CS 4700: Foundations of Artificial Intelligence

Spring 2020

Prof. Haym Hirsh

CS 4701: Practicum in Artificial Intelligence

Organizational Meeting:
Fri, Jan 25, 4:30pm, Gates G01
Makeup time TBA
(see website http://www.cs.cornell.edu/courses/cs4701/)
Email: FAI-Practicum-l@cornell.edu

CS 4700 and CS 4701 are uncoupled
4701 is not synced up with 4700

CS 4700: Foundations of Artificial Intelligence

Spring 2020

Prof. Haym Hirsh

Today

Overview of AI

Overview of the course

Today

Overview of AI

Overview of the course

What is Artificial Intelligence (AI)?

“AI” is an over-used term

Different people mean different things

Historical

John McCarthy (1927-2011)

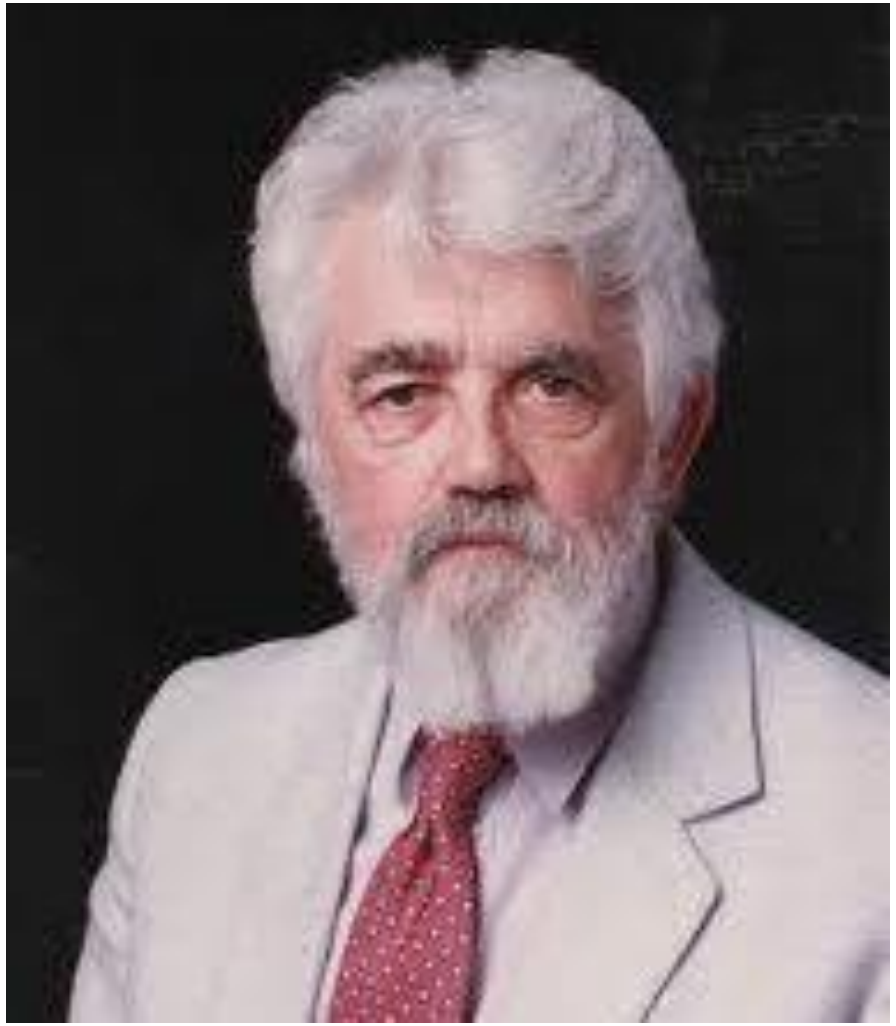

A PROPOSAL FOR THE

DARTMOUTH SUMMER RESEARCH PROJECT

ON ARTIFICIAL INTELLIGENCE

J. McCarthy, Dartmouth College
M. L. Minsky, Harvard University
N. Rochester, I. B. M. Corporation
C. E. Shannon, Bell Telephone Laboratories

August 31, 1955

A PROPOSAL FOR THE

DARTMOUTH SUMMER RESEARCH PROJECT

ON ARTIFICIAL INTELLIGENCE

J. McCarthy, Dartmouth College
M. L. Minsky, Harvard University
N. Rochester, I. B. M. Corporation
C. E. Shannon, Bell Telephone Laboratories

August 31, 1955

A Proposal for the

DARTMOUTH SUMMER RESEARCH PROJECT ON ARTIFICIAL INTELLIGENCE

We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. 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. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.

The following are some aspects of the artificial intelligence problem:

1) Automatic Computers

If a machine can do a job, then an automatic calculator can be programmed to simulate the machine. The speeds and memory capacities of present computers may be insufficient to simulate many of the higher functions of the human brain, but the major obstacle is not lack of machine capacity, but our inability to write programs taking full advantage of what we have.

2) How Can a Computer be Programmed to Use a Language

It may be speculated that a large part of human thought consists of manipulating words according to rules of reasoning

A Proposal for the

DARTMOUTH SUMMER RESEARCH PROJECT ON ARTIFICIAL INTELLIGENCE

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.

solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.

The following are some aspects of the artificial intelligence problem:

1) Automatic Computers

If a machine can do a job, then an automatic calculator can be programmed to simulate the machine. The speeds and memory capacities of present computers may be insufficient to simulate many of the higher functions of the human brain, but the major obstacle is not lack of machine capacity, but our inability to write programs taking full advantage of what we have.

2) How Can a Computer be Programmed to Use a Language

It may be speculated that a large part of human thought consists of manipulating words according to rules of reasoning

(1950s)

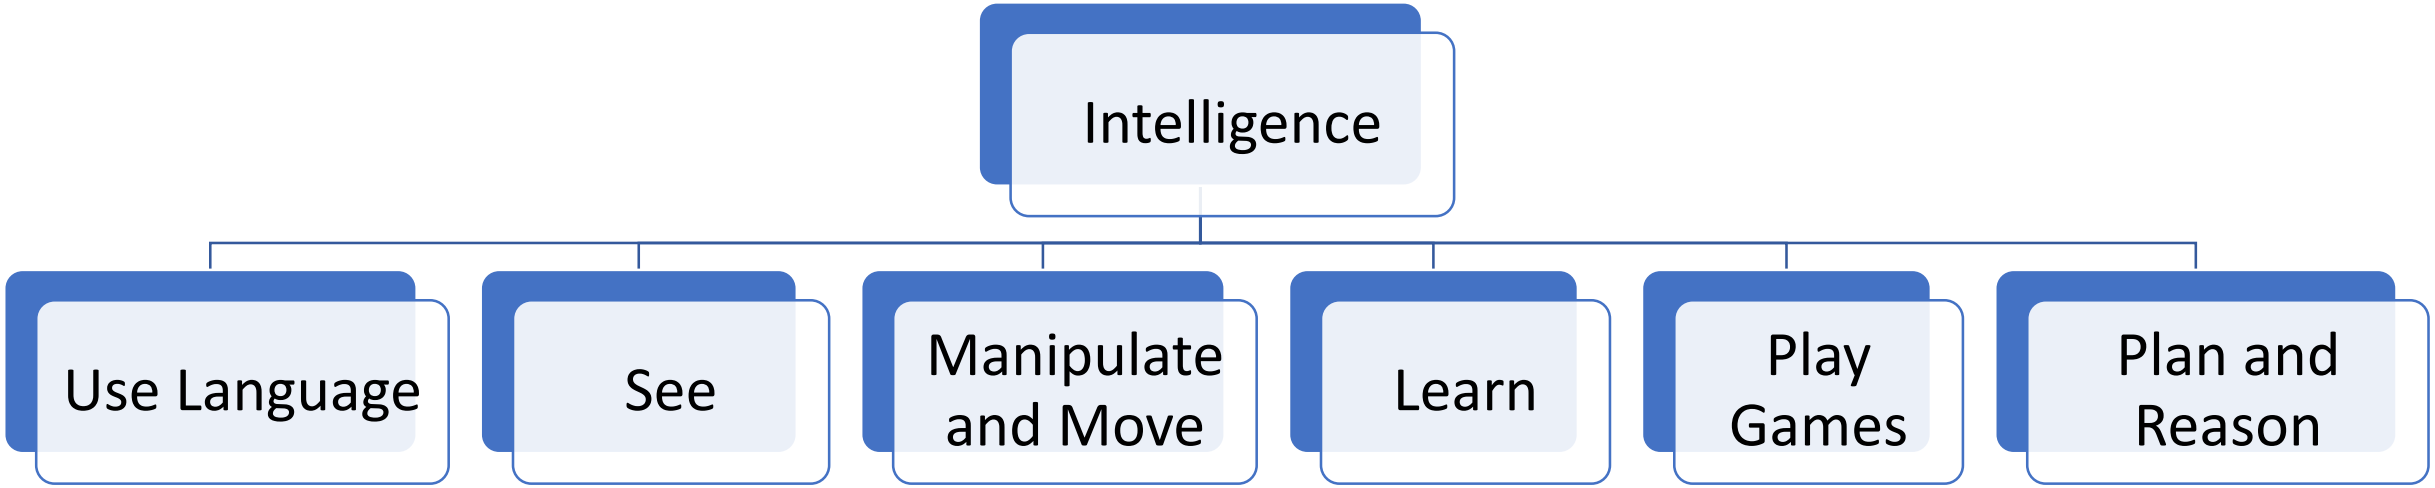
Intelligence
Use Language
See
Manipulate and Move
Learn
Play Games
Plan and Reason

(1950s)

