


---

# CSE 573 PMP: Artificial Intelligence

Hanna Hajishirzi

slides adapted from  
Emma Pierson, Dan Klein, Pieter Abbeel [ai.berkeley.edu](http://ai.berkeley.edu)  
And Dan Weld, Luke Zettlemoyer



# Resources

---

## Website

tentative schedule

lecture slides

course policies, etc.

<https://courses.cs.washington.edu/courses/csep573/25au/>


All resources (slides, notes, recordings, assignments, etc.) posted here

Ed: Discussion forum

Gradescope: Submit assignments here

# Course Staff (TAs)

---



**TA:** Tianhua Tao  
**Office hours:** W 7-8pm  
**Location:** TBD



**TA:** Hamish Ivison  
**Office hours:** TBD  
**Location:** TBD



**TA:** Mengyi Shan  
**Office hours:** M 6:30-7:30pm  
**Location:** TBD



**TA:** Jingwei Ma  
**Office hours:** Tu 7-8pm  
**Location:** TBD

# Course Staff

---

- Office hours
  - Schedule on the website
  - **TAs**: concepts, projects, homework
  - **Hanna**: concepts, high level guidance



# Course Format

---

## Programming Assignments (50%)

4 projects

Python

Autograded

Give you hands-on experience with the algorithms

I expect you to get 100% on projects

## ○ Written homeworks (50%)

- 2 written homeworks

- Gives you a more conceptual understanding of the material

# Prerequisites

---

Data Structure or Equivalent:  
CSE 332

Math:

Basic exposure to probability and data structures


Programming – Familiar with Python

There is a 0<sup>th</sup> project (P0)

# Textbook

Not required, but for students who want to read more we recommend

Russell & Norvig, AI: A Modern Approach, 3<sup>rd</sup> Ed.



Warning: Not a course textbook, so our presentation does not necessarily follow the presentation in the book.

# Course Policies

---

## Grade:

Your grade will be:

50% programming assignments,

50% written homeworks,

Assignments should be done individually unless otherwise specified.

Late Policy: **6** penalty-free late days for the whole quarter.


# What is AI?

---

What is AI?

What can it do?

What is this course?



# AI is having real-world impact

## Public imagination

Text assistants






# AI is having real-world impact

## Public imagination

Text assistants

Image generation



vibrant portrait painting of Salvador Dalí with a robotic half face



a shiba inu wearing a beret and black turtleneck



a close up of a handpalm with leaves growing from it



an espresso machine that makes coffee from human souls, artstation



panda mad scientist mixing sparkling chemicals, artstation



a corgi's head depicted as an explosion of a nebula

# AI is having real-world impact

---

## Public imagination

Text assistants

Image generation

Film/Game Industry






# AI is having real-world impact

Public imagination

Economy

\$4.8 trillion by 2033  
(25x increase in a decade)



<https://unctad.org/publication/technology-and-innovation-report-2025>

# AI is having real-world impact

Public imagination

Economy

Politics



# AI is having real-world impact

Public imagination

Economy

Politics



<https://www.nytimes.com/2025/02/11/world/europe/vance-speech-paris-ai-summit.html>



# AI is having real-world impact

Public imagination


Economy

Politics

The New York Times

OPINION

## How A.I. Chatbots Become Political



# AI is having real-world impact

Public imagination

Economy

Politics

Law

Aug. 18, 2023, 12:18 PM; Updated: Aug. 18, 2023, 12:48 PM

## AI-Generated Art Lacks Copyright Protection, D.C. Court Says (1)



**Riddhi Setty**  
Reporter



**Isaiah Poritz**  
Legal Reporter



Bloomberg Law, 2023

# AI is having real-world impact

Public imagination

Economy

Politics

Law

Labor

## The human labor behind AI chatbots and other smart tools

Data labeling is an important step in developing artificial intelligence but also exposes the people doing the work to harmful content.

MarketWatch, 2023



# AI is having real-world impact

---

Public imagination

Economy

Politics

Law

Labor

Science

Chemistry and biology

Science & technology | The 2024 Nobel prizes

## AI wins big at the Nobels

Awards went to the discoverers of micro-RNA, pioneers of artificial-intelligence models and those using them for protein-structure prediction

The Economist, 2024

# AI is having real-world impact

Public imagination

Economy

Politics

Law

Labor


Science

Physics

AMIT KATWALA SCIENCE FEB 16, 2022 11:00 AM

## DeepMind Has Trained an AI to Control Nuclear Fusion

The Google-backed firm taught a reinforcement learning algorithm to control the fiery plasma inside a tokamak nuclear fusion reactor.



