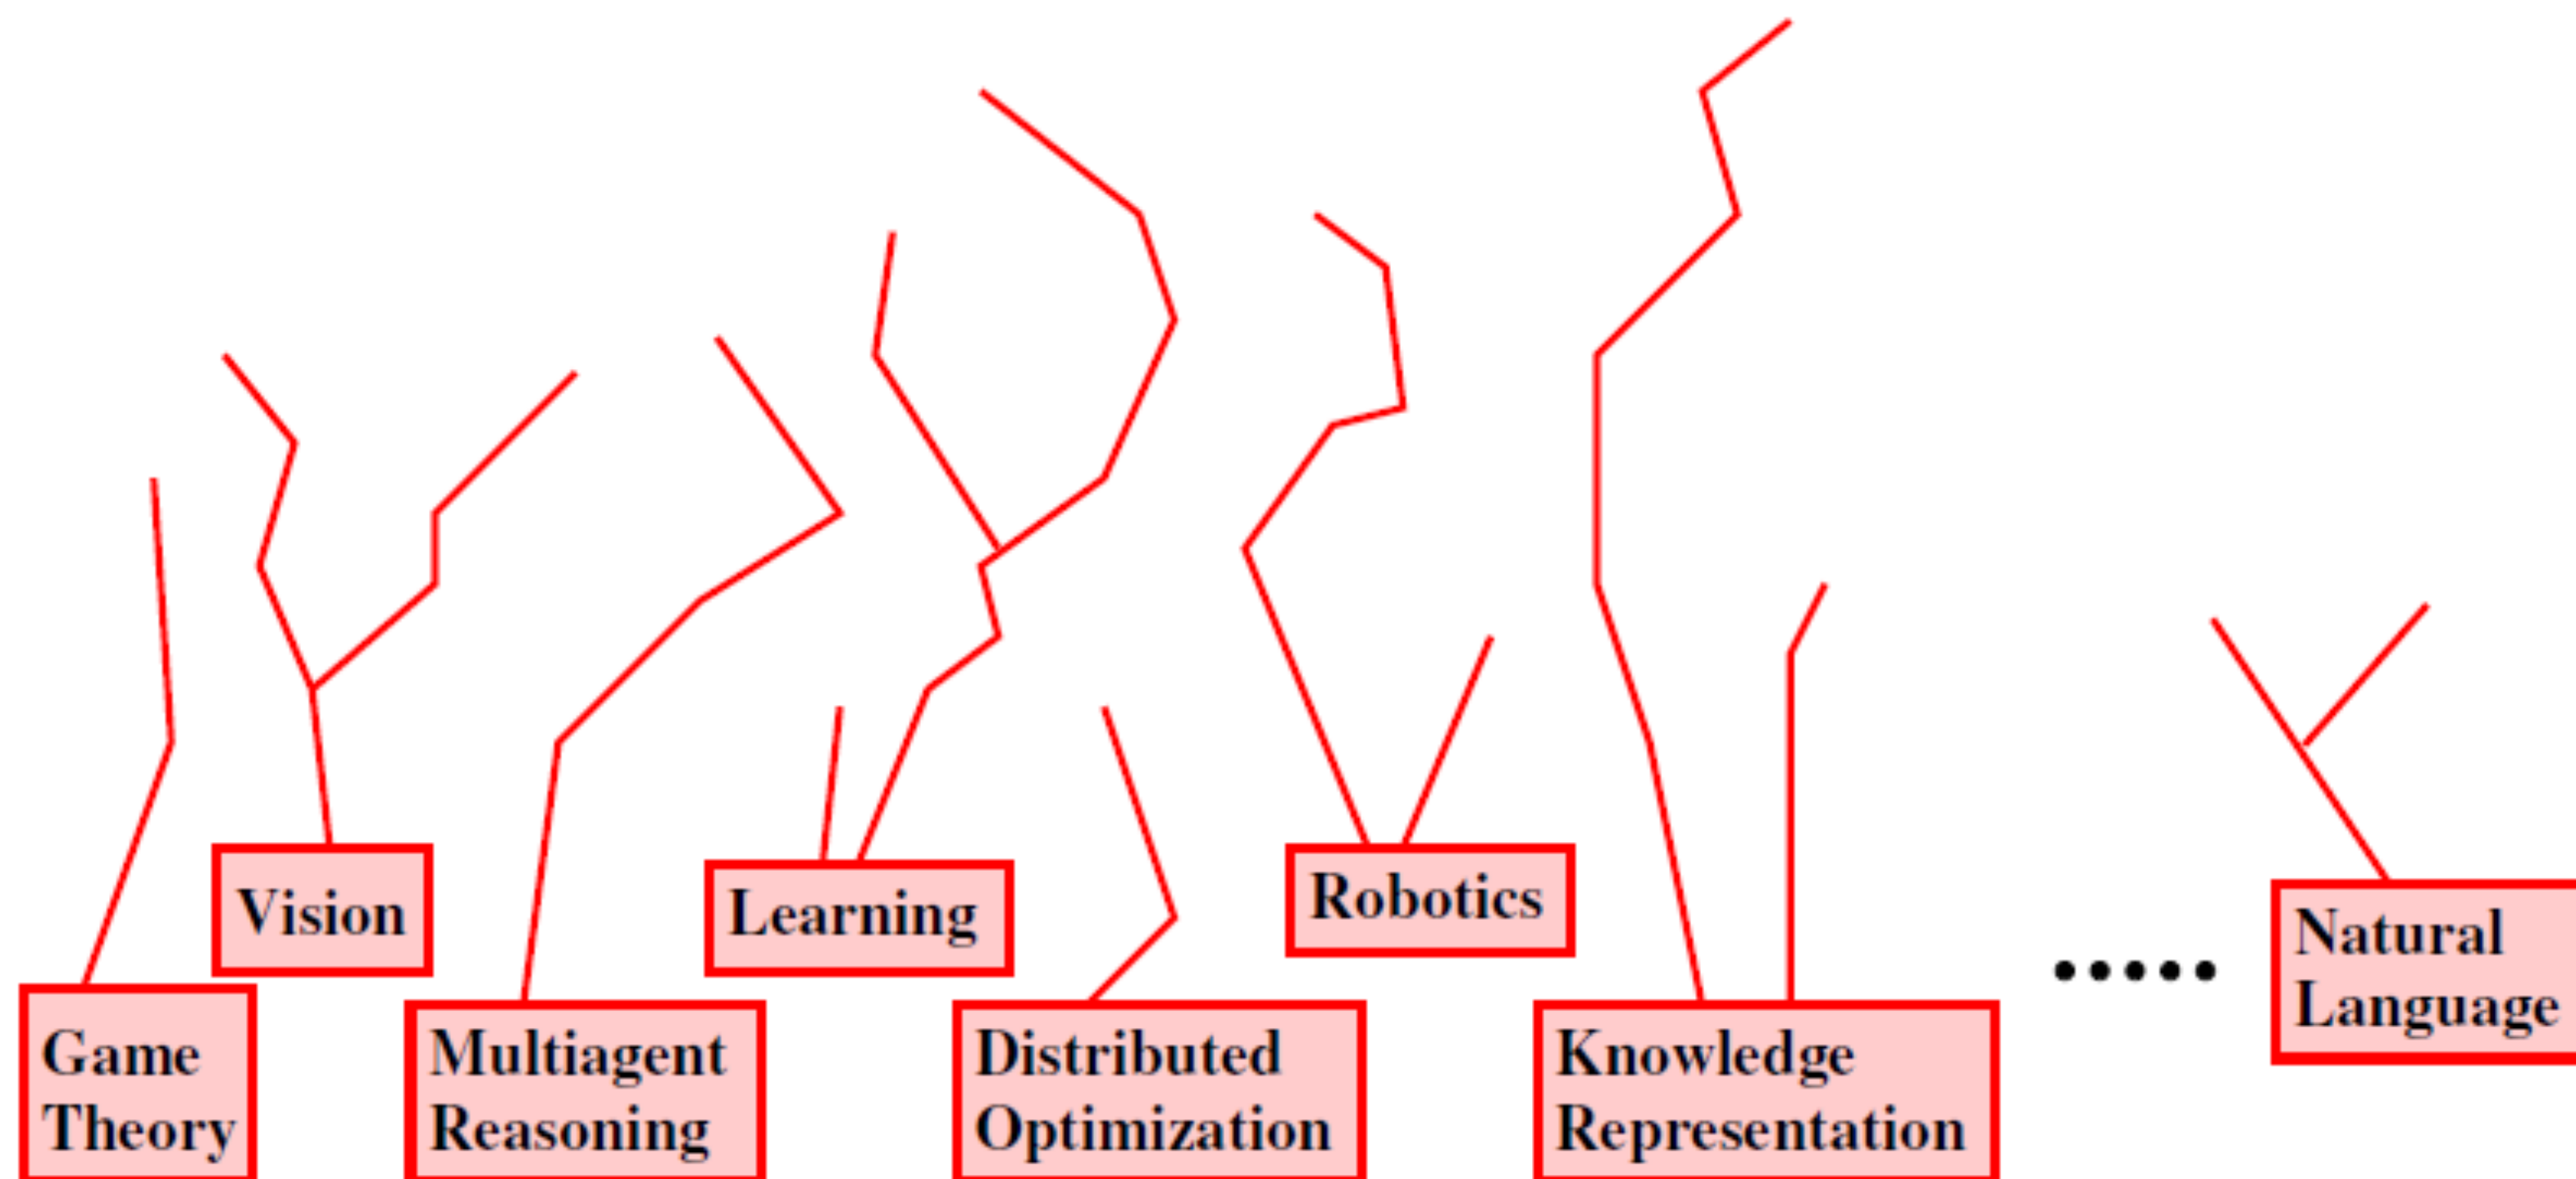
Perception Uncertainty