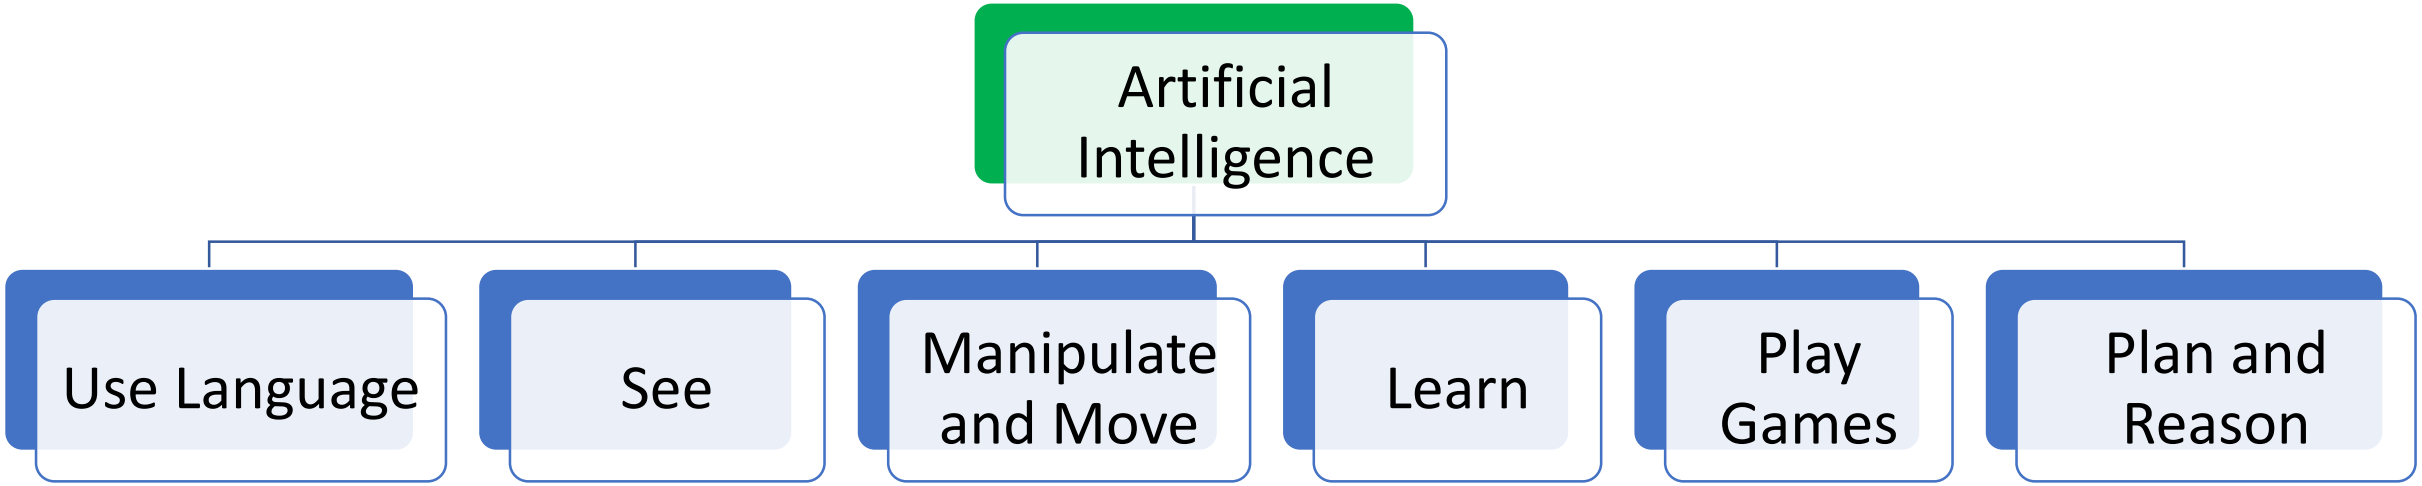
Artificial Intelligence
Use Language
See
Manipulate and Move
Learn
Play Games
Plan and Reason

(1950s)

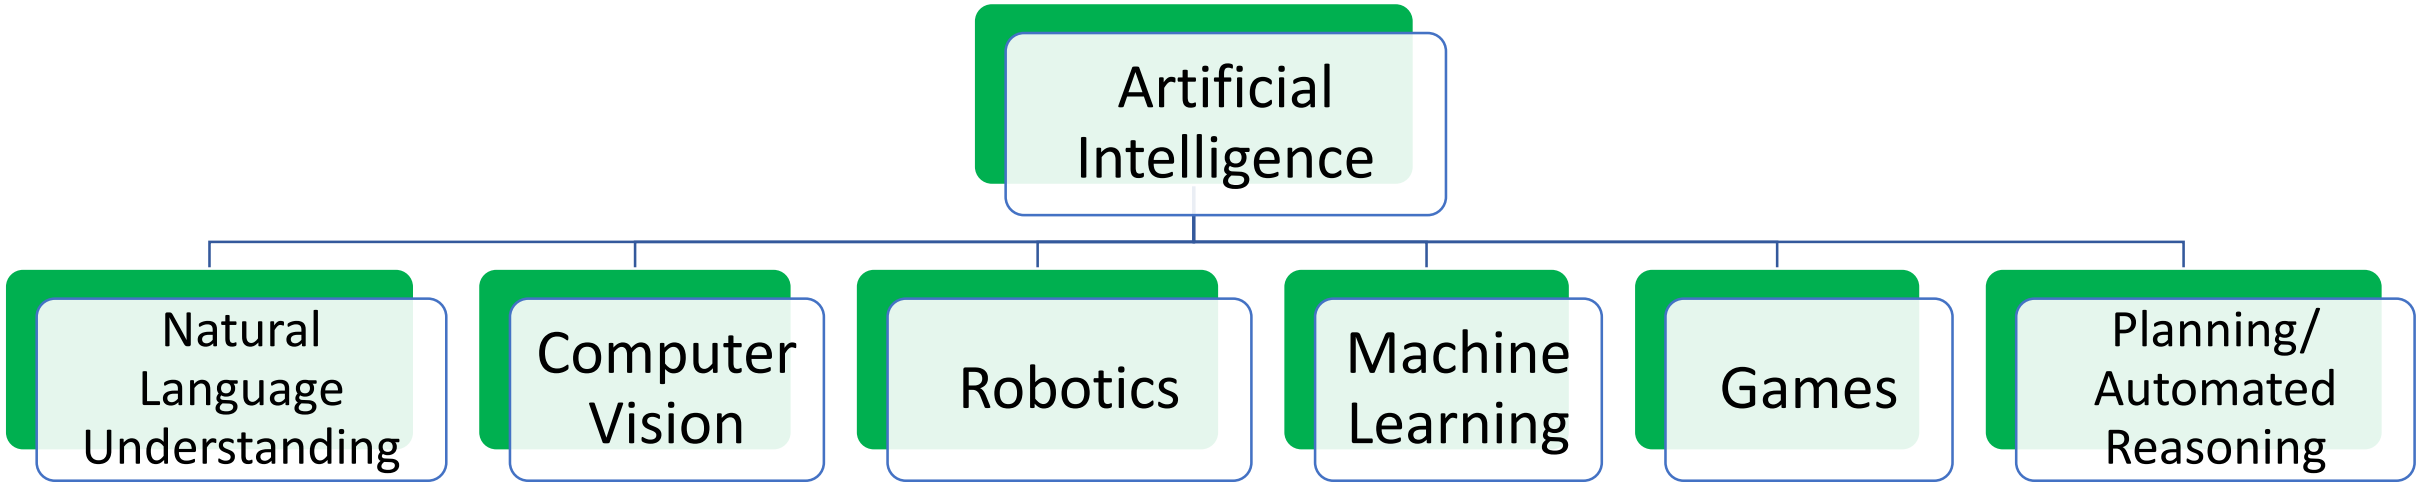
Artificial Intelligence
Natural Language Understanding
Computer Vision
Robotics
Machine Learning
Games
Planning/ Automated Reasoning

(1950s)

The AI Spectrum

Narrow “cognitive” skills

Broad capabilities

The AI Spectrum

Narrow “cognitive” skills
“Weak AI”

Broad capabilities
“Strong AI”

The AI Spectrum

Narrow “cognitive” skills
“Weak AI”

Broad capabilities
“Strong AI”
“Artificial General Intelligence”
AGI

The AI Spectrum

Narrow “cognitive” skills
“Weak AI”

Broad capabilities
“Strong AI”