PHOTOGRAPH: CURDIN WÜTHRICH, SPC/EPFL

Wired, 2022



# AI is having real-world impact

Public imagination

Economy

Politics

Law

Labor

Science

Conservation



Nature Communications, 2022

# AI is having real-world impact

---

Public imagination

Economy


Politics

Law

Labor

Science

Resource use



The New York Times, 2025

# AI is having real-world impact

Public imagination

Economy

Politics

Law

Labor

Science

Health



The Boston Globe, 2018

# AI is having real-world impact

Public imagination

Economy

Politics

Law

Labor

Science

Health

Education

BREAKING

## ChatGPT In Schools: Here's Where It's Banned—And How It Could Potentially Help Students

**Arianna Johnson** Forbes Staff

*I cover the latest trends in science, tech and healthcare.*

Follow

 2

Jan 18, 2023, 02:31pm EST

Forbes, 2023

# AI is having real-world impact

---

Public imagination

Economy

Politics

Law

Labor

Science

Health

Education

Social interaction

**The Washington Post**  
*Democracy Dies in Darkness*

---

## **‘It’s almost like we never even spoke’: AI is making everyone on dating apps sound charming**

As more people use AI to doctor their messages, others are left wondering exactly who they’re falling for.

Updated July 3, 2025

# AI is having real-world impact

---

Public imagination

Economy

Politics

Law

Labor

Science

Health

Education


Social interaction

Ok, but what actually is AI???

# What should we build?

---

Should we make machines that...





# Rational Decisions


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

Rational: *maximally achieving pre-defined goals*

Goals are expressed in terms of the **utility** of outcomes

World is uncertain, so we'll use **expected** utility

Being rational means acting to **maximize your expected utility**





# Rational Decisions

---

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

Rational: *maximally achieving pre-defined goals*

Goals are expressed in terms of the **utility** of outcomes

World is uncertain, so we'll use **expected** utility

Being rational means acting to **maximize your expected utility**

A better title for this course might be:

**Computational Rationality**

# Perspectives on Intelligence

---

## Skills-based perspective

“A system is only intelligent if it can do [X].”

Play chess?

Learn from experience?

Use words properly?

Make mistakes?

Not make mistakes?

# What About the Brain?

Brains (human minds) are very good at making rational decisions, but not perfect

Brains aren't as modular as software, so hard to reverse engineer!

AI may be better than brains at some tasks

*“Brains are to intelligence as wings are to flight”*

We can't yet build AI on the scale of the brain

~100T synapses in the human brain vs ~1.8T weights in GPT4

Still, the brain can be a great inspiration for AI!



# A (Short) History of AI

## 1940-1950: Early days: neural and computer science meet

1943: McCulloch & Pitts: Perceptron–boolean circuit model of brain


1950: Turing's "Computing Machinery and Intelligence"

## 1950—70: Excitement! Logic-driven

1950s: Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine

1956: Dartmouth meeting: "Artificial Intelligence" adopted

*"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."*



# A (Short) History of AI

## 1940-1950: Early days: neural and computer science meet

1943: McCulloch & Pitts: Perceptron–boolean circuit model of brain

1950: Turing's "Computing Machinery and Intelligence"

## 1950—70: Excitement! Logic-driven

1950s: Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine

1956: Dartmouth meeting: "Artificial Intelligence" adopted

1969: Minsky & Papert: perceptrons can't learn XOR/parity!

## 1970—90: Knowledge-based approaches

1969—79: Early development of knowledge-based systems

1980—88: Expert systems industry booms; backpropagation makes it feasible to train multi-layer neural networks

1988—93: Expert systems industry busts: "AI Winter"

## 1990—2010: Statistical approaches, agents


Resurgence of probability, focus on uncertainty

Agents and learning systems... "AI Spring"?