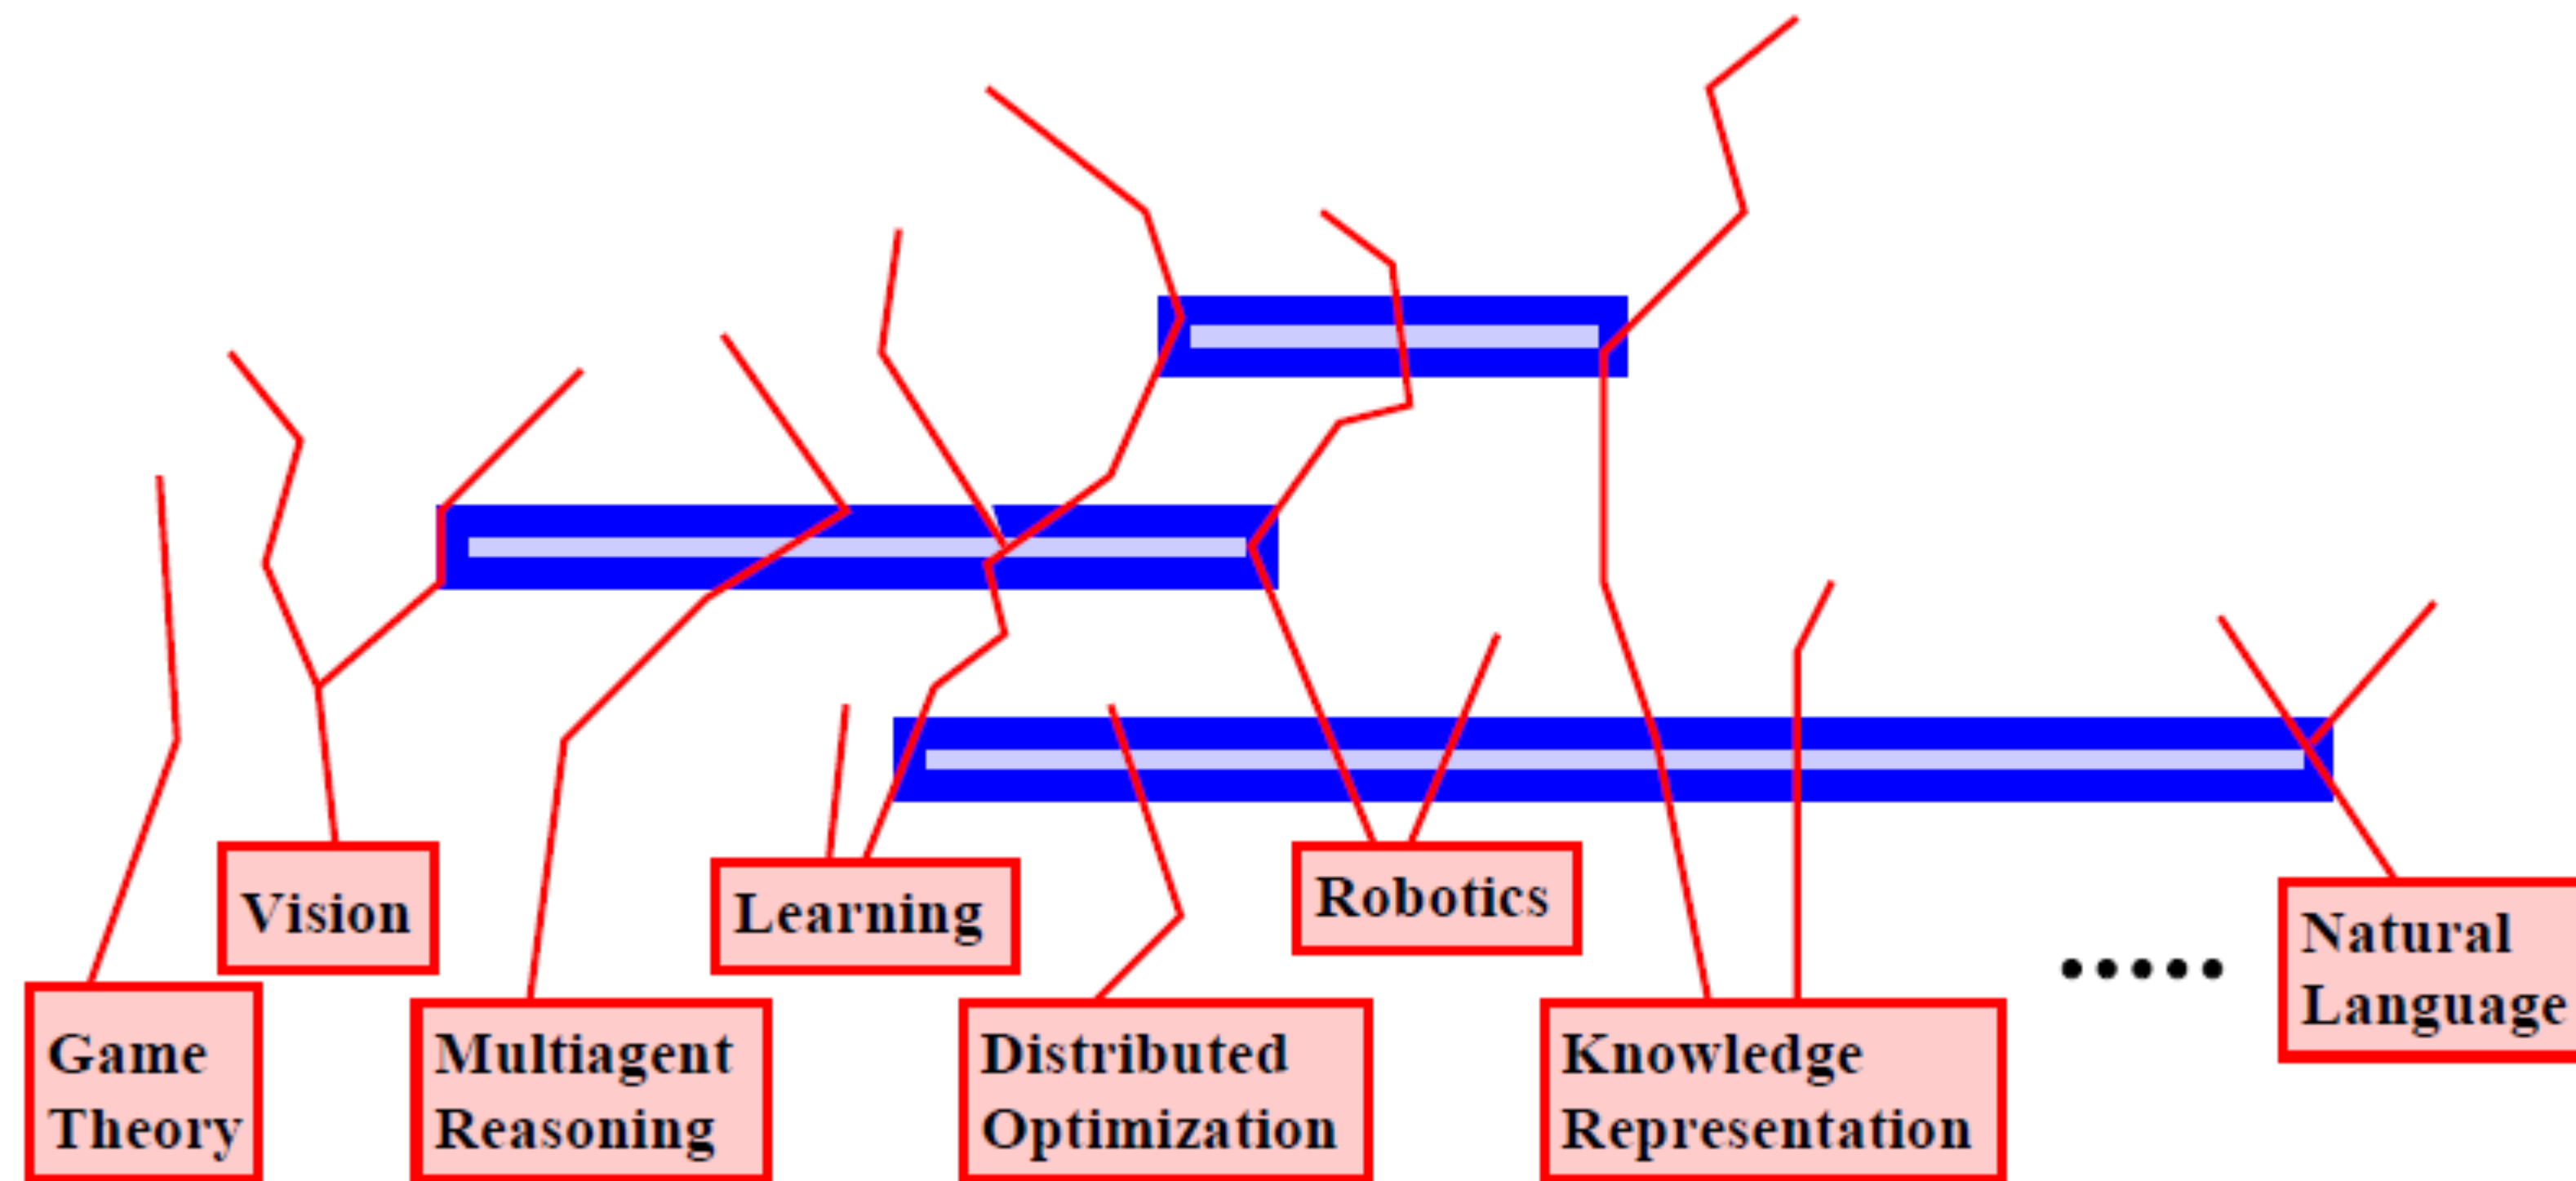
Bottom-up approach