The AI Spectrum

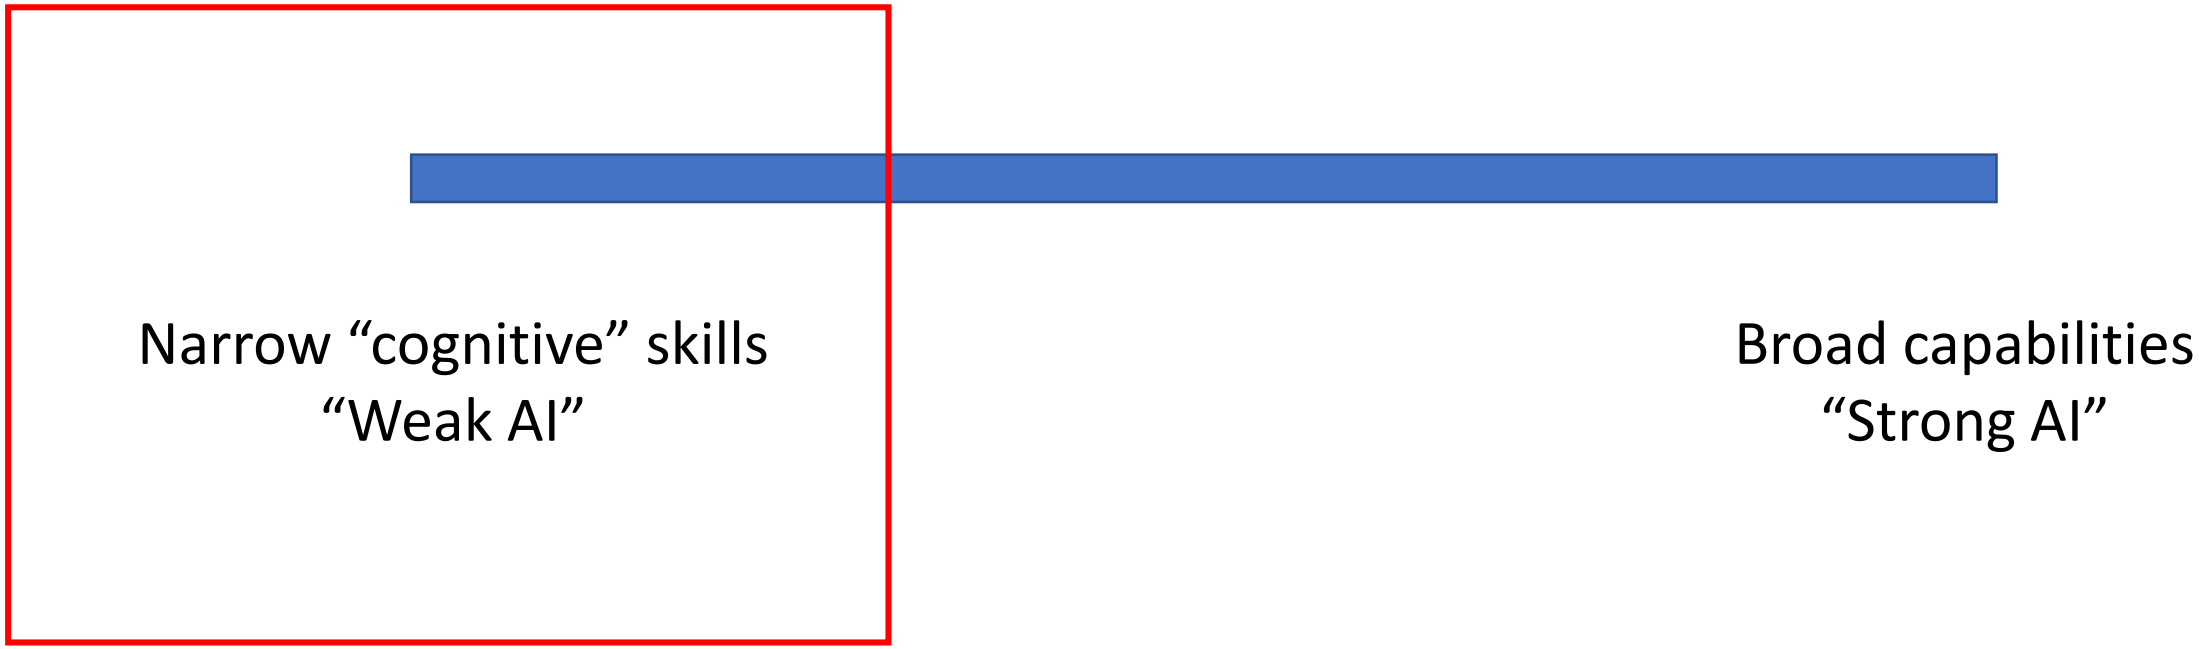

The AI Spectrum

Narrow “cognitive” skills
“Weak AI”

Broad capabilities
“Strong AI”

Important successes

The AI Spectrum

Collateral “Successes”

Narrow “cognitive” skills
“Weak AI”

Broad capabilities
“Strong AI”

Important successes

Collateral Successes

- Time sharing
- Functional programming languages
- Hardware verification
- …
- Web search engines
- Recommendation systems
- Language technologies
- Machine learning
- …
- Autonomous vehicles?
- Face recognition?

The AI Spectrum

Collateral “Successes”

Narrow “cognitive” skills
“Weak AI”

Broad capabilities
“Strong AI”

Important successes

The AI Spectrum

Collateral “Successes”

Narrow “cognitive” skills
“Weak AI”

Important successes

Broad capabilities
“Strong AI”

The AI Spectrum

Collateral “Successes”

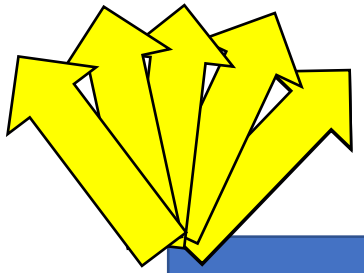

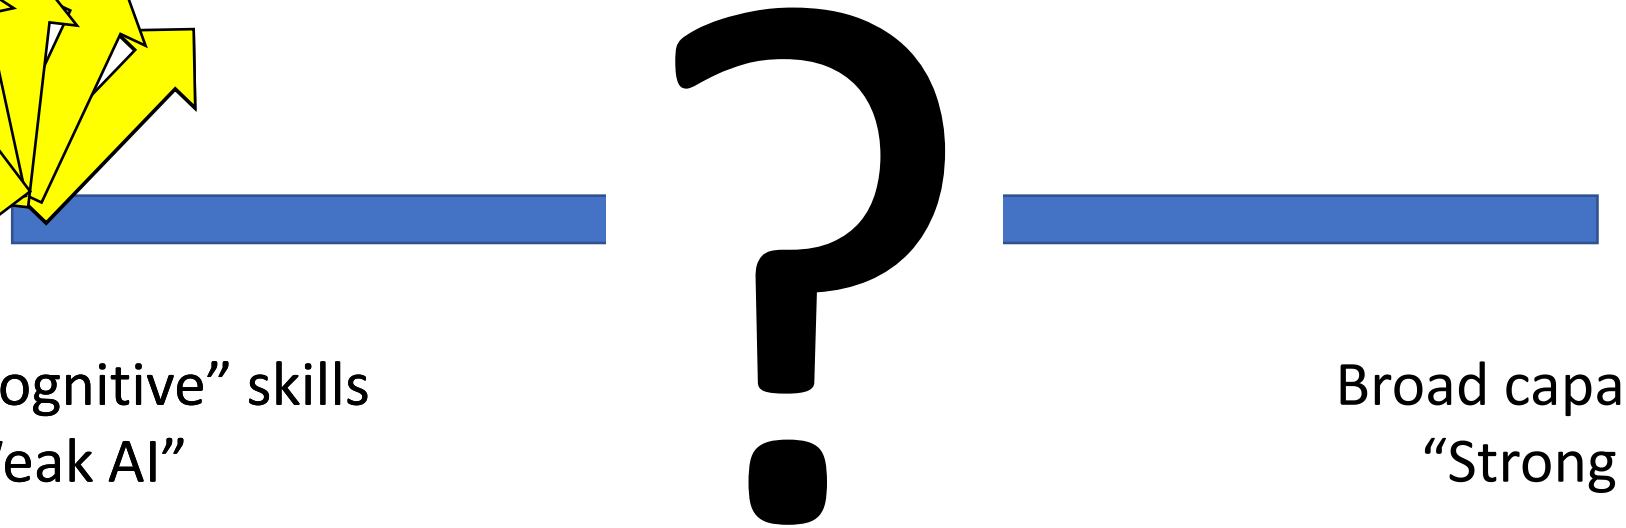

Narrow “cognitive” skills
“Weak AI”

Broad capabilities
“Strong AI”

Important successes


Robocup19 - Soccer Simulation 3D - Final
UTAustinVilla 0:0 magmaOffenburg 00:36
Playmode: PlayOn
Server Speed: 99%
Watch later
Share
4:48 / 16:09
HD
YouTube

The AI Spectrum

Collateral “Successes”

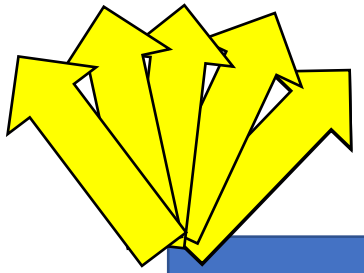

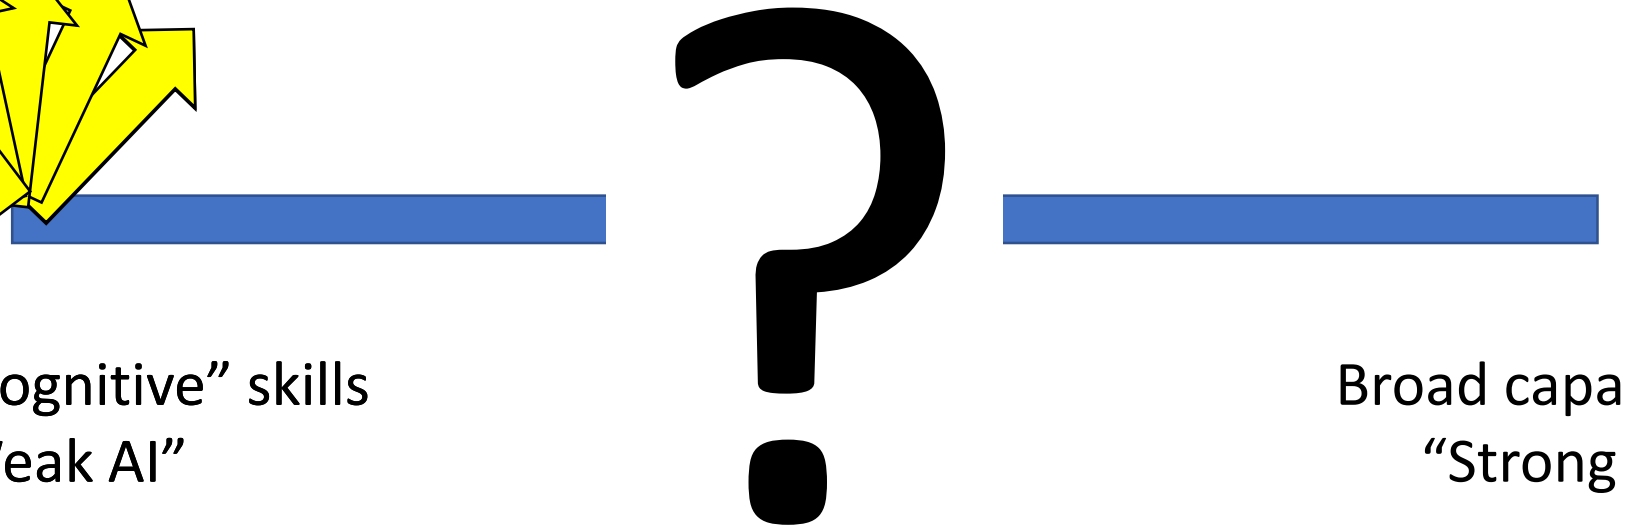

Narrow “cognitive” skills
“Weak AI”