1992: TD-Gammon achieves human-level play at backgammon

1997: Deep Blue defeats Gary Kasparov at chess

2002: Embodied AI; Roomba vacuum invented




# A (Short) History of AI

## 2010—2017: Big Data, GPUs, Deep Learning

- 2011: Apple releases Siri
- 2012: AlexNet wins ImageNet competition
- 2015: DeepMind achieves “human-level” control in Atari games
- 2016: DeepMind’s AlphaGo defeats Lee Sedol at Go
- 2016: Google Translate migrates to neural networks

## 2017—: Scaling Up, Large Language Models

- 2017: Google invents Transformer architecture
- 2017: DeepStack/Libratus defeat humans at poker
- 2018-2020: AlphaFold predicts protein structure from amino acids
- 2021-2022: Modern text-to-image generation
- 2022: OpenAI releases ChatGPT
- 2023: Every other company also releases a chatbot
- 2024: Nobel prizes in physics and chemistry go to AI advances



# A (Short) History of AI

---

## Zooming out, some general patterns:

Overclaiming followed by disillusionment (AI winters)

Simultaneous discovery, or rediscovery, of ideas

Interdisciplinary


Math, statistics, control theory, neuroscience, psychology, economics,  
philosophy, computer engineering, linguistics

And increasingly: law, medicine, education...

# What Can AI Do?

**Quiz:** Which of the following can be done at present?

- ✓ Win against any human at chess?
- ✓ Win against the best humans at Go?
- Play a decent game of table tennis?
- Unload any dishwasher in any home?
- Drive safely along the highway?
- Drive safely along streets of San Francisco?
- Buy a week's worth of groceries on the web?
- Buy a week's worth of groceries at UDistrict Farmer's Market?
- Discover and prove a new mathematical theorem?
- Perform a surgical operation?
- Translate spoken Chinese into spoken English in real time?
- Win an art competition?
- Write an intentionally funny story?
- Construct a building?





# What Can AI Do?

**Quiz:** Which of the following can be done at present?

- ✓ Win against any human at chess?
- ✓ Win against the best humans at Go?
- Play a decent game of table tennis?
- Unload any dishwasher in any home?
- Drive safely along the highway?
- Drive safely along streets of San Francisco?
- Buy a week's worth of groceries on the web?
- Buy a week's worth of groceries at UDistrict Farmer's Market?
- Discover and prove a new mathematical theorem?
- Perform a surgical operation?
- Translate spoken Chinese into spoken English in real time
- Win an art competition?
- Write an intentionally funny story?
- Construct a building?

Google's AlphaGo beats Lee Sedol at Go in 2016



But ... a plot twist in 2023!

DEEP BLUE WAS JUST THE START —

## Man beats machine at Go in human victory over AI

Amateur exploited weakness in systems that have otherwise dominated grandmasters.

RICHARD WATERS, FINANCIAL TIMES - 2/19/2023, 4:51 AM

# What Can AI Do?

**Quiz:** Which of the following can be done at present?

- ✓ Win against any human at chess?
- ✓ Win against the best humans at Go?
- ✓ Play a decent game of table tennis?
- Unload any dishwasher in any home?
- Drive safely along the highway?
- Drive safely along streets of San Francisco?
- Buy a week's worth of groceries on the web?
- Buy a week's worth of groceries at UDistrict Farmer's Market?
- Discover and prove a new mathematical theorem?
- Perform a surgical operation?
- Translate spoken Chinese into spoken English in real time?
- Win an art competition?
- Write an intentionally funny story?
- Construct a building?



# What Can AI Do?

**Quiz:** Which of the following can be done at present?

- ✓ Win against any human at chess?
- ✓ Win against the best humans at Go?
- ✓ Play a decent game of table tennis?
- ✗ Unload any dishwasher in any home?
- ✓ Drive safely along the highway?
- ? Drive safely along streets of San Francisco?
- Buy a week's worth of groceries on the web?
- Buy a week's worth of groceries at UDistrict Farmer's Market?
- Discover and prove a new mathematical theorem?
- Perform a surgical operation?
- Translate spoken Chinese into spoken English in real time?
- Win an art competition?
- Write an intentionally funny story?
- Construct a building?

## California just opened the floodgates for self-driving cars

After a state regulatory board ruling, San Francisco will have 24/7 robotaxis on its streets. A Washington Post analysis shows how it's a pivotal moment for the industry.