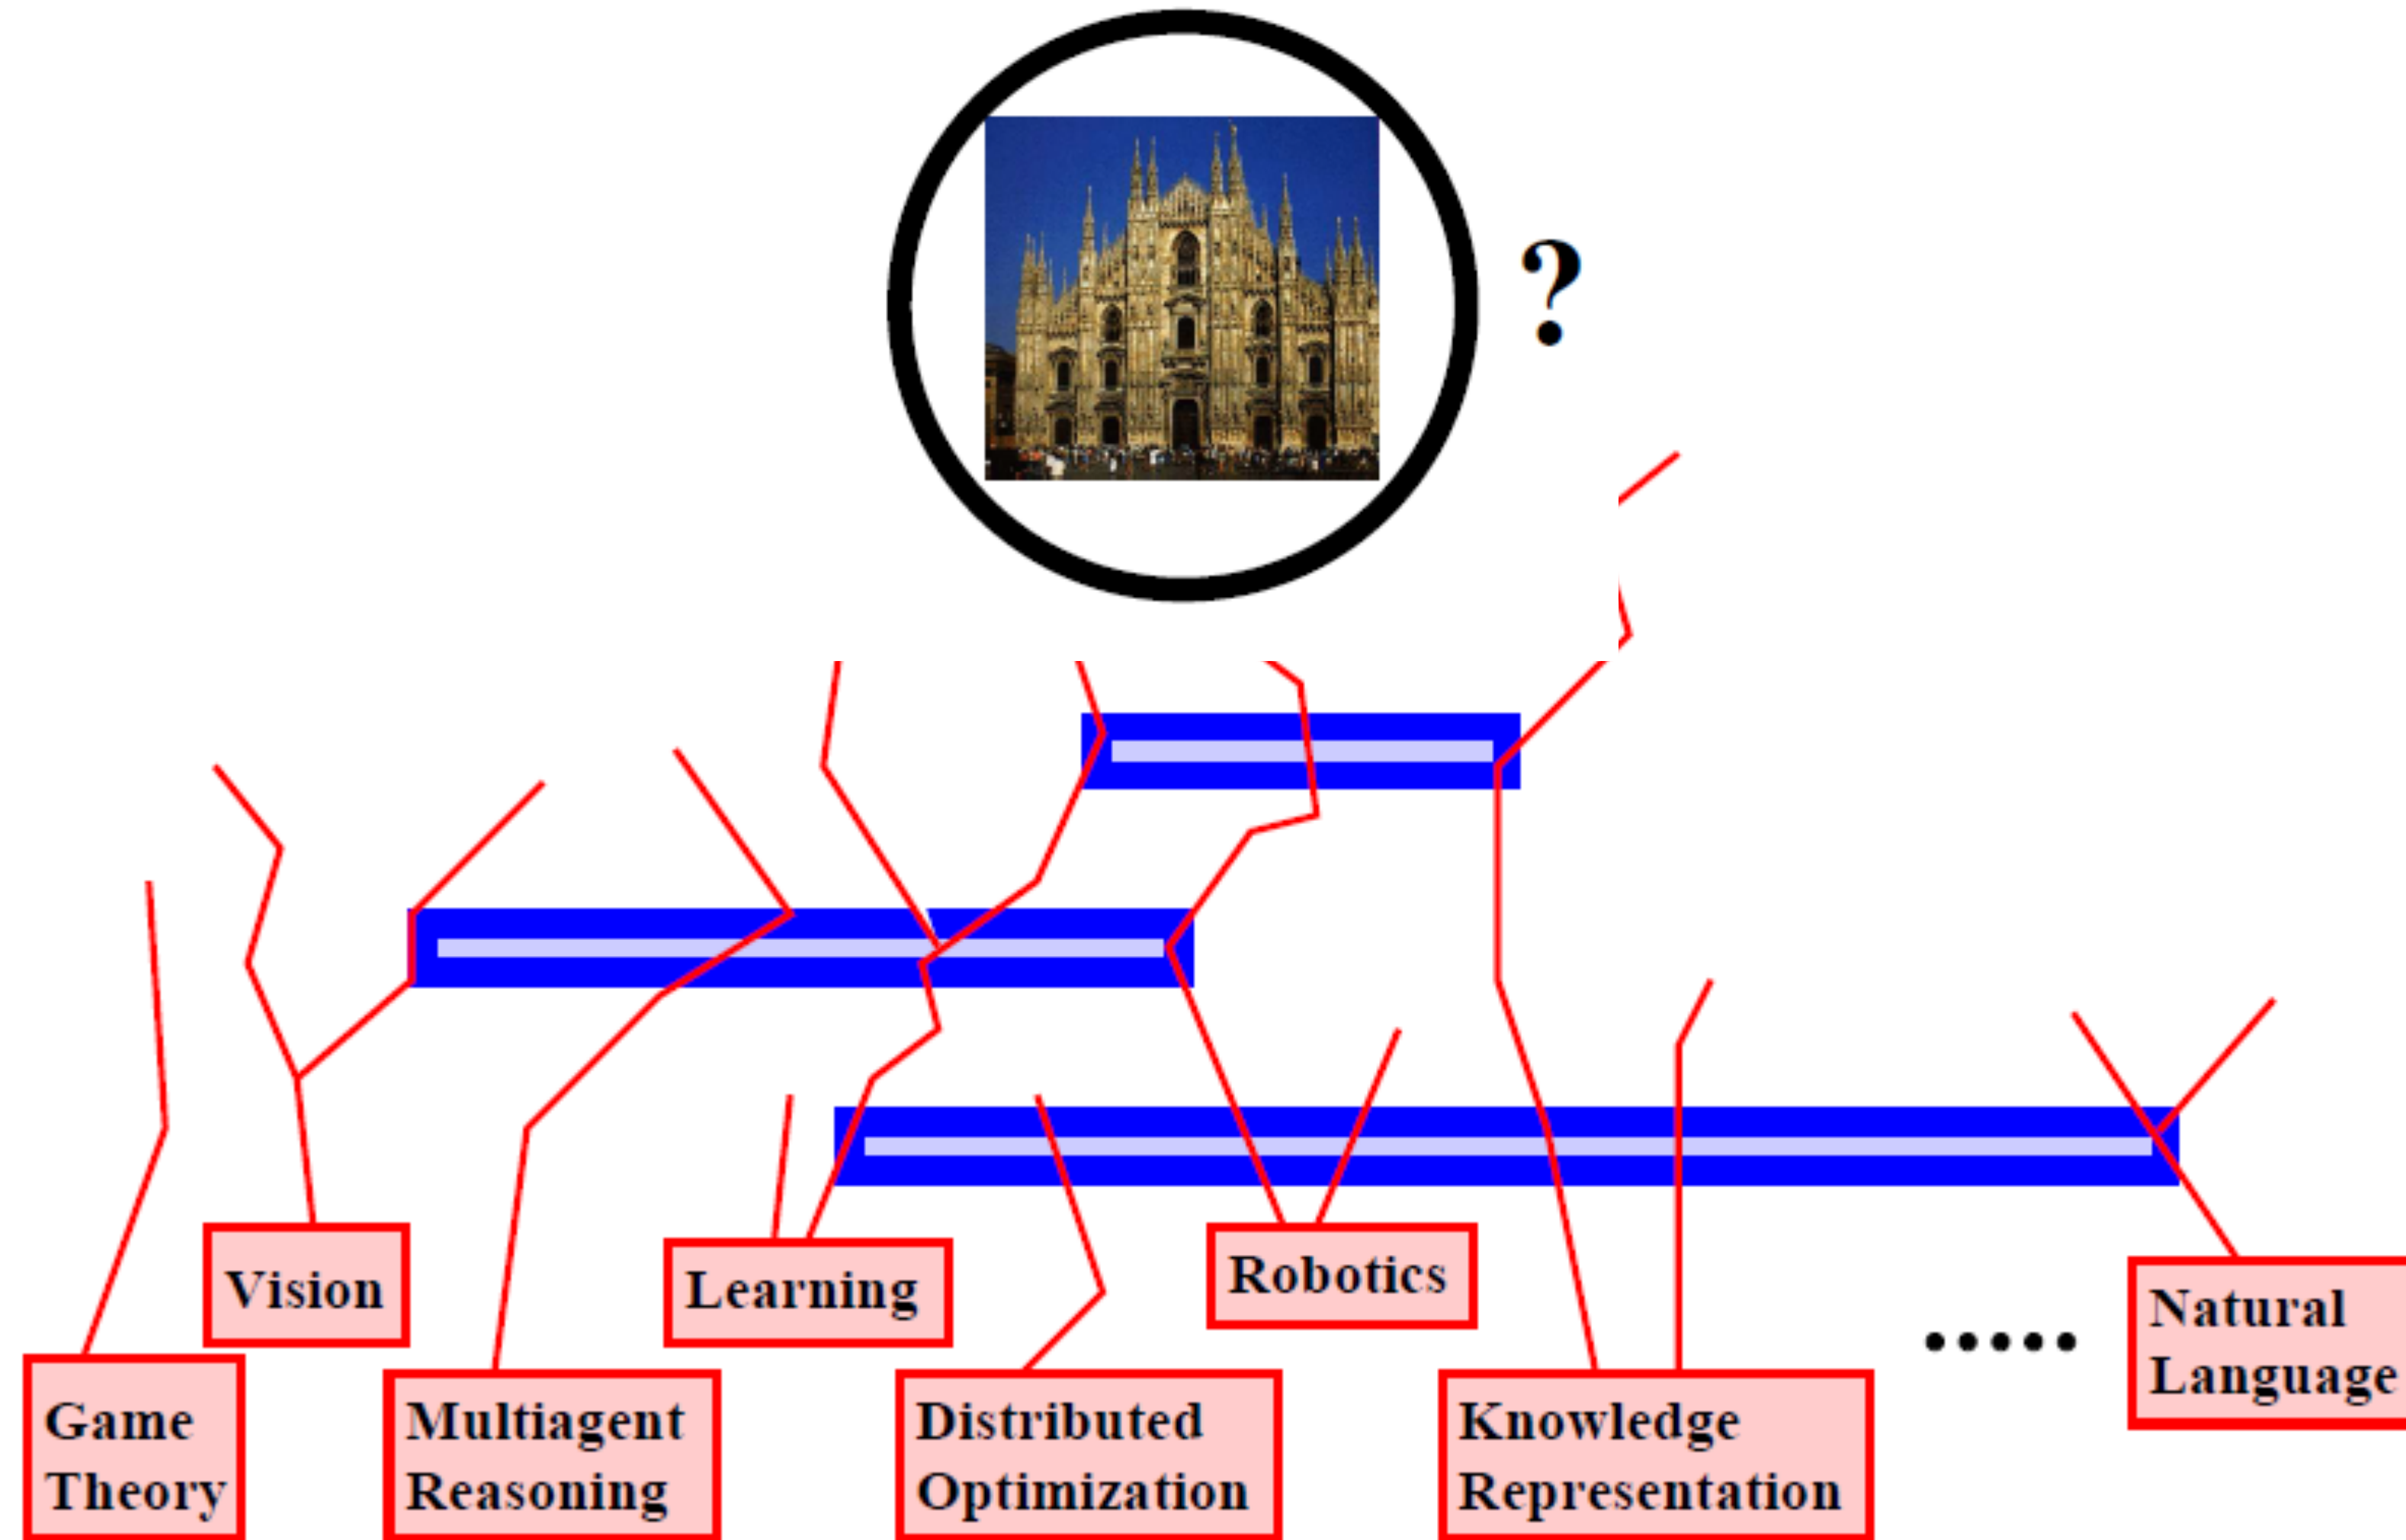
The bricks