Broad capabilities
“Strong AI”

Important successes

The AI Spectrum

Collateral “Successes”

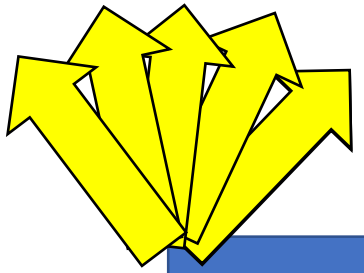

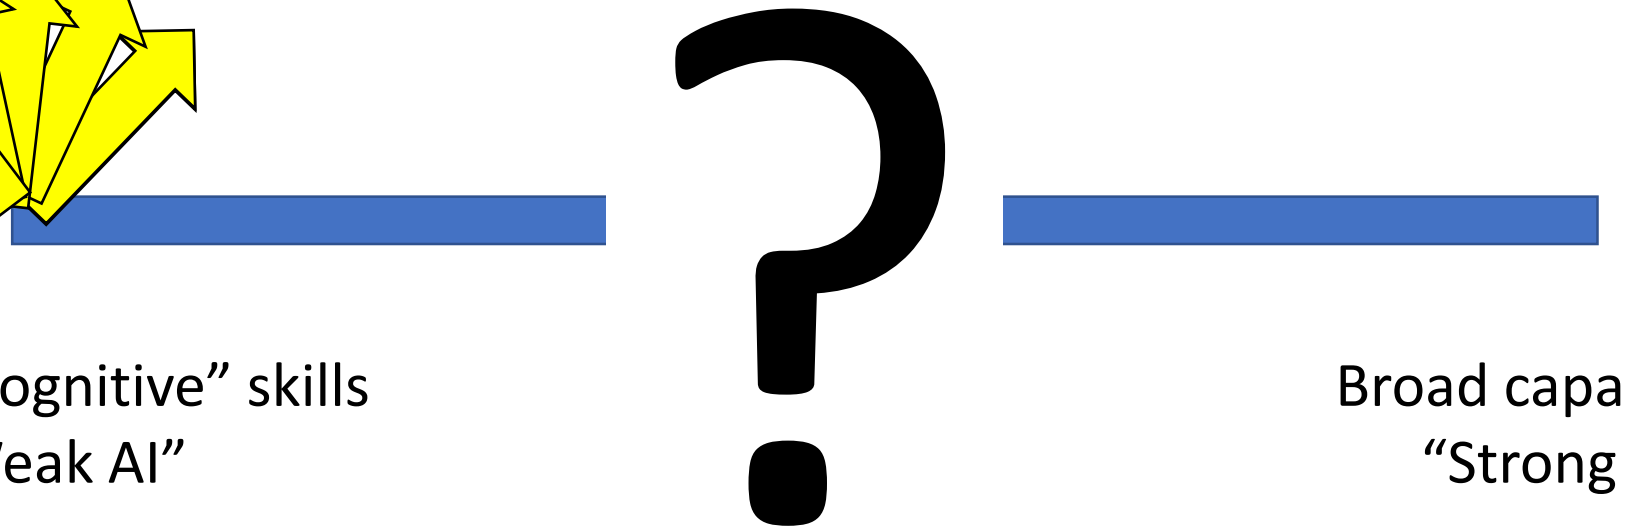

Narrow “cognitive” skills
“Weak AI”

Important successes

Broad capabilities
“Strong AI”

Fearmongering
Utopian idealism
Bad predictions
Science Fiction

The AI Spectrum

Collateral “Successes”

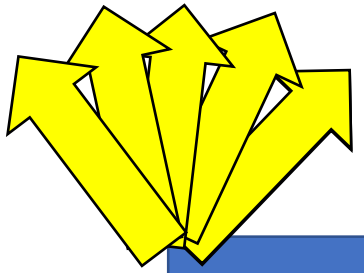

Narrow “cognitive” skills
“Weak AI”

Important successes

?

Broad capabilities
“Strong AI”

Fearmongering
Utopian idealism
Bad predictions
Science Fiction

Legitimate concerns

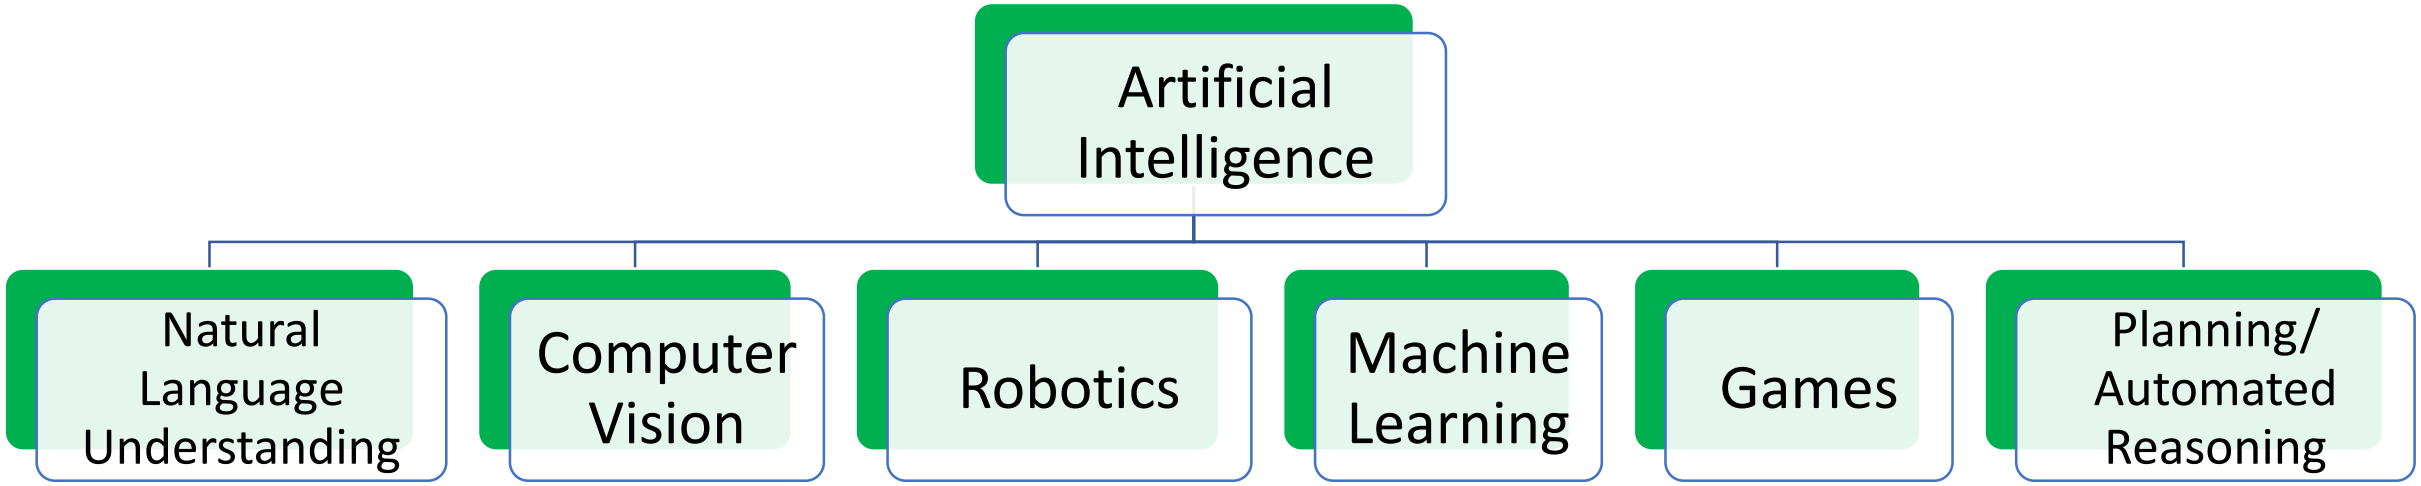

1990s: Common ideas arising in separate areas:
Probabilistic modeling
Machine learning, mathematical optimization of error on training data