By [Trisha Thadani](#) and [Jeremy B. Merrill](#)

Updated August 10, 2023 at 10:06 p.m. EDT | Published August 10, 2023 at 9:15 p.m. EDT

**The Register**

## California DMV hits brakes on Cruise's SF driverless fleet after series of fender benders

50% chop effective 'immediately' as department investigates traffic and safety issues

[Jude Karabus](#)

Mon 21 Aug 2023 // 14:28 UTC

**MONEYWATCH**

## Waymo recalls roughly 1,200 self-driving vehicles prone to hitting road barriers

**MONEYWATCH**

By [Anne Marie D. Lee](#)  
Edited By [Alain Sherter](#)  
May 14, 2025 / 4:36 PM EDT / CBS News



# What Can AI Do?

**Quiz:** Which of the following can be done at present?

- ✓ Win against any human at chess?
- ✓ Win against the best humans at Go?
- ✓ Play a decent game of table tennis?
- ✗ Unload any dishwasher in any home?
- ✓ Drive safely along the highway?
- ? Drive safely along streets of San Francisco?
- ✓ Buy a week's worth of groceries on the web?
- ✗ Buy a week's worth of groceries at UDistrict Farmer
- ? Discover and prove a new mathematical theorem?
- ✗ Perform a surgical operation?
- ✓ Translate spoken Chinese into spoken English in real time?
- ✓ Win an art competition?
- Write an intentionally funny story?
- Construct a building?





# What Can AI Do?

**Quiz:** Which of the following can be done at present?

- ✓ Win against any human at chess?
- ✓ Win against the best humans at Go?
- ✓ Play a decent game of table tennis?
- ✗ Unload any dishwasher in any home?
- ✓ Drive safely along the highway?
- ? Drive safely along streets of San Francisco?
- ✓ Buy a week's worth of groceries on the web?
- ✗ Buy a week's worth of groceries at UDistrict Farmer's Market?
- ? Discover and prove a new mathematical theorem?
- ✗ Perform a surgical operation?
- ✓ Translate spoken Chinese into spoken English in real time?
- ✓ Win an art competition?
- ? Write an intentionally funny story?
- Construct a building?

please write me a very funny, and extremely short story

A penguin walks into a bar, looks around, and says,  
"Has anyone seen my brother?"

The bartender squints and asks,  
"What does he look like?"

ChatGPT, 2025

# What Can AI Do?

**Quiz:** Which of the following can be done at present?

- ✓ Win against any human at chess?
- ✓ Win against the best humans at Go?
- ✓ Play a decent game of table tennis?
- ✗ Unload any dishwasher in any home?
- ✓ Drive safely along the highway?
- ? Drive safely along streets of San Francisco?
- ✓ Buy a week's worth of groceries on the web?
- ✗ Buy a week's worth of groceries at UDistrict Farmer's Market?
- ? Discover and prove a new mathematical theorem?
- ✗ Perform a surgical operation?
- ✓ Translate spoken Chinese into spoken English in real time?
- ✓ Win an art competition?
- ? Write an intentionally funny story?
- ✗ Construct a building?





# This Course: Designing Rational Agents

An **agent** is an entity that perceives and acts.



A **rational agent** selects actions that maximize its (expected) **utility**.

Characteristics of the **percepts**, **environment**, and **action space** dictate techniques for selecting rational actions



This course is about:

General AI techniques for a variety of problem types

Learning to recognize when and how a new problem can be solved with an existing technique



# Pac-Man as an Agent





# Course Topics

---

Core Components of Rational Agents:

Search &  
Planning

Probability &  
Inference

Supervised  
Learning

Reinforcement  
Learning


# Course Topics

Search &  
Planning

Probability &  
Inference

Supervised  
Learning

Reinforcement  
Learning



How can I use my ***model*** of the world to find a ***sequence of actions*** to achieve my ***goal***?


# Course Topics

Search &  
Planning

Probability &  
Inference

Supervised  
Learning

Reinforcement  
Learning



How can I make sense of *uncertainty*?




# Course Topics

Search &  
Planning

Probability &  
Inference

Supervised  
Learning

Reinforcement  
Learning



How can I learn a ***model*** of the world from ***data***?