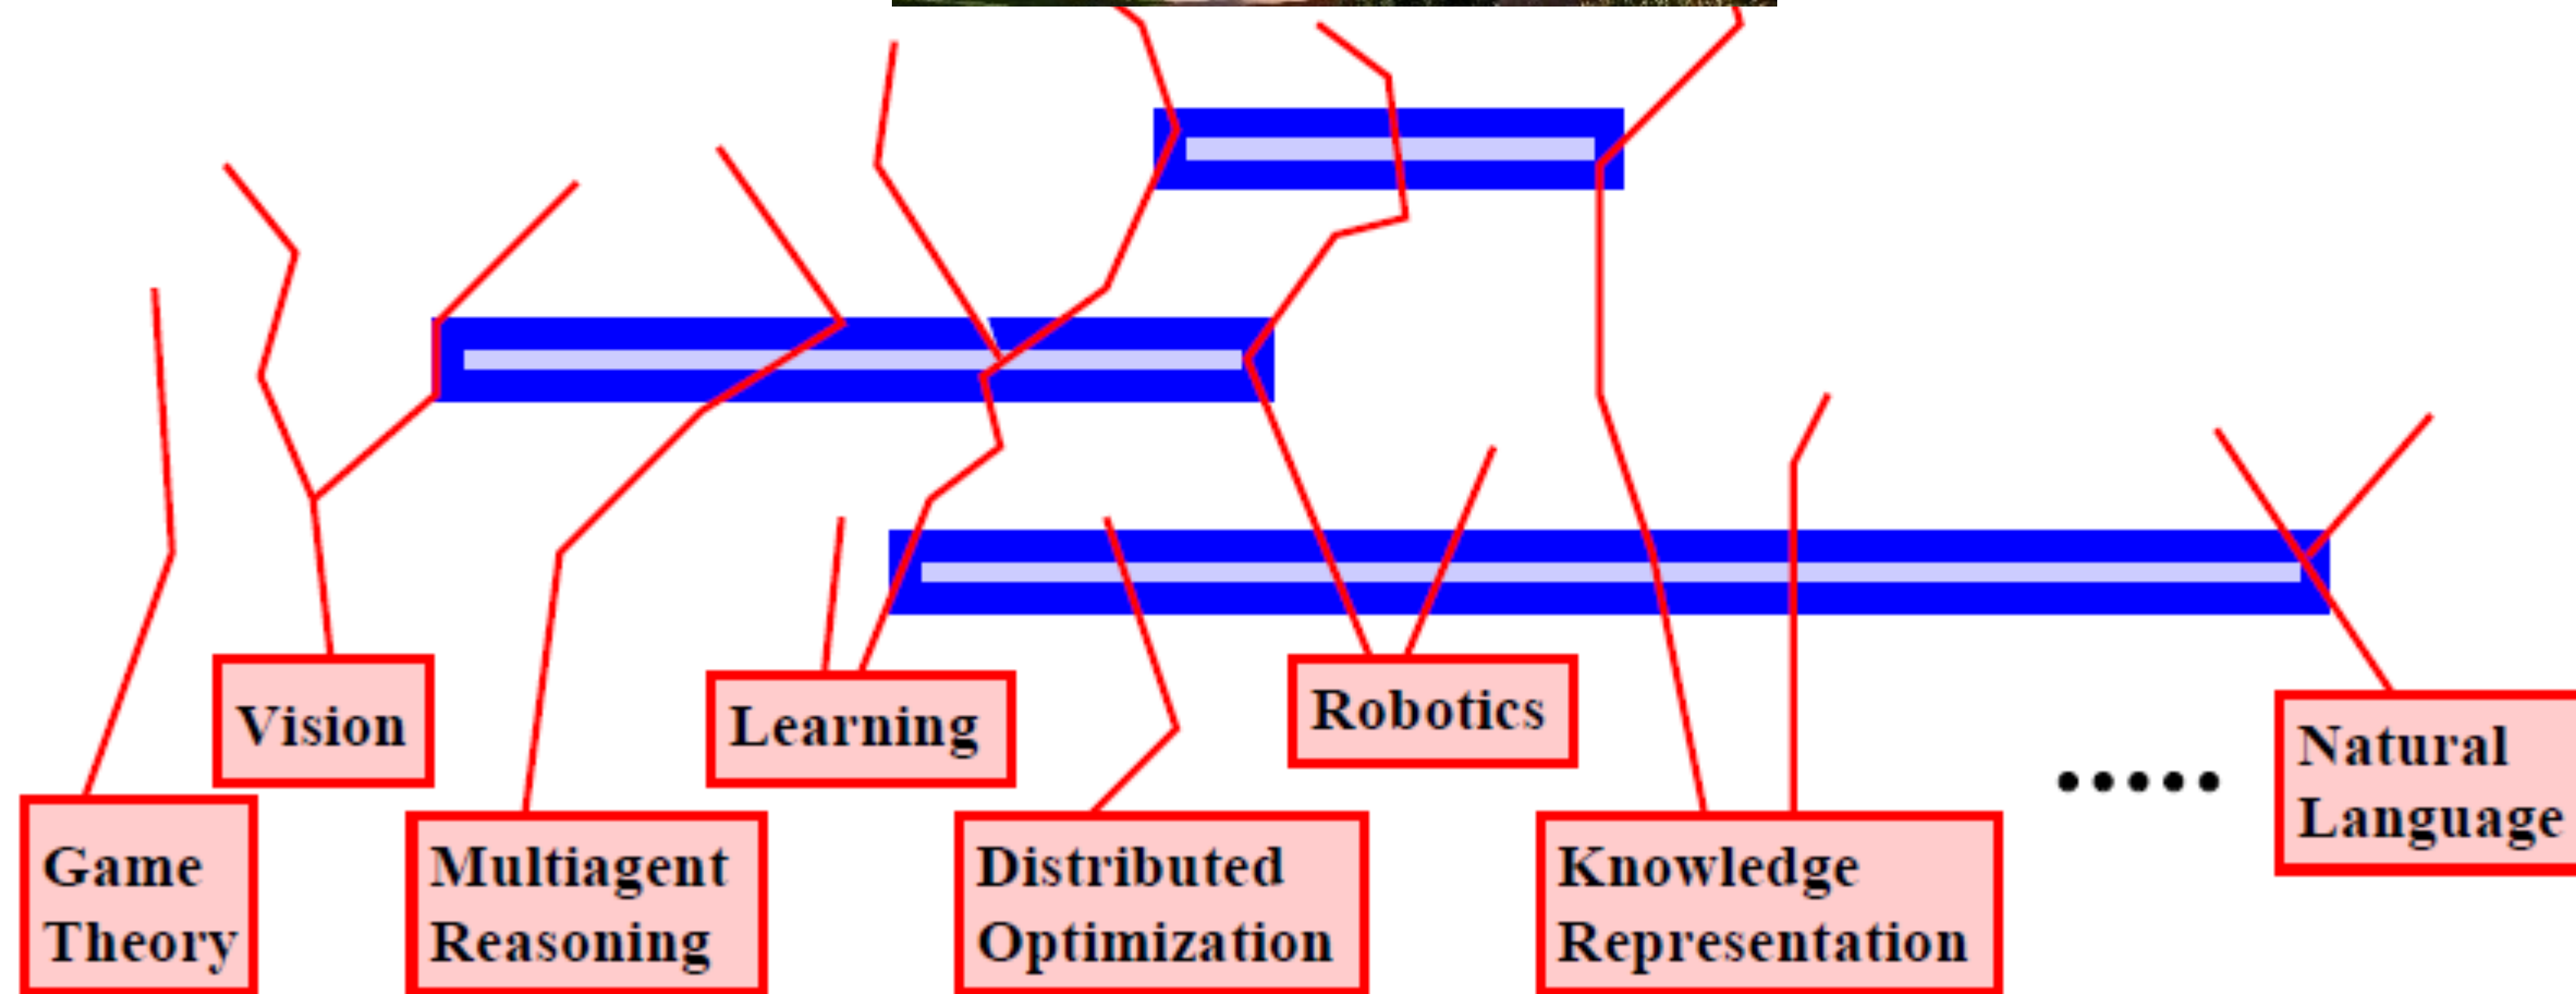
The beams and mortar