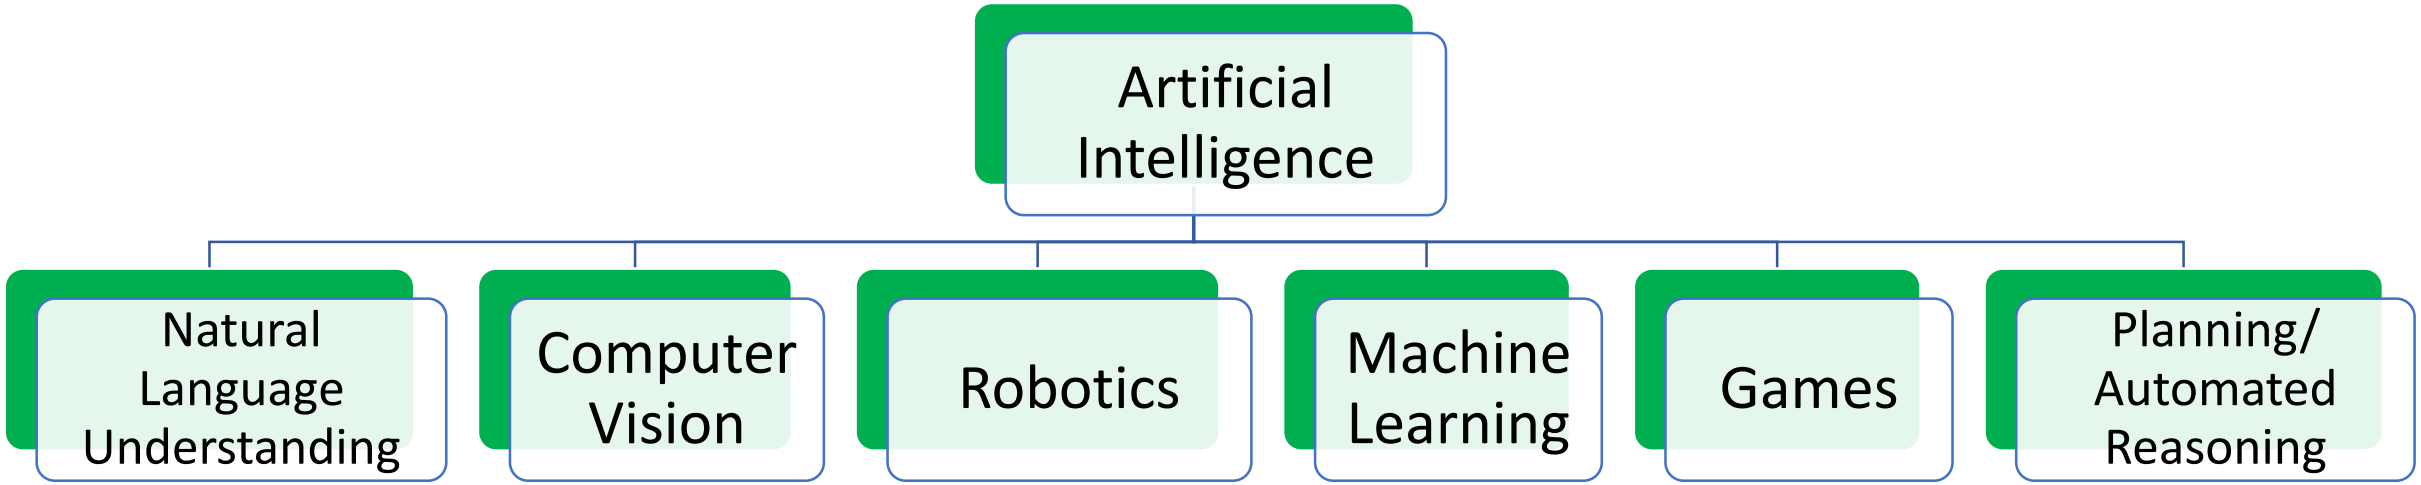

2000-present: Successes due to

- "Standing on the shoulders of giants"
- Moore's Law
- Machine learning/lots of data

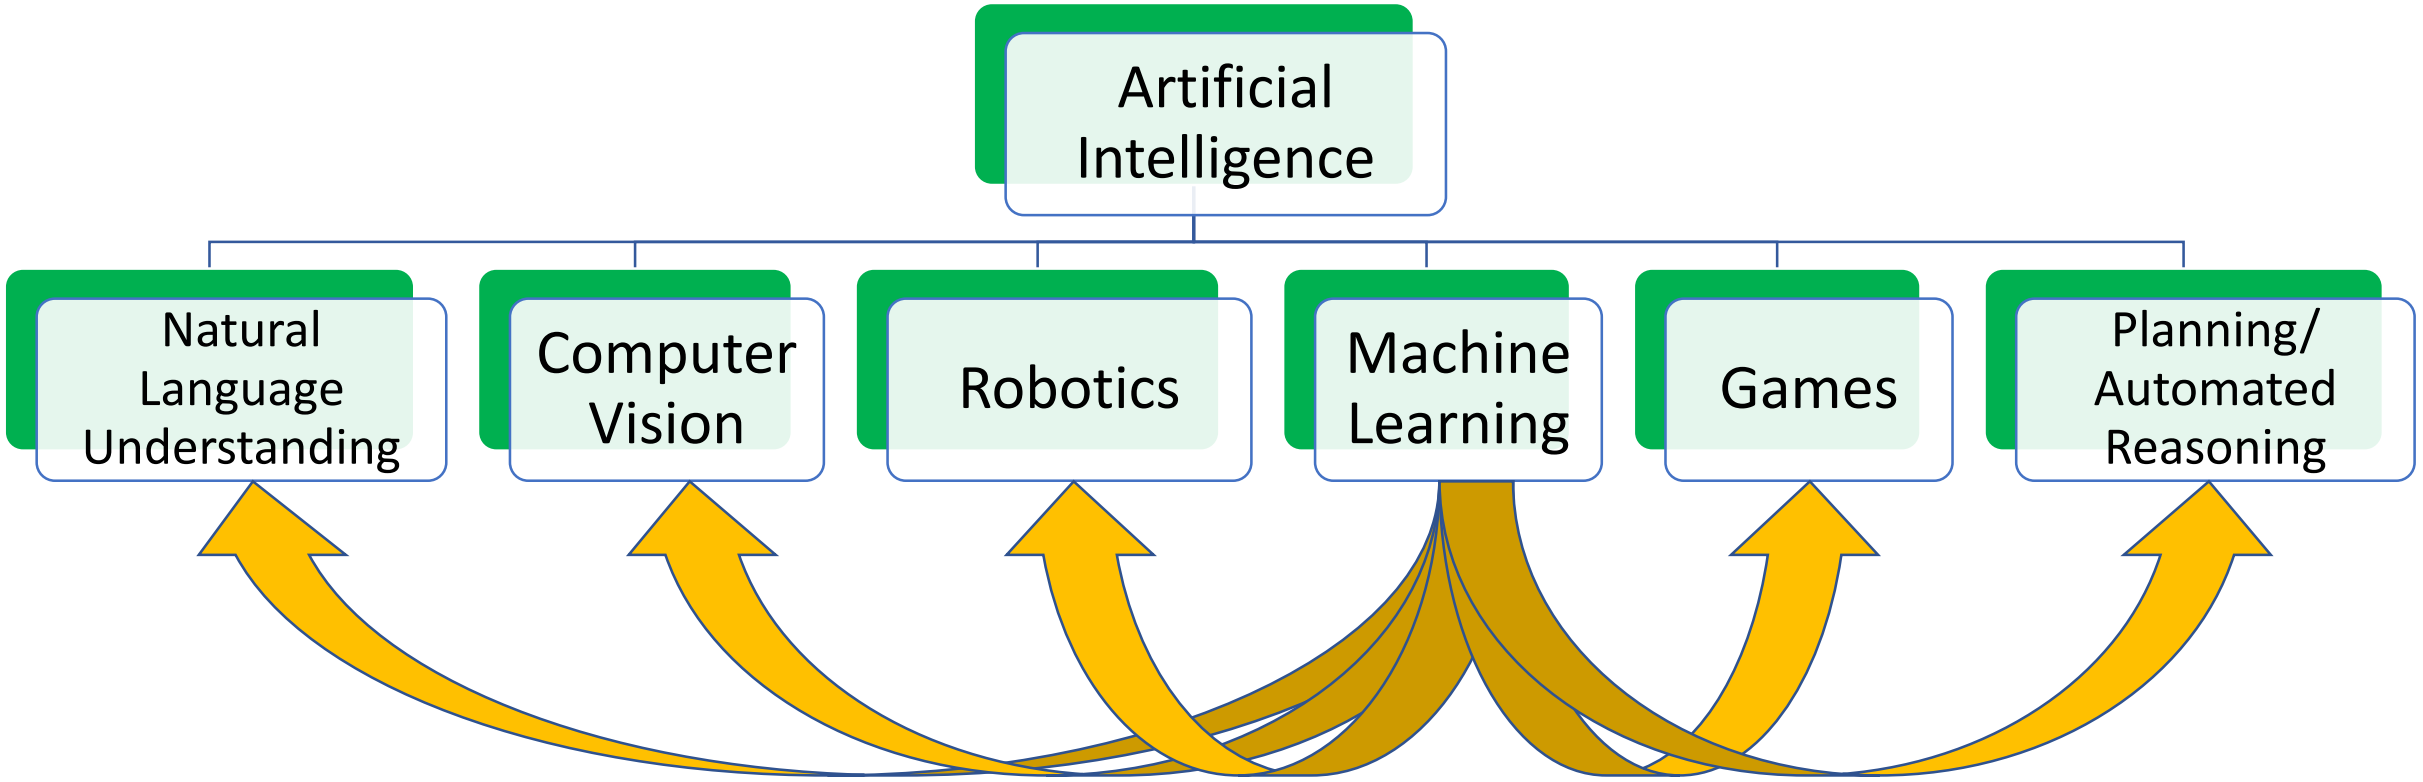
Artificial Intelligence
Natural Language Understanding
Computer Vision
Robotics
Machine Learning
Games
Planning/ Automated Reasoning

The AI Spectrum

Collateral “Successes”

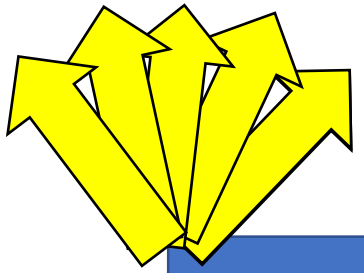

Narrow “cognitive” skills
“Weak AI”

Important successes

Broad capabilities
“Strong AI”

Fearmongering
Utopian idealism
Bad predictions
Science Fiction

Legitimate concerns

The AI Spectrum

Collateral “Successes”

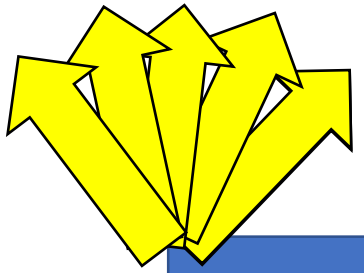

Non-AI “AI”

Narrow “cognitive” skills
“Weak AI”

Important successes

?