# Course Topics

Search &  
Planning

Probability &  
Inference

Supervised  
Learning

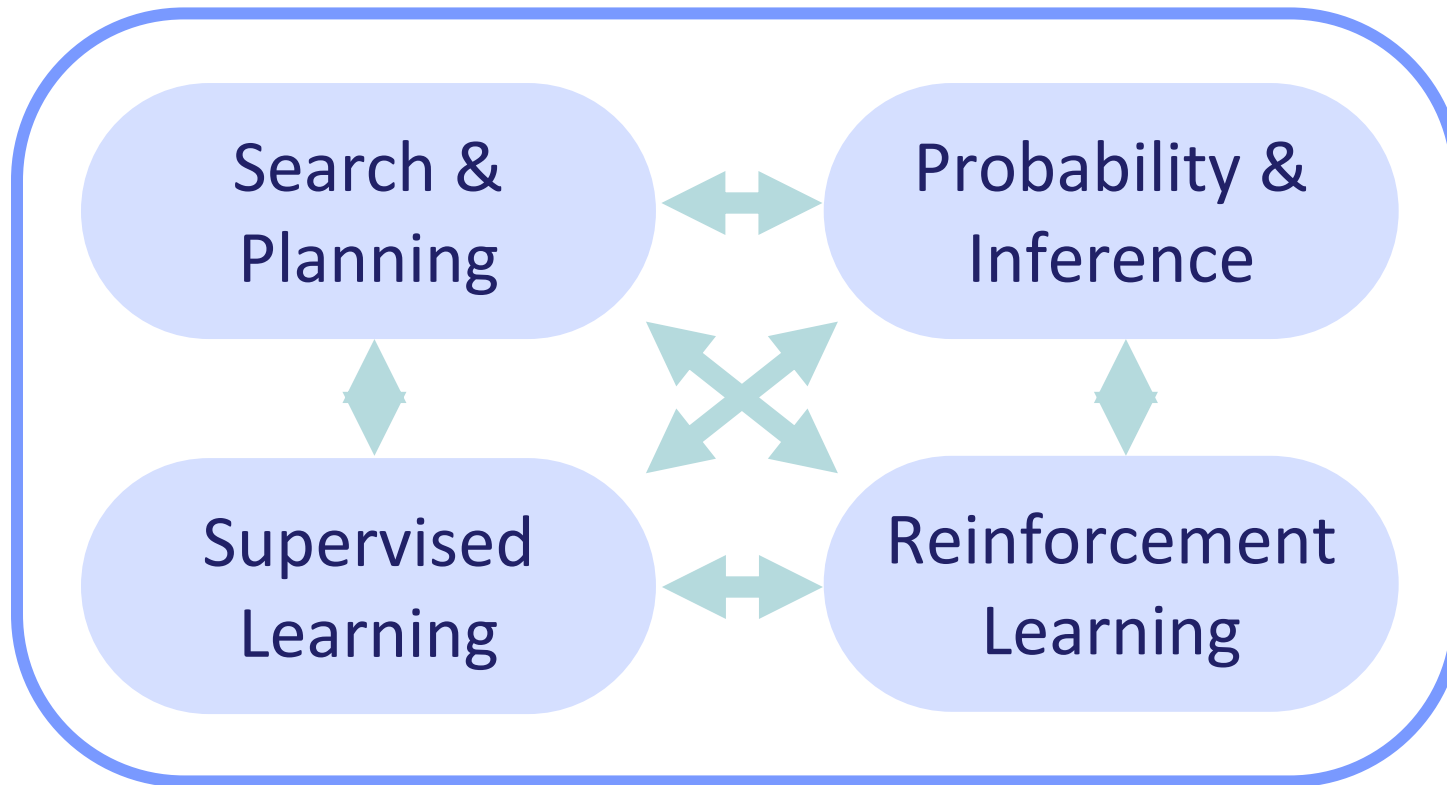
Reinforcement  
Learning



How can I learn a ***policy*** for any situation  
so that I can ***maximize utility***?