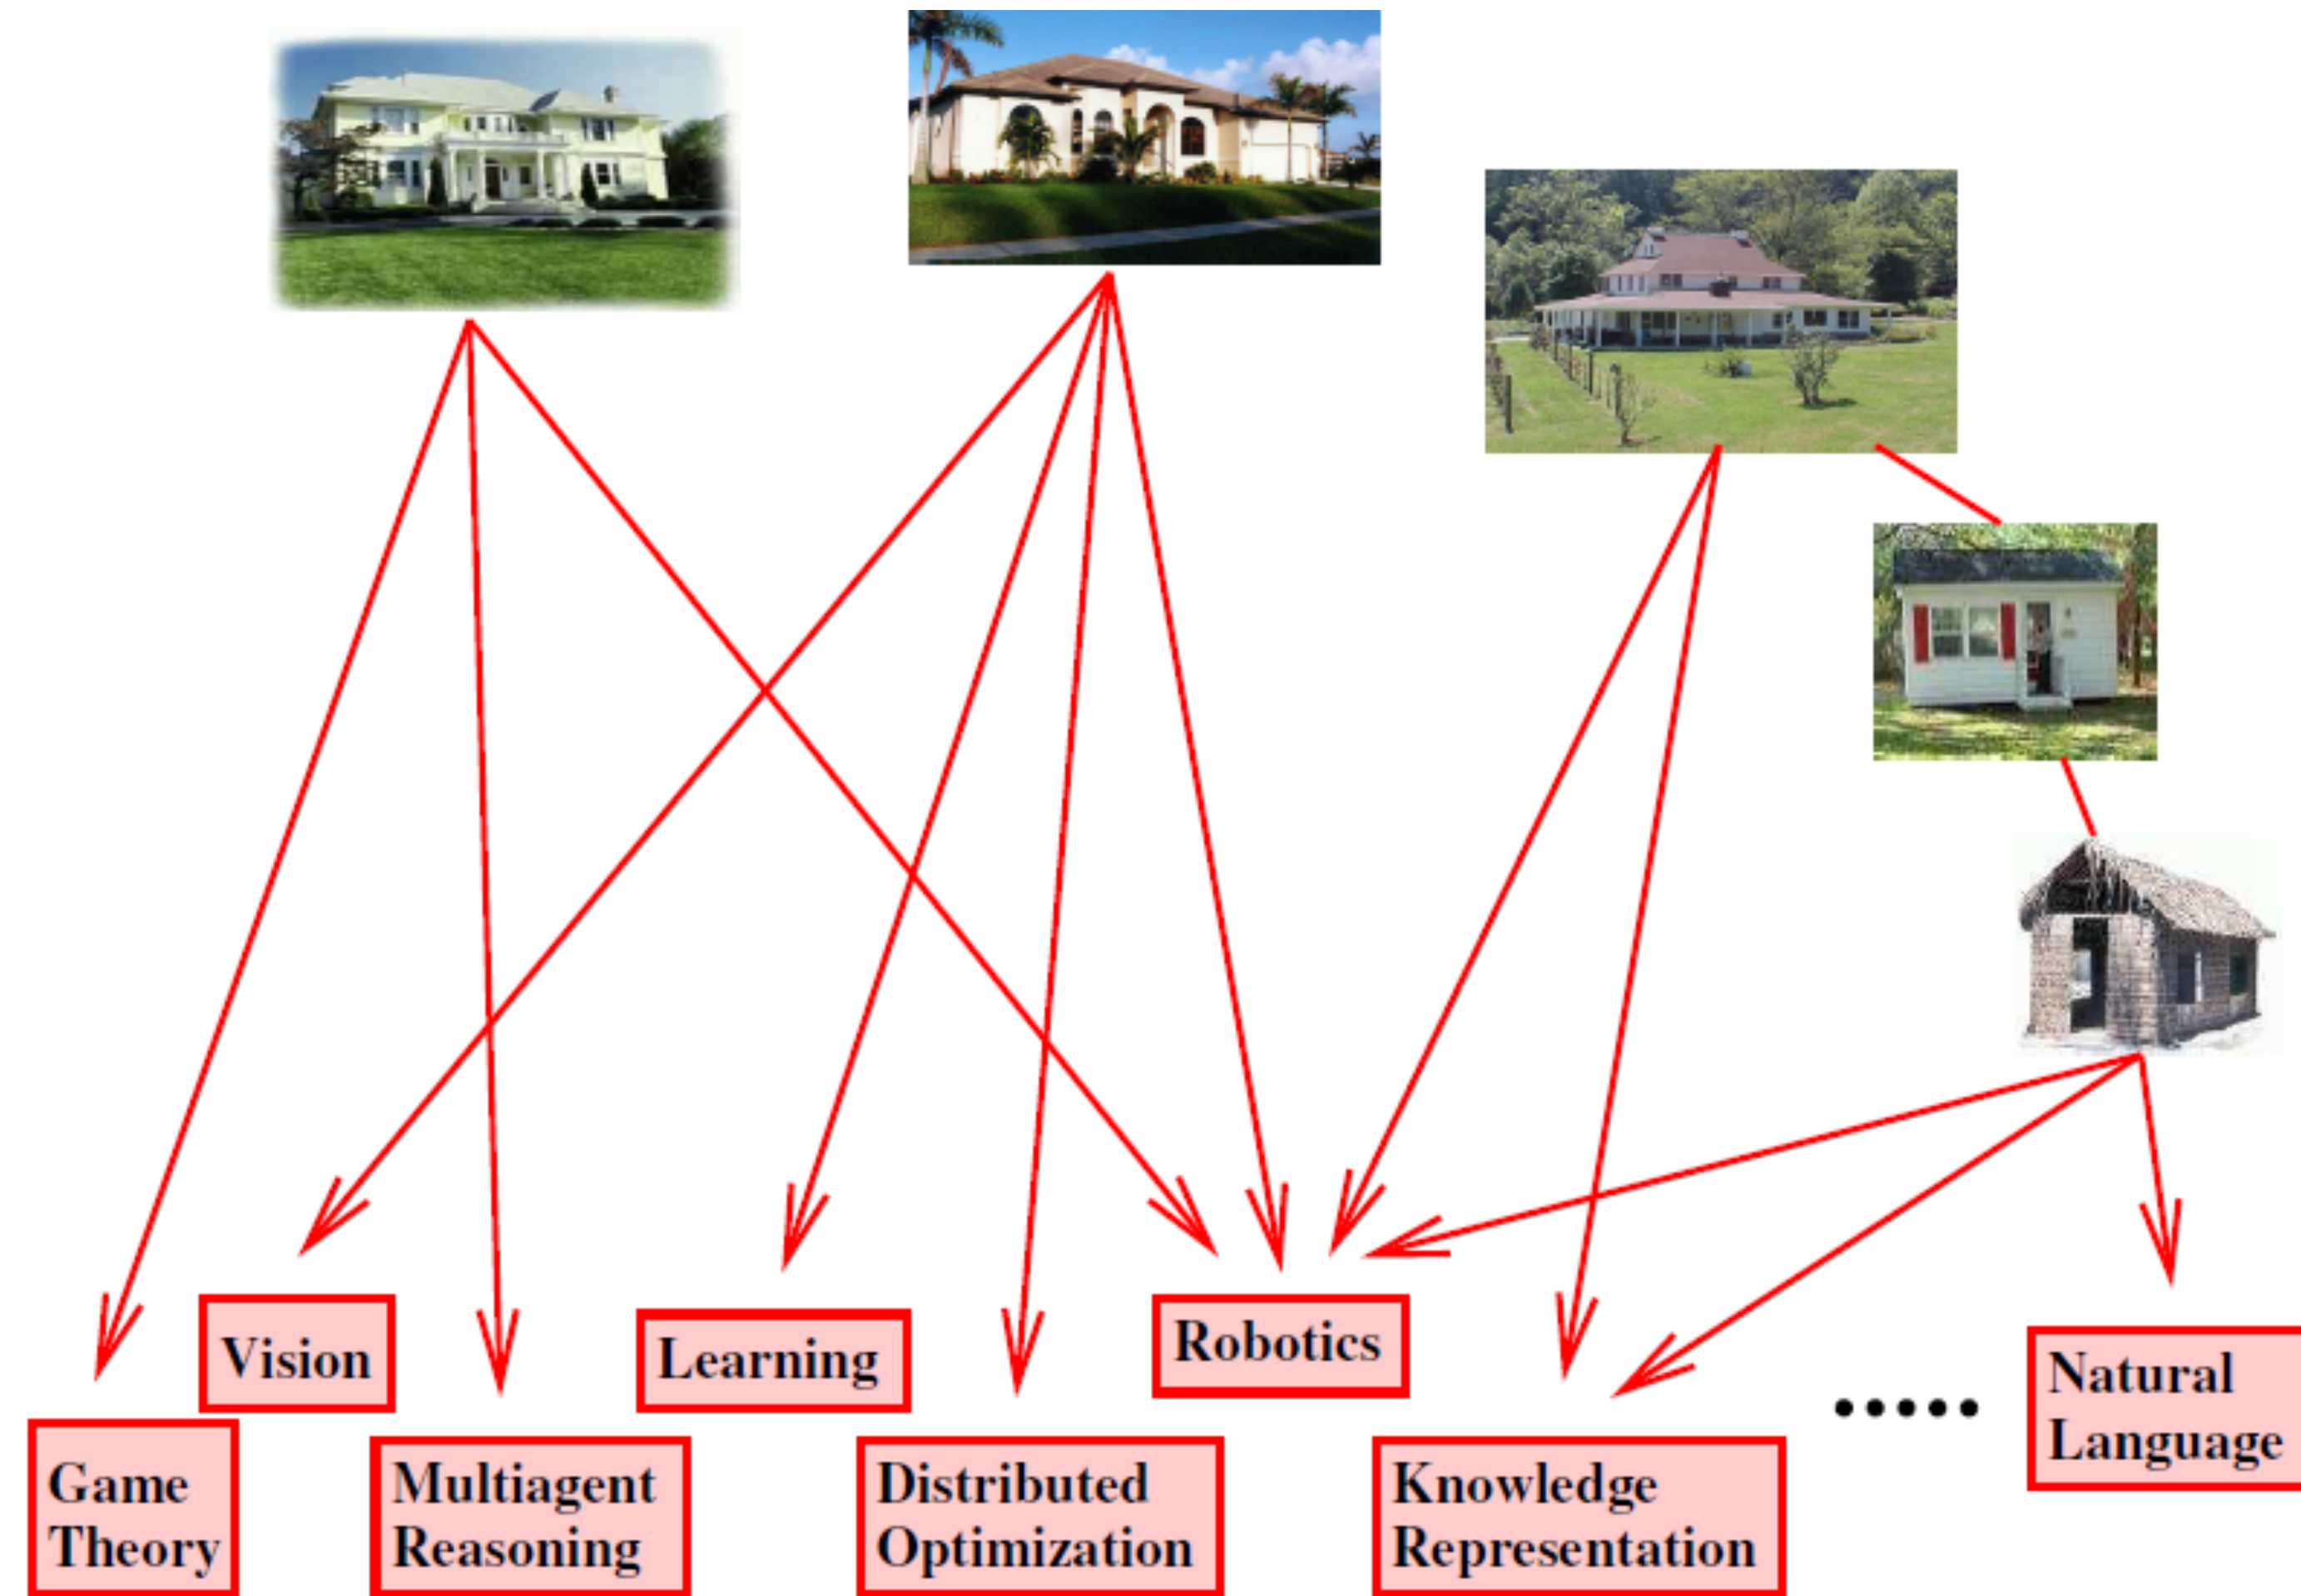
Towards a cathedral?