Legitimate concerns

Broad capabilities
“Strong AI”

Fearmongering
Utopian idealism
Bad predictions
Science Fiction

The AI Spectrum

Collateral “Successes”

All of these are called AI

Narrow “cognitive” skills
“Weak AI”

Broad capabilities
“Strong AI”

Important successes

Fearmongering
Utopian idealism
Bad predictions
Science Fiction

Legitimate concerns

The AI Spectrum

Collateral “Successes”

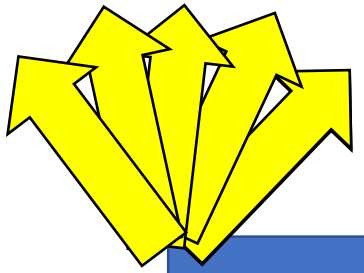

Non-AI “AI”

Narrow “cognitive” skills
“Weak AI”

Important successes

?

Legitimate concerns

Broad capabilities
“Strong AI”

Fearmongering
Utopian idealism
Bad predictions
Science Fiction

The AI Spectrum

Collateral “Successes” – many are also called AI

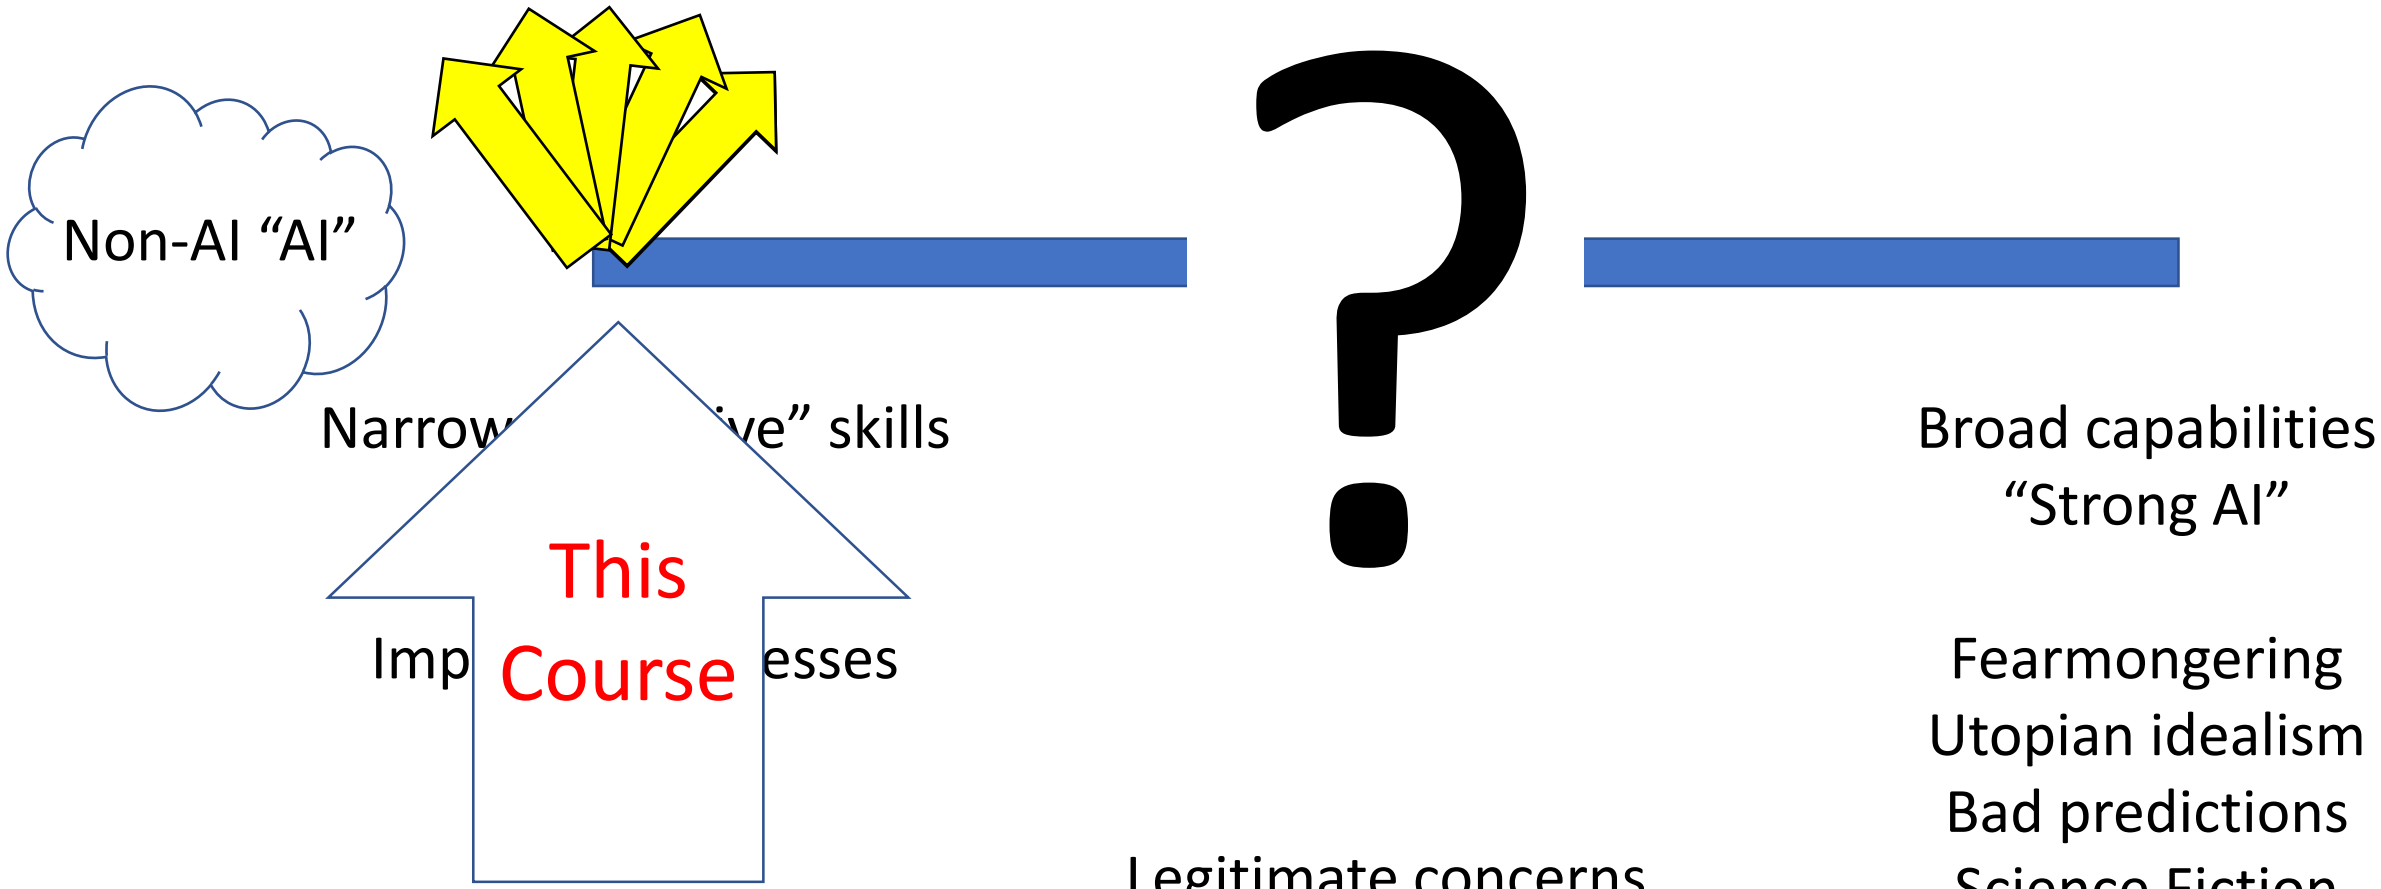

This Course

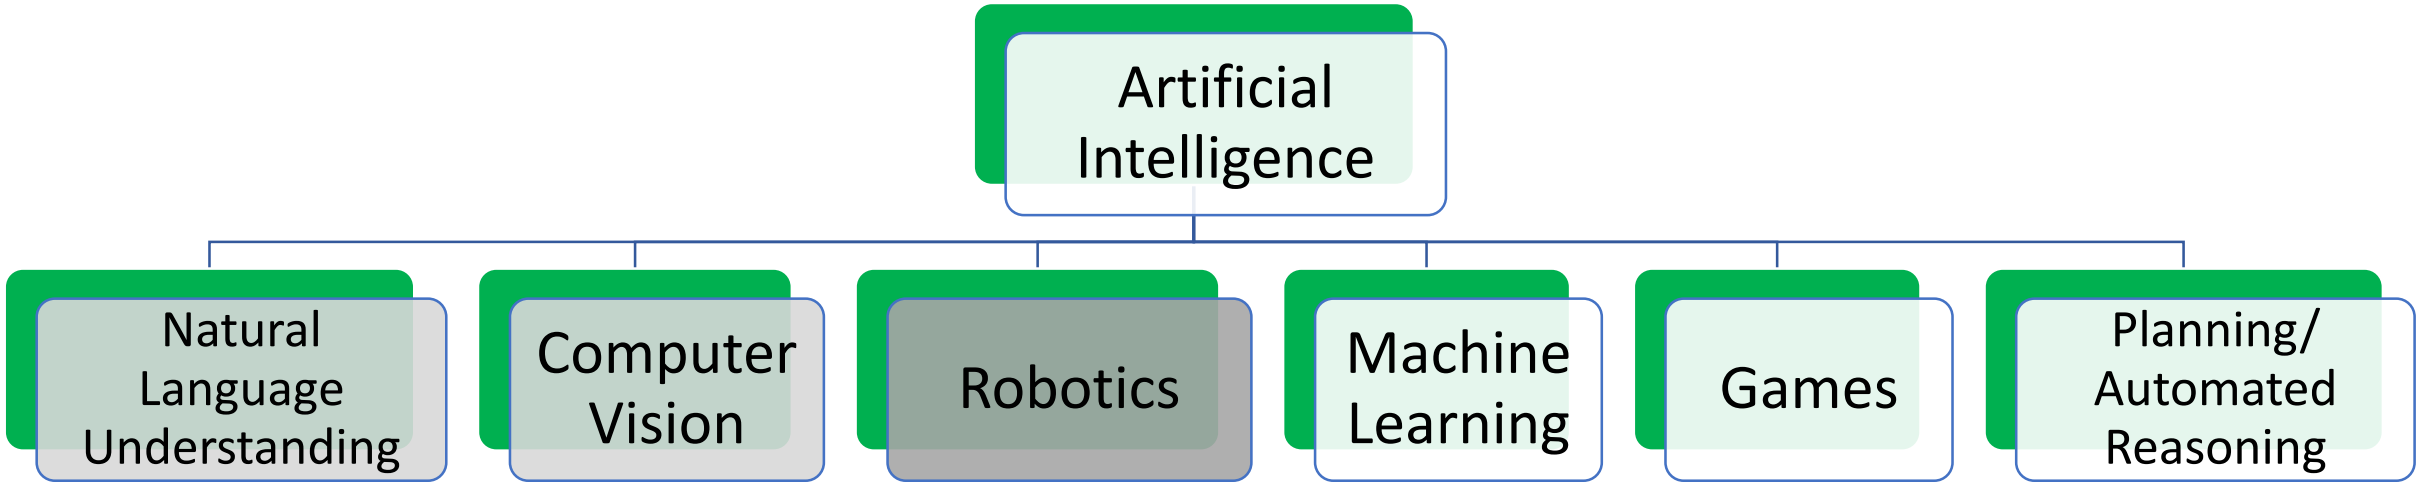
Artificial Intelligence
Natural Language Understanding
Computer Vision
Robotics
Machine Learning
Games
Planning/ Automated Reasoning

This Course

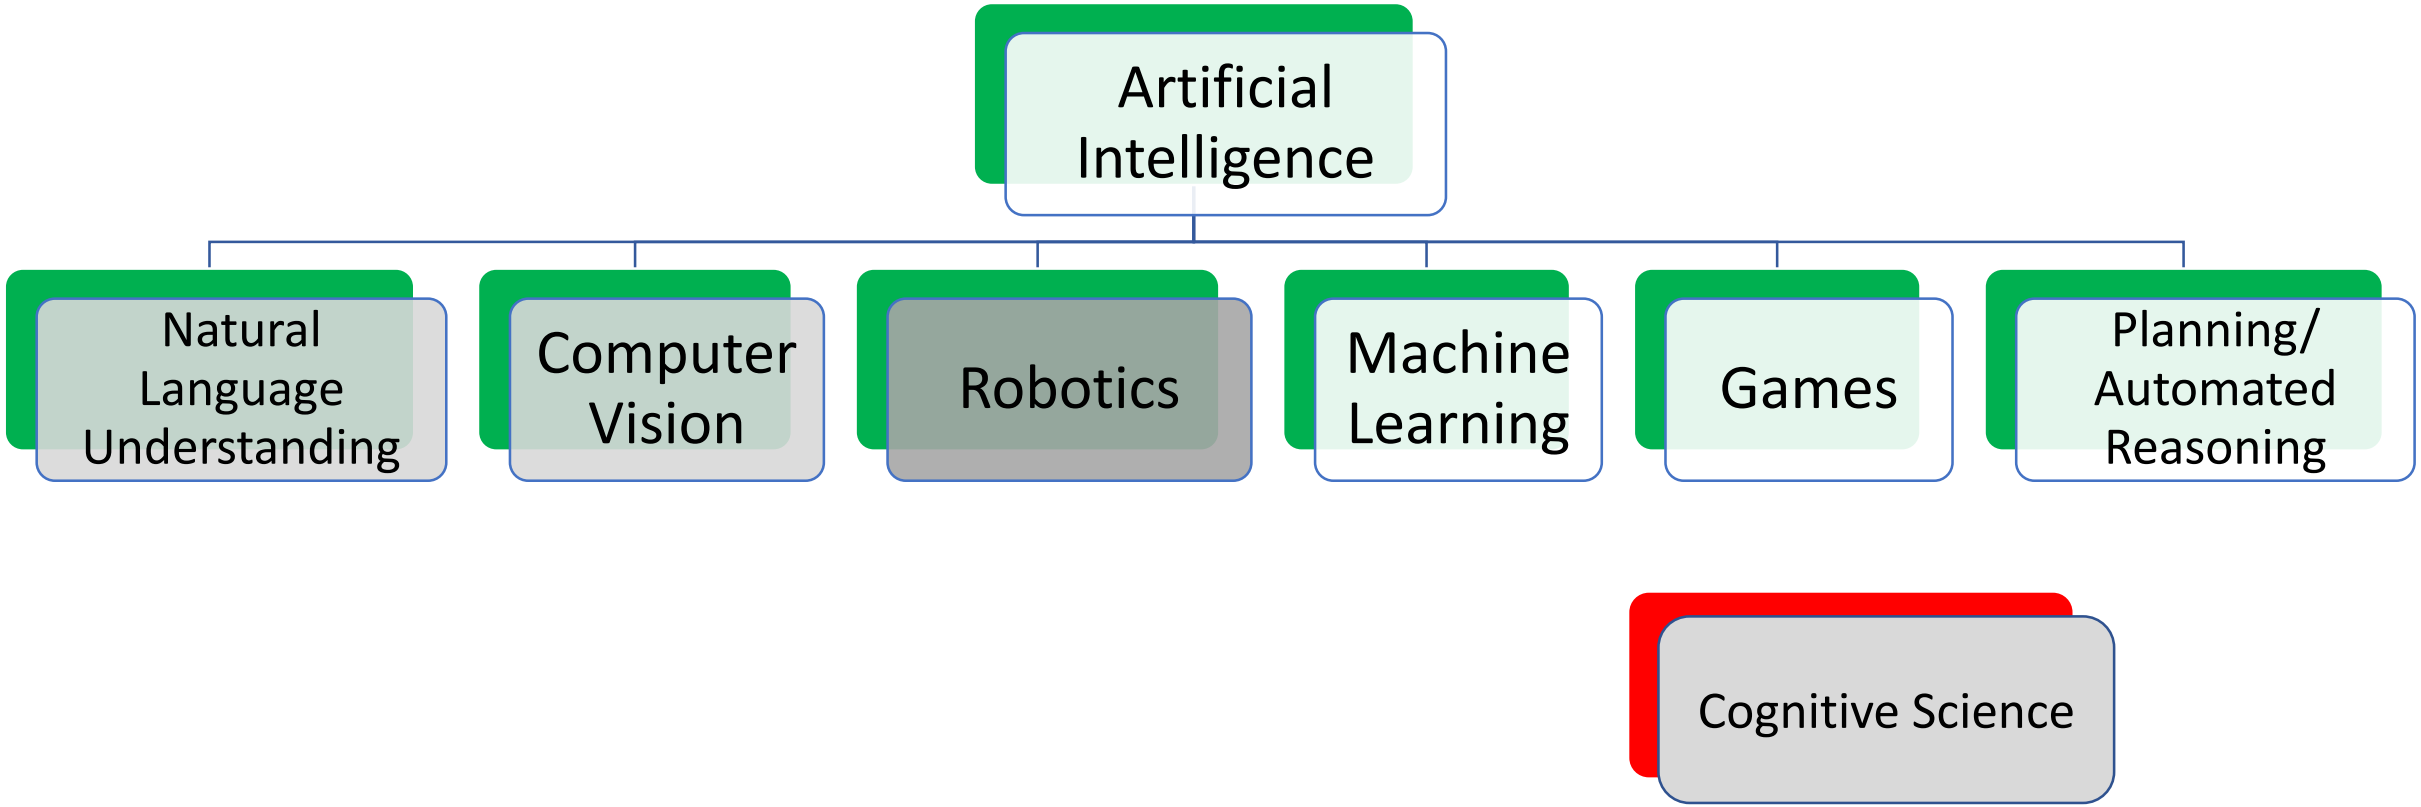
Artificial Intelligence
Natural Language Understanding
Computer Vision
Robotics
Machine Learning
Games
Planning/ Automated Reasoning
Cognitive Science