# Course Topics

---



# Course Topics

---

Search &  
Planning

Probability &  
Inference

Supervised  
Learning

Reinforcement  
Learning

Applications

Impact on Sciences, Technology, Society

# Topics in This Course

---

## Part I: Intelligence from Computation

- Fast search

- Adversarial and uncertain search

## Part II: Reasoning under Uncertainty

- Decision theory: Reinforcement Learning, Markov Decision Processes

- Machine learning

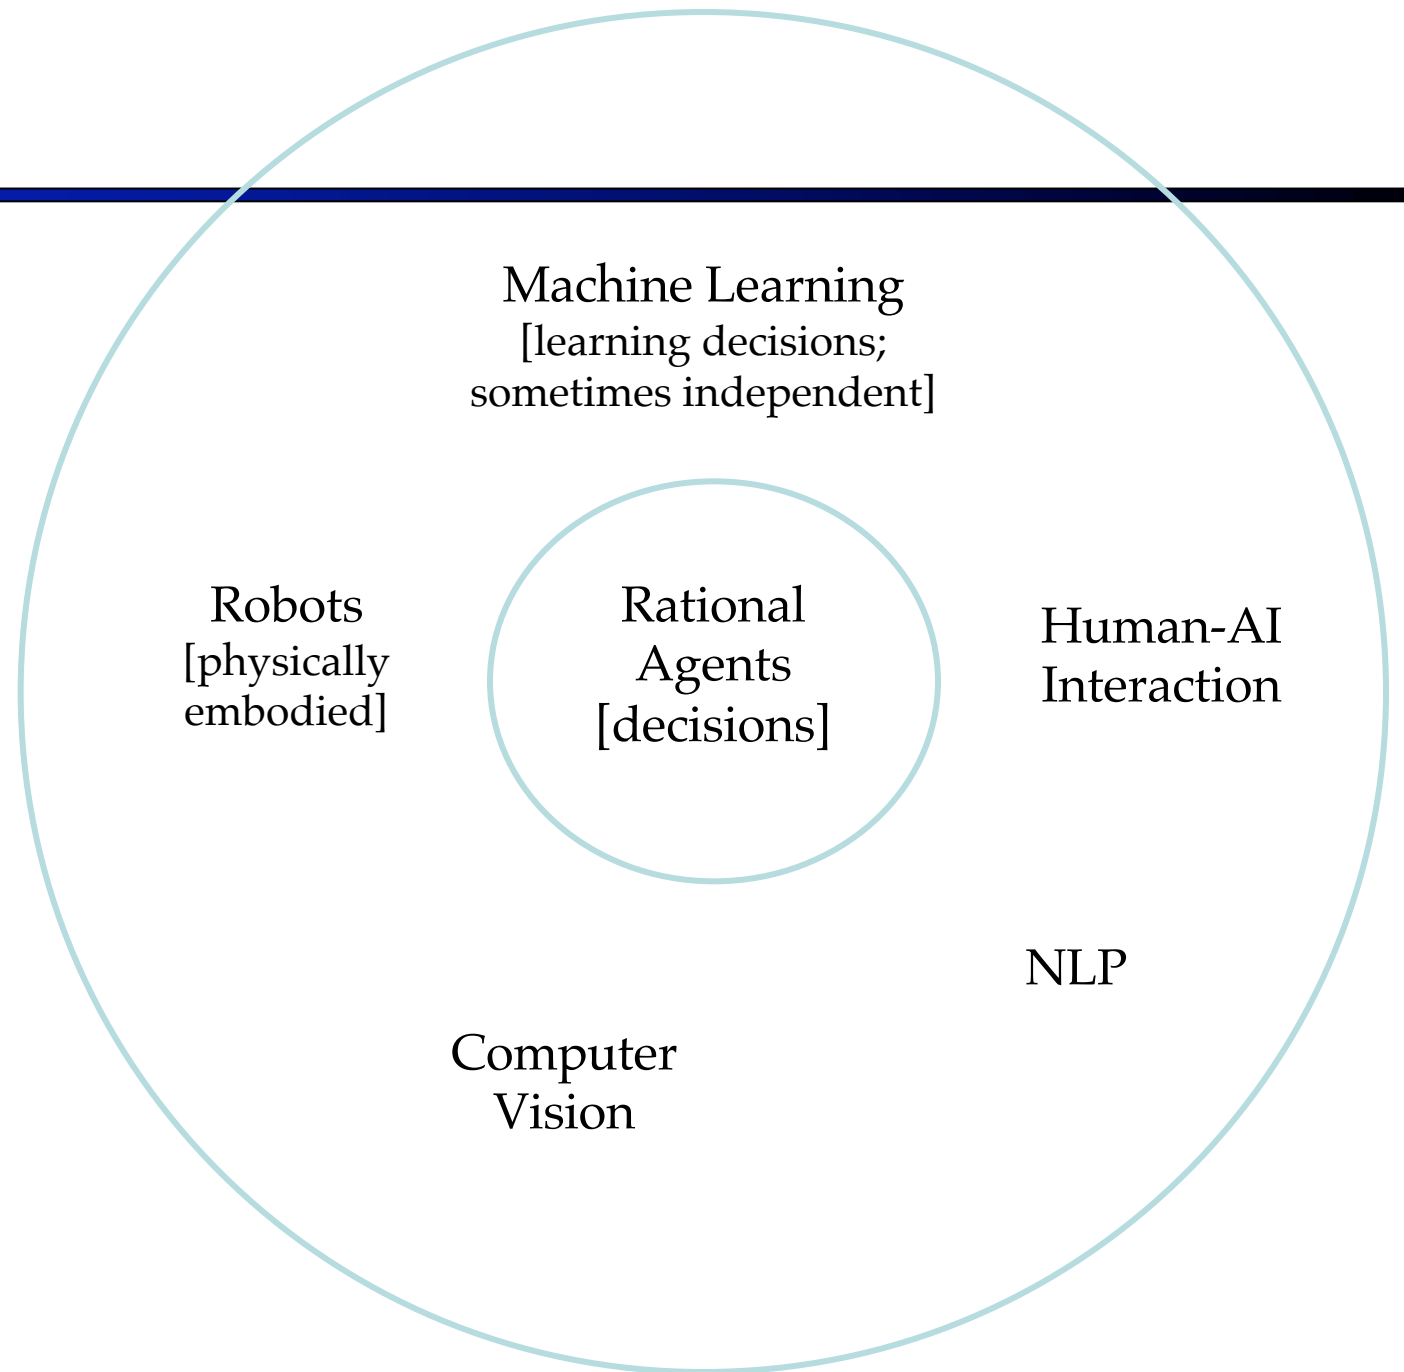
- Graphical Models - Bayes Nets; HMMs

## Throughout: Applications

- Natural language, vision, robotics, games, ...