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

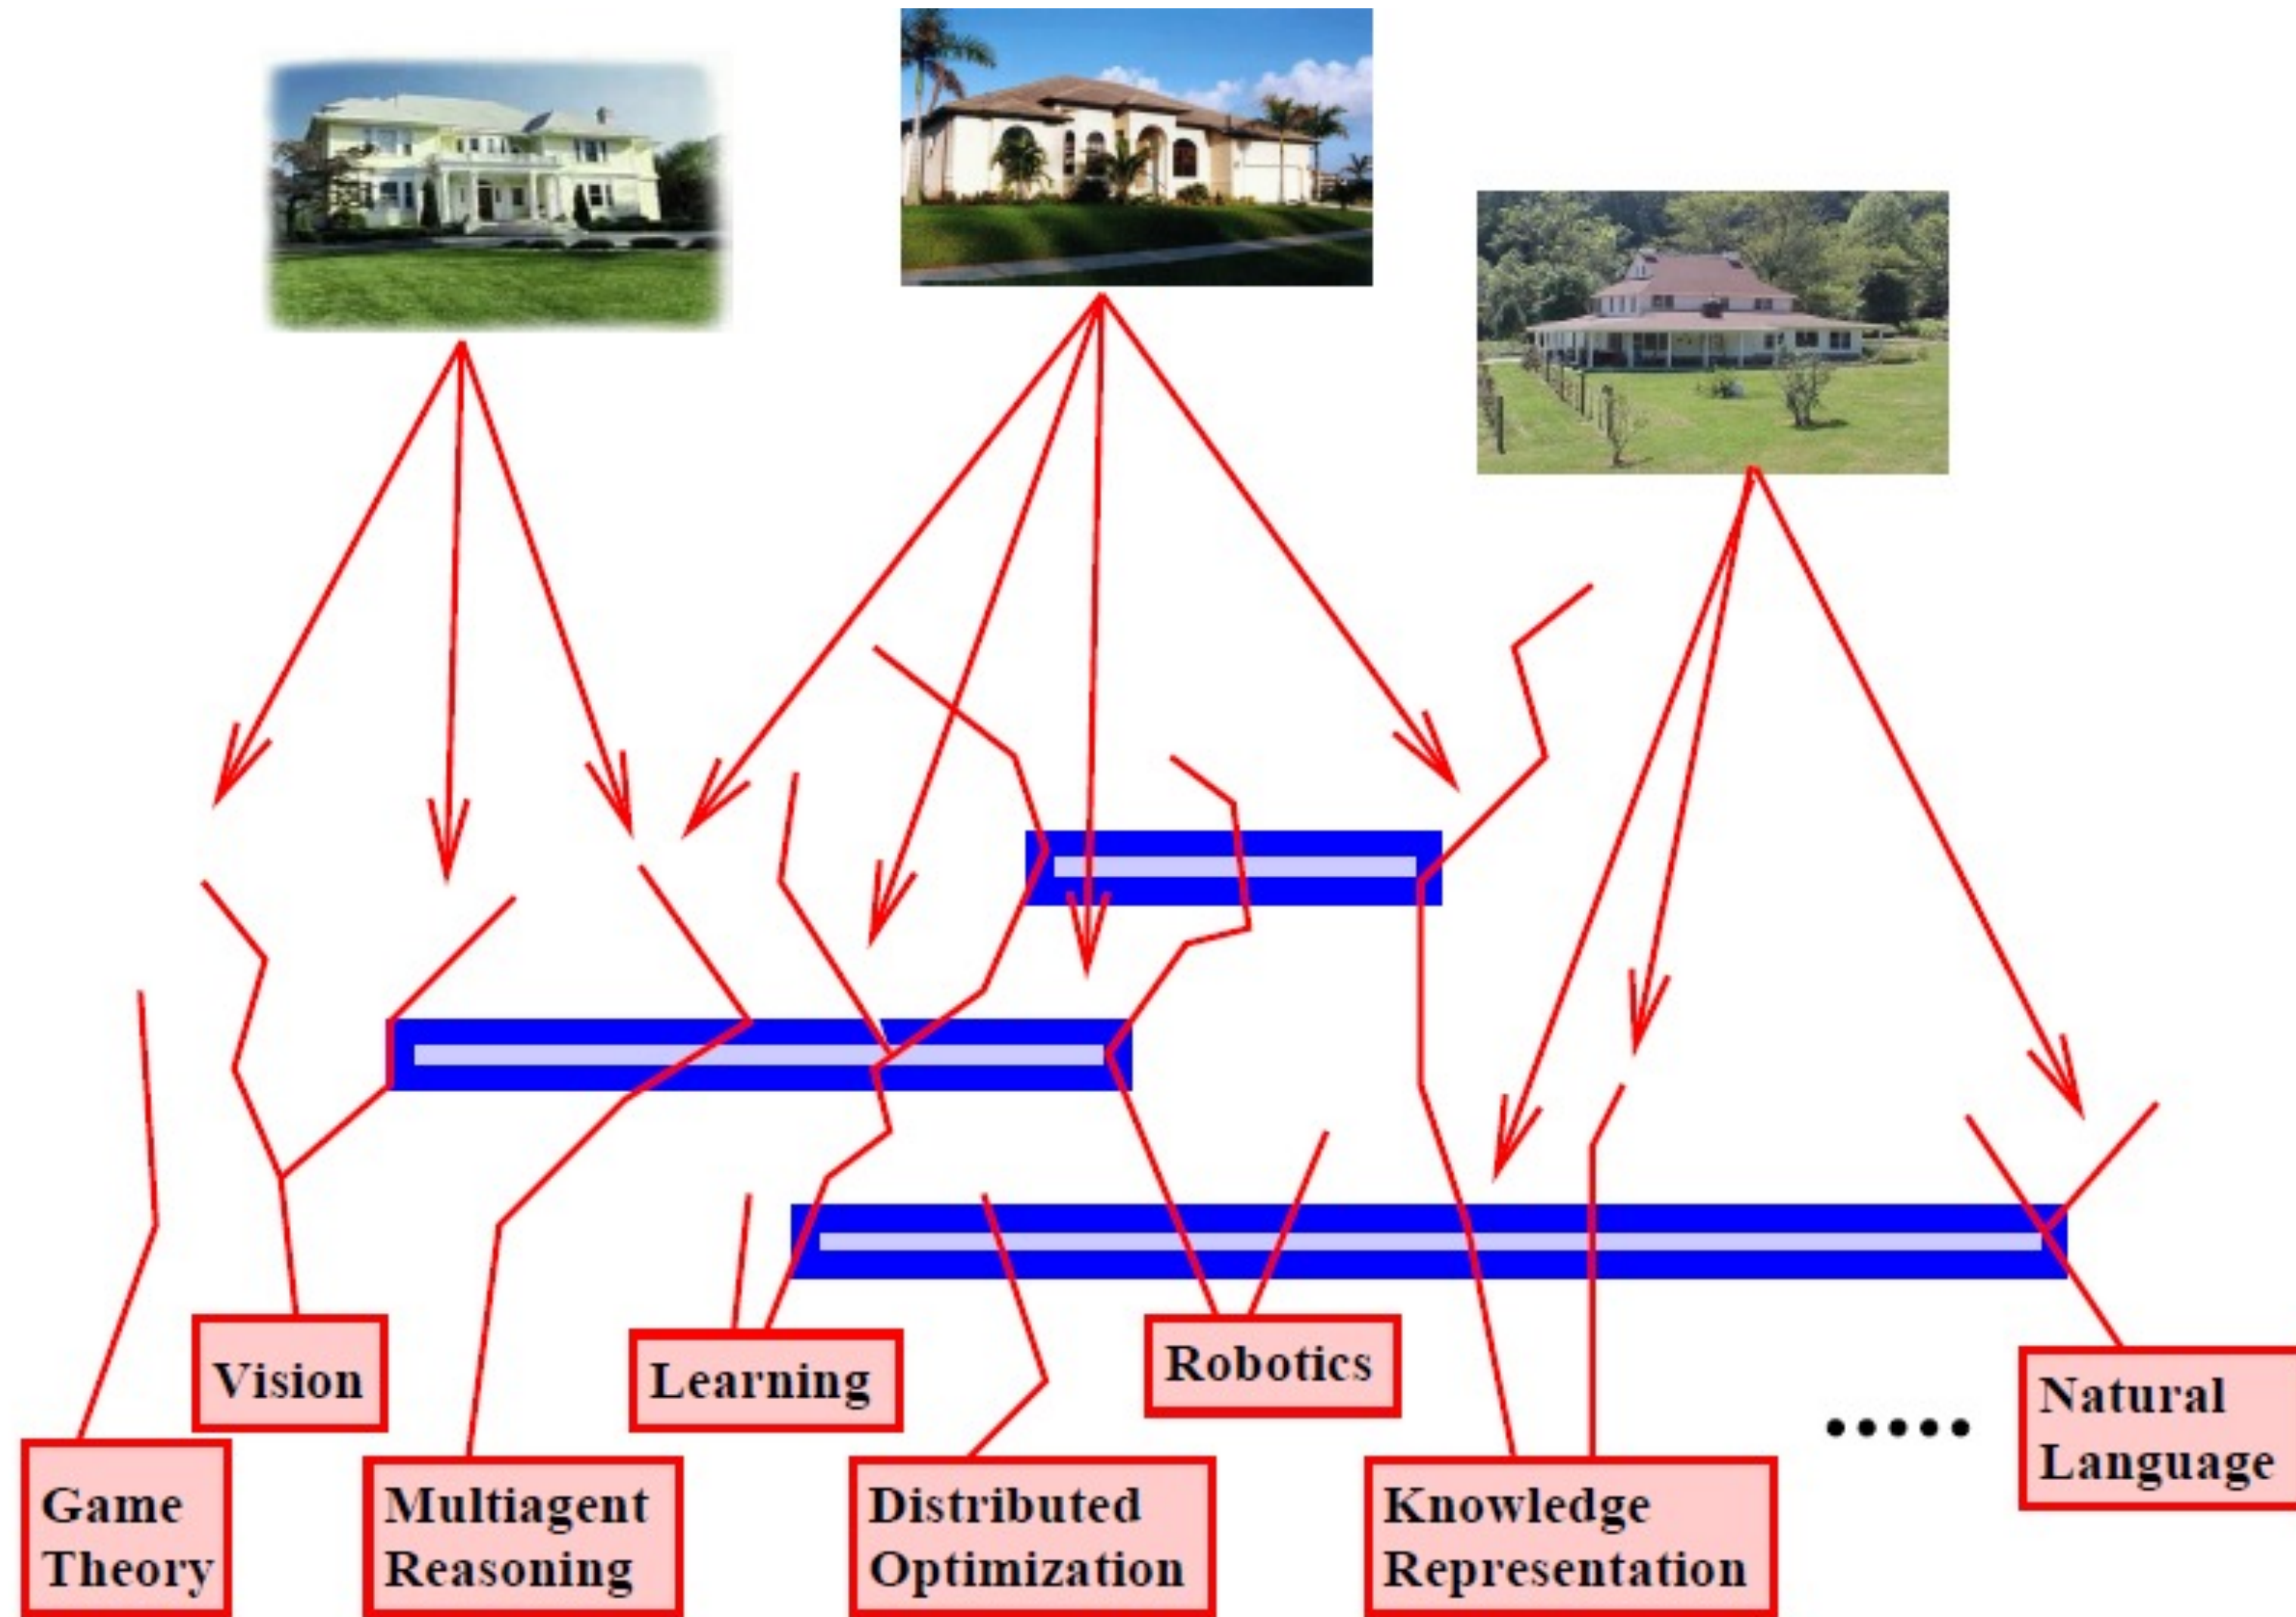


RoboCup soccer



Assistive robots

Meeting in the middle



Exciting Time to Study AI!



Stable Diffusion

Exciting Time to Study AI!

SA Can you summarize the story told in the "Acts" albums made by the band The Dear Hunter?

The Dear Hunter is an American rock band that has released several concept albums that tell a story known as "The Act" series. The story follows the life of a character known as "The Boy," who is born in the late 19th century and follows his journey through a series of events that are influenced by his relationships, struggles, and choices. The story is told through the lyrics of the songs and is divided into six albums, each of which corresponds to a different act in the story.

Act I: The Lake South, The River North (2006) - This album tells the story of The Boy's birth and childhood, including his relationship with his family and the events that shape his early life.

Act II: The Meaning of, and All Things Regarding Ms. Leading (2007) - This album follows The Boy as he leaves home and embarks on a journey of self-discovery, during which he meets a woman named Ms. Leading and becomes involved in a love triangle.

Act III: Life and Death (2009) - This album tells the story of The Boy's involvement in World War I and the impact it has on his life and relationships.

Act IV: Rebirth in Reprise (2013) - This album follows The Boy as he tries to rebuild his life after the events of Act III and faces new challenges and struggles.

Act V: Hymns with the Devil in Confessional (2016) - This album tells the story of The Boy's descent into despair and his efforts to find redemption.

Act VI: Regeneration (2020) - This album concludes the story of The Boy and the events of The Act series. It follows him as he comes to terms with his past and looks towards the future.

JC write a poem about the estate tax

On the wealthy, the estate tax
A burden they do not want to face
For they've worked hard and earned their place
In the upper echelon of society's race

Their wealth, a result of their labor and strife
Now subject to the government's knife
To take a slice of what they've amassed
And redistribute it to those less fortunate, alas

But the wealthy cry out in protest
Saying it's unfair and needs to be addressed
That their money should be passed on to their heirs
Not taken by the government in unfair shares

But the estate tax is a necessary tool
To balance the scales and play the game fair and cool
For without it, the gap between the rich and poor
Would only continue to grow, and that's no good for sure

So let the wealthy pay their fair share
And help to level the playing field out there
For in the end, we're all in this together
And the estate tax ensures a better future for all, forever.

M What is the best burger recipe?

There are many different burger recipes out there, and the "best" one will depend on personal preference. Some people prefer their burgers to be simple and classic, while others like to experiment with different ingredients and flavors.


Here is a basic burger recipe that can serve as a starting point:

Ingredients:

- 1 pound ground beef