This Course

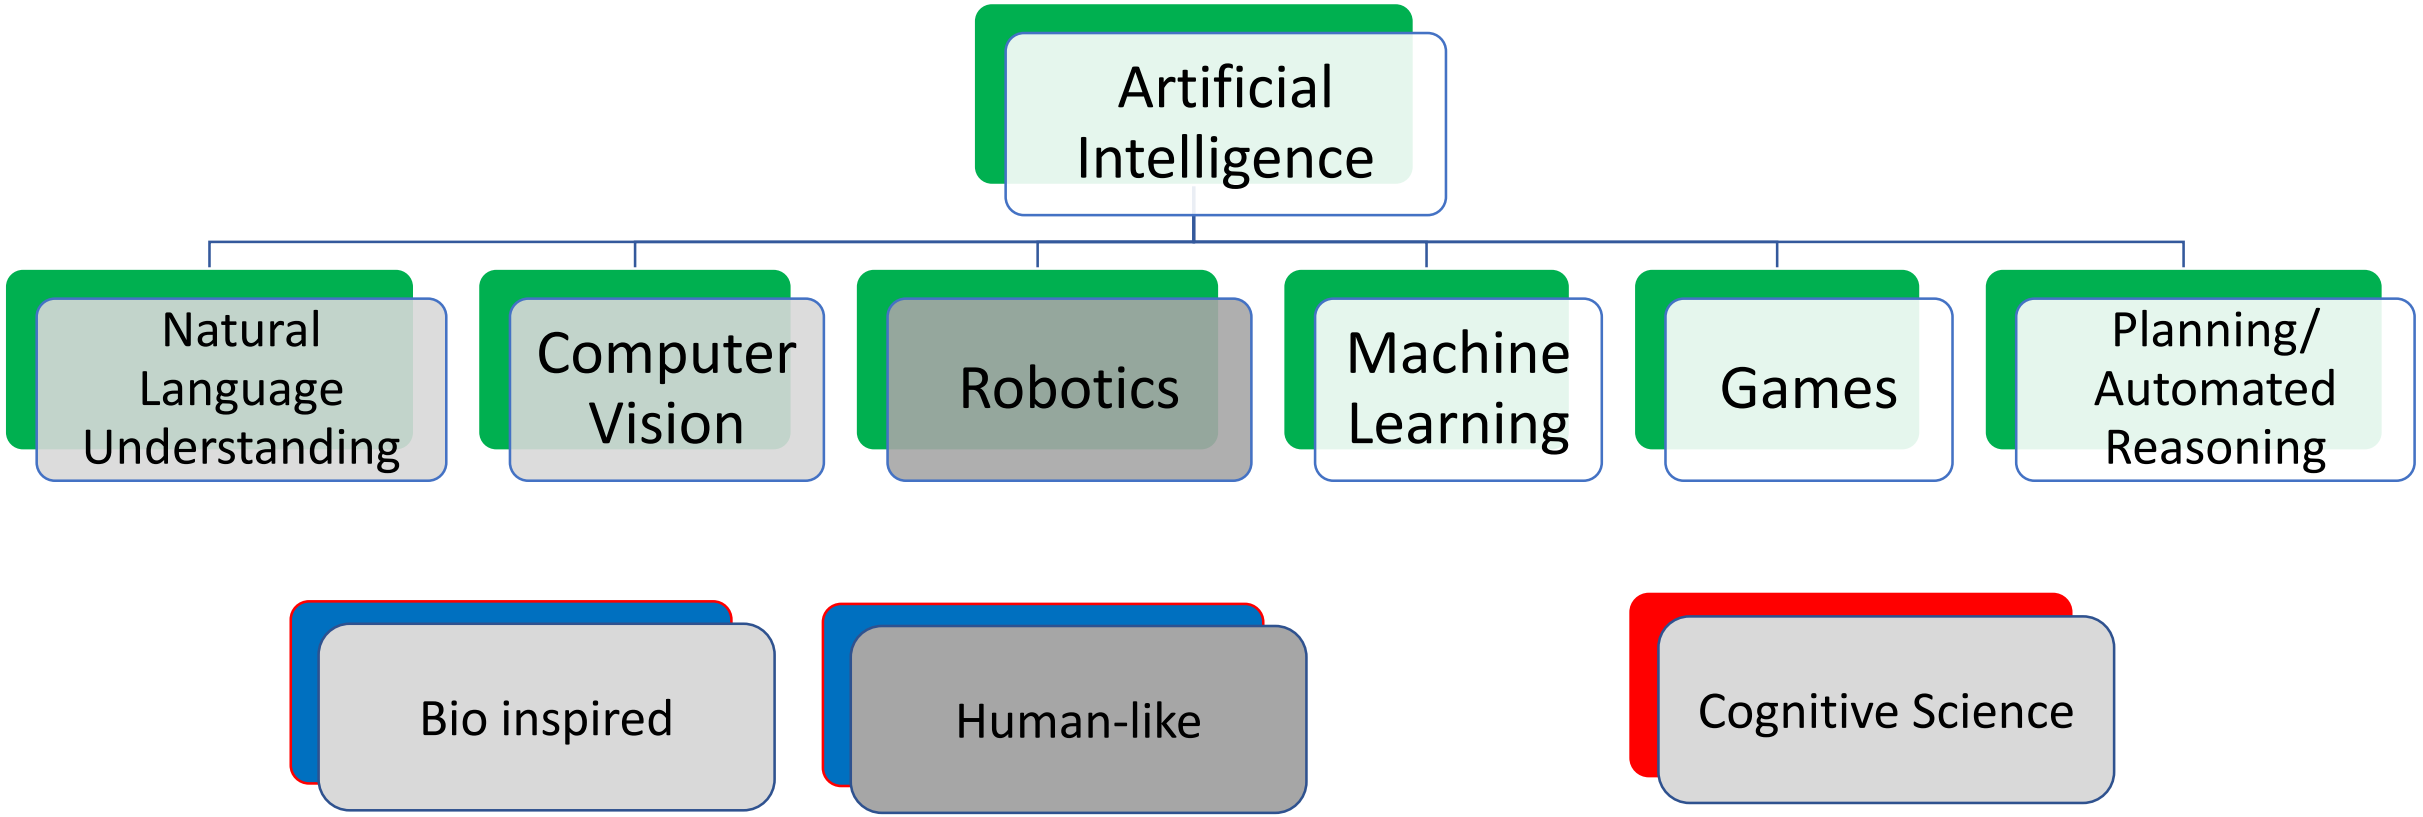
Artificial Intelligence
Natural Language Understanding
Computer Vision
Robotics
Machine Learning
Games
Planning/ Automated Reasoning
Bio inspired
Human-like
Cognitive Science

Today

Overview of AI

Overview of the course

Today

Overview of AI

Overview of the course

Course Details

- Instructor: Prof. Haym Hirsh, Gates 352 (Office Hours TBA)
- Head TA: Molly Feldman
- Course website: http://www.cs.cornell.edu/courses/cs4700/
- Course email: FAI-L@cornell.edu
- Discussions: Piazza
- Assignment submissions: Gradescope (Entry Code: 9B85X2)
- No official auditing option

Course Details

- Instructor: Prof. Haym Hirsh, Gates 352 (Office Hours TBA)
- Head TA: Molly Feldman
- Course website: http://www.cs.cornell.edu/courses/cs4700/
- Course email: FAI-L@cornell.edu
- Discussions: Piazza
- Assignment submissions: Gradescope (Entry Code: 9B85X2)
- No official auditing option

READ THE WEBSITE

Prerequisites

- CS 2110/ENGRD 2110
- CS 2800

- Main items:
  - Tree and graph algorithms
  - Probability
  - Propositional and first-order logic
  - Big-O notation
  - Ability to program

Grading

- 30%: Homeworks
- 30%: Prelim
  - March 17, 7:30pm, Baker Laboratory 200
  - This is the drop deadline
    – you will not know your prelim grade by then
- 40%: Final
  - May 13, 2:00pm

Deadline for
notification of conflicts:
February 4
Email FAI-L@cornell.edu

Homeworks

- ~6 over the semester (roughly every two weeks)
- Must be typeset (LaTeX, MS Word, etc.)
- Some will involve programming
- Submissions on Gradescope (Entry Code: 9B85X2)
- Late policy: Up to 2 days late for 50% credit (except for Homework 1)
- Collaboration policy: Writeup must be your own
  - Acknowledge collaborators, if in doubt please ask! – see website

Programming

- Python
  - This is a 4000 level course, if you don’t know it, learn it
  - It’s good for you to know

- Jupyter Notebooks
  - This is a 4000 level course, if you don’t know it, learn it
  - It’s good for you to know
  - Introduction later next week (TBA)

Grading

- Regrade requests:
  - Within 7 days, through Gradescope
  - Should be about mistakes in grading, not “why was this wrong”
  - Reserve the right to regrade other questions
  - Negative Karma (more on this in 4 slides)
- +/- 3%: Percentages will be adjusted to decrease whichever is lowest and increase whichever is highest by 3%

Technology Policy

No technology except for first five rows of left and right sides

(Unless as part of an in-class exercise)

Special Accommodations

Scan documentation letter and email to FAI-L@cornell.edu

Textbook

We will be using draft chapters of

Artificial Intelligence: A Modern Approach, 4th Edition
by Stuart Russell and Peter Norvig

Chapters will be available off of the course website

Feedback solicited

Karma Points

- Used for borderline students – a way of measuring engagement
- Examples of ways to earn karma points:
  - Answering questions and writing helpful posts on Piazza
  - Attending AI seminars that I share with you - you will be asked to submit a list of seminars that you attended at the end of the semester
  - Providing helpful feedback for the authors on the textbook
  - Submitting a course evaluation at the end of the semester.
- Examples of ways to lose karma points:
  - Posting a question on Piazza that has already been asked and answered
  - Emailed course staff a question that is already answered on the website
  - Abusing the regrade process – too many, unprofessional language

First Karma Lecture

Dylan Hadfield-Menell, UC Berkeley

Jan 28 11:40am-12:40pm
Gates G01

Attendance

- You are responsible for what is covered in class
  - If it’s in the book and not covered in class, it’s not required
  - If it’s in the slides and not covered in class, it not required
  - If it’s not in the book but covered in class, it’s required
    - even if it’s not on the slides

Readings

- Role of textbook: Augments what's covered in class

- Skim Chapters 1 and 2
  - Chapter 1: broad overview of AI complementary to this lecture
    - History
    - Relationship to other fields
  - Chapter 2: frames how the textbook presents subjects

Next lecture:

Search

Sections 3.1-3.4

Homework 1

- Due: Friday, Jan 31, 11:59pm
- No late submissions!
- Tests background material
- Graded in time for add deadline

If you’re not enrolled,
get on the waitlist

Permission numbers given on Tuesdays and Thursdays

I’m optimistic about people getting in