---

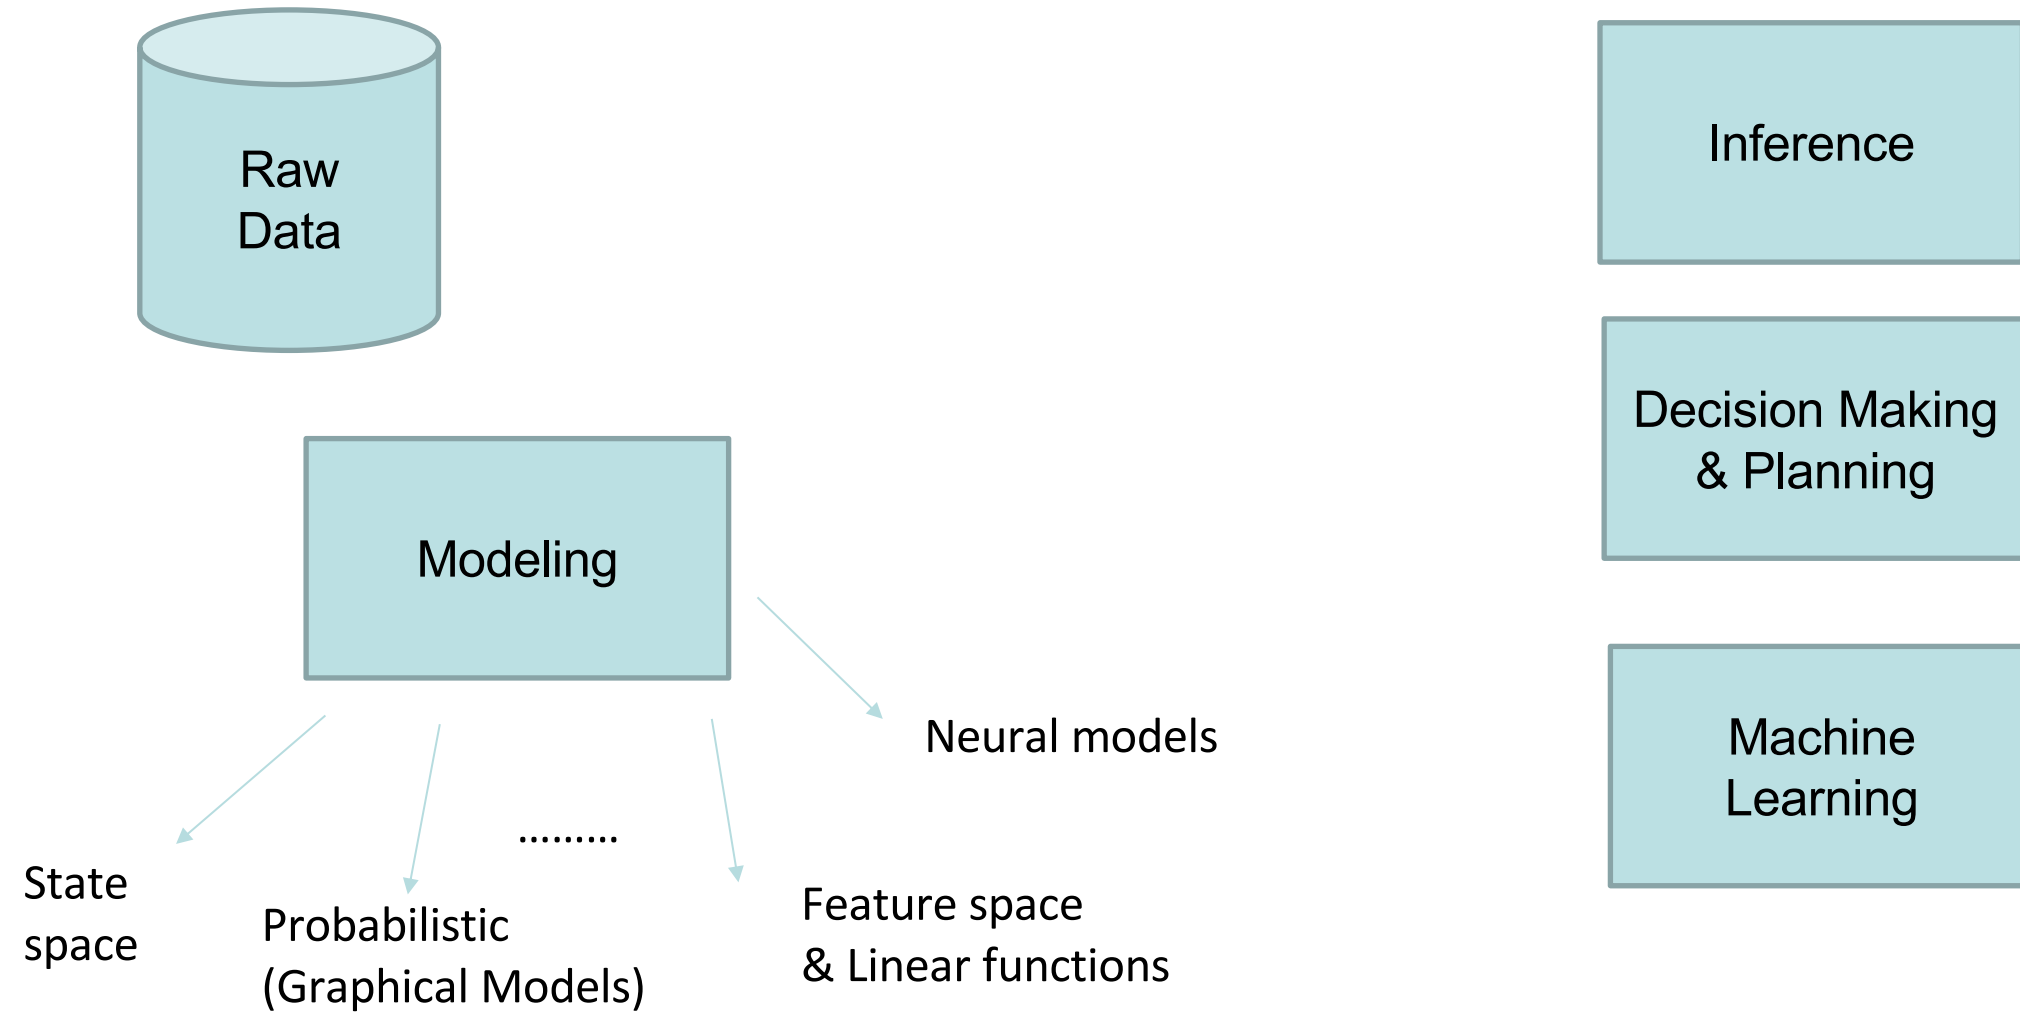
# This course w.r.t. Current AI Research





# This course w.r.t. Current AI Research

---



# Should I take CSEP 573?

---

Yes, if you want to know how to design rational agents!

CSEP 573 gives you extra mathematical and conceptual maturity

CSEP 573 gives you a survey of other non-CS fields that interact with AI  
(e.g. robotics, cognitive science, economics)

# By the end of this course you'll:


---

Build and understand math of rational, learning agents

Select and apply the right AI methods for wide range of problems

Recognize how these methods are used in modern AI systems

Be prepared to make decisions on how AI is used in society



# Important This Week

---

- Important this week:
  - **Check out** canvas--- our main resource for assignments and grades
  - **Check out** website – for schedule and slides
  - **Check out** Ed – for discussions; we have added everyone to Ed
  - **Check out** Gradescope -- for written assignments
  - **P0: Python tutorial** is out

# Next Lecture: Search

---