- 4 hamburger buns
- Lettuce, tomato, onion, and other toppings of your choice
- Salt and pepper


Instructions:

Exciting Time to Study AI!




Today

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




Course Staff



Yuke Zhu

yukez@cs.utexas.edu



Zhenyu Jiang

zhenyu@cs.utexas.edu

OH: 10-11am, Wednesdays

Location: GDC 3.416



Pranav Atreya

pranavatreya@utexas.edu

OH: 3-4pm, Tuesdays

Location: GDC 3.416

Course Information

■ Communication:

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

Class website:

https://rpl.cs.utexas.edu/cs343_spring2023

(or Google “Yuke Zhu” and go to the Teaching tab)

■ Course technology:

- Gradescope for interactive homework (allow multiple submissions!)
- Autograded programming projects (submit via Gradescope)
- Make sure you have a CS Unix account IMMEDIATELY!
- Create a Gradescope account


The screenshot displays the Gradescope interface for course C S 343, Spring 2023. The course ID is 479050. The sidebar on the left includes navigation options: Dashboard (selected), Assignments, Roster, Extensions, Course Settings, Instructors (Yuke Zhu, Pranav Atreya, Zhenyu Jiang), and Course Actions (Leave Course). The main content area shows the 'Project 0' assignment, which is released on Jan 09 at 12:00AM and has a due date of Jan 12 at 11:59PM. The assignment is 100% graded and has 3 submissions. The interface also indicates 'No Published Grades' and a 'Things To Do' section with a task to 'Review and publish grades for Project 0 now that you're all done grading.'

Course Information

- Prerequisites:
 - Upper division standing
 - No formal class pre-reqs
 - **There will be a lot of math (and programming)**
- Coursework
 - ~6 homework assignments:
 - ~2 weeks for each, but overlapping
 - Online, autograded, allows multiple attempts
 - 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)
 - Capture the Flag contest
 - One midterm
 - One final

Textbook

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



After classes we'll post slides

Video recordings will be provided through Lectures Online (Canvas)

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 & Deep learning
 - AI research trends
- **Throughout: Applications**
 - Natural language, vision, robotics, games, ...



Homework Exercises

- Online on Gradescope
- Autograded text boxes / multiple choice
- Try as many times as you want!
 - May need to come to TA office hours to
 - unlock after too many attempts
- Goal: self-assess and prepare for exams
- 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 CTF 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:

- Class attendance and participation (10%)
- Homework Exercises (20%)
- Programming Assignments (25%)
- Capture the Flag contest (5%)
- 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** AI 100 report
 - **Read** the syllabus
 - **Register** for the class on Gradescope
 - **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 1/12 at 11:59pm via Gradescope)
- 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 readings and attend lectures! – 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 and be ready to interact
 - Office hours
 - 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?

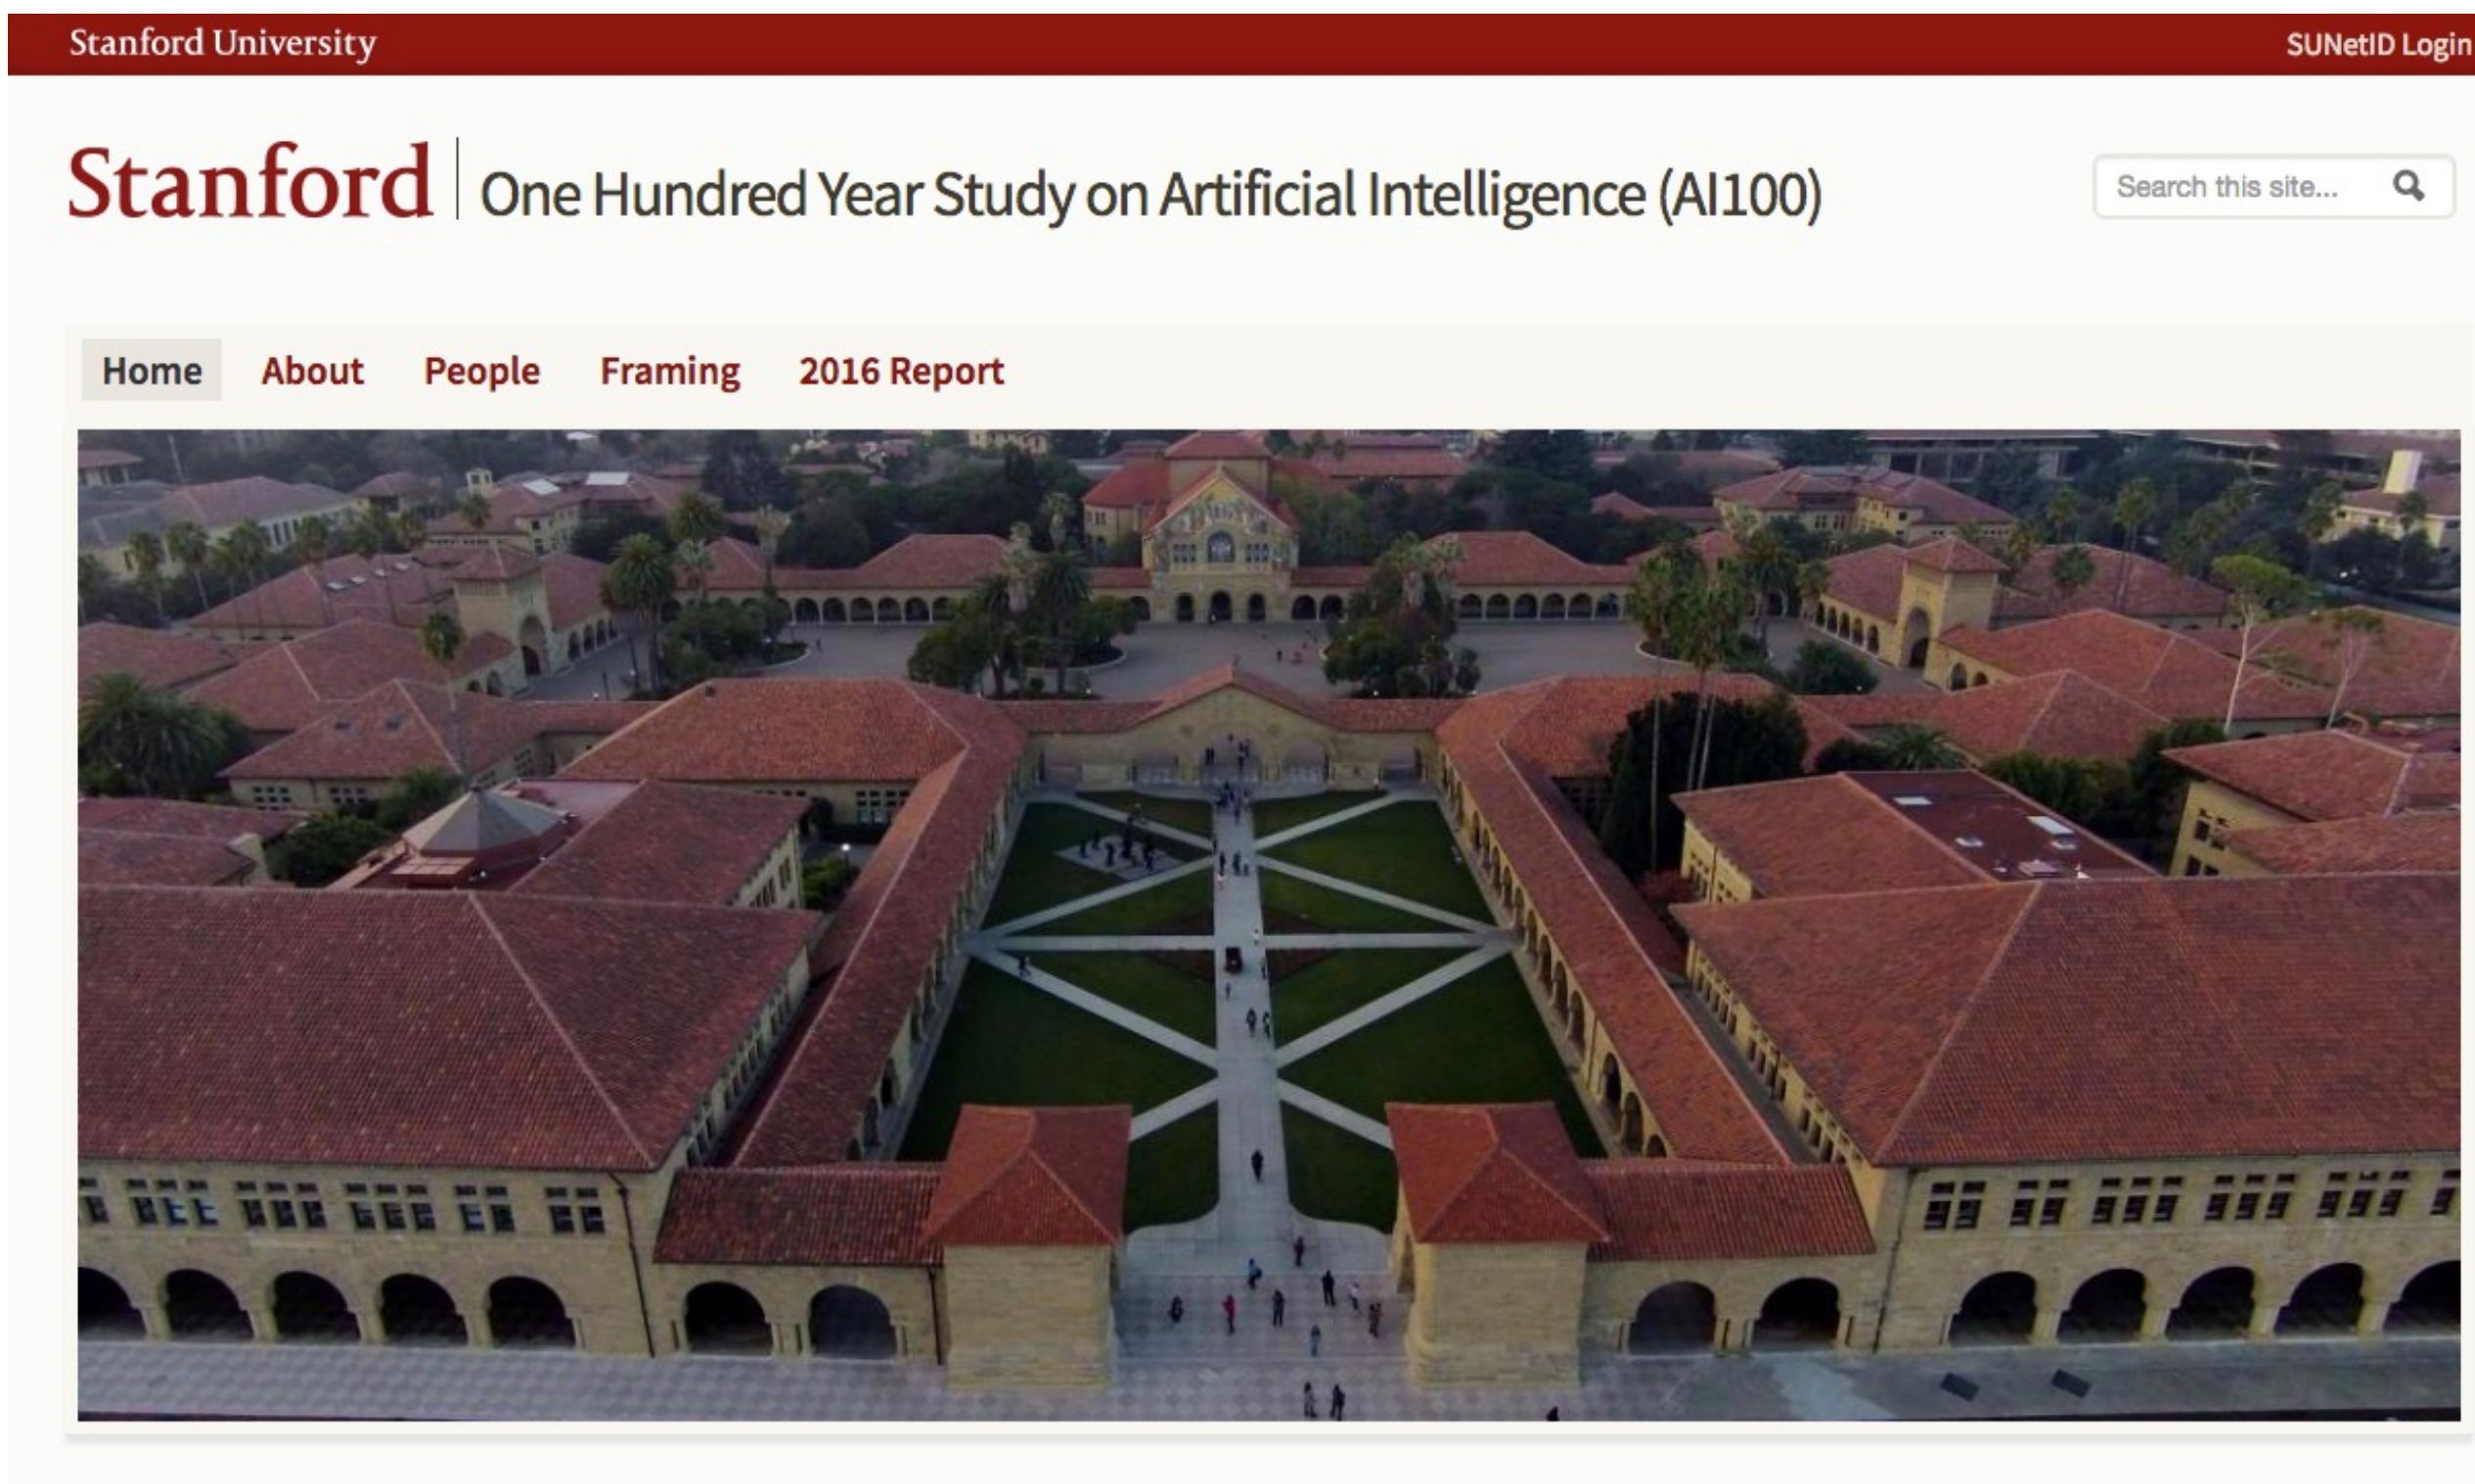
Discussion Question: Would you rather have been born:
– 50 years earlier? – 50 years later?

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!**

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.”


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



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**

Maximize Your Expected Utility

