# Artificial Intelligence

CS 165A

Apr 4, 2023

Prof. Yu-Xiang Wang

Computer Science

About myself

Yu-Xiang Wang 王宇翔

Eugene Aas Chair Assistant Professor

Department of Computer Science,
and Department of Statistics (By Courtesy)

Co-Director, Center for Responsible Machine Learning

UC Santa Barbara

Office: Henley Hall 2013

E-mail: yuxiangw AT cs.ucsb.edu

Curriculum Vita: [link]

Yu-Xiang is pronounced approximately as [' juː'ʃiːʌŋ],
namely, *y~eu~ee – sh~ih~ah~ng* .

Follow me on Twitter @yuxiangw_cs

Research area: Statistical Machine Learning. Optimization, reinforcement learning, differential privacy, deep learning.

Short biography:

China => Singapore
⇒ PhD from Carnegie Mellon University
⇒ Scientist at Amazon AI
⇒ Professor at UCSB

Homepage:
https://cs.ucsb.edu/~yuxiangw/

Call me: Yu-Xiang or Professor Wang. **Don't** call me: Wang, Yu, or Professor.

Teaching Assistants

- Esha Singh
  - esingh@ucsb.edu

- Vihaan Akshaay Rajendiran
  - vihaanakshaay@ucsb.edu

- Zihan Ma
  - zihan_ma@ucsb.edu

Roadmap today

- Introduction (AIMA Ch 1)
  - Why should you care?
  - History of AI
  - Example of the State-of-the-Art

- Course logistics
  - Schedule, grades, evaluation

- Intelligent Agents (AIMA Ch 2)

Artificial Intelligence

- AI in the media
  - Popular movies
    - 2001: A Space Odyssey
    - Star Wars
    - The Terminator
    - The Matrix
    - Artificial Intelligence: AI
    - and many …
  - Popular press, novels
- Often portrayed as
  - A property of evil computers
  - Computers doing impossible things
- Public view
  - Books and movies have inspired many AI researchers
  - Books and movies have raised the public's expectations

5

AI / Machine Learning has revolutionized almost every aspect of our daily life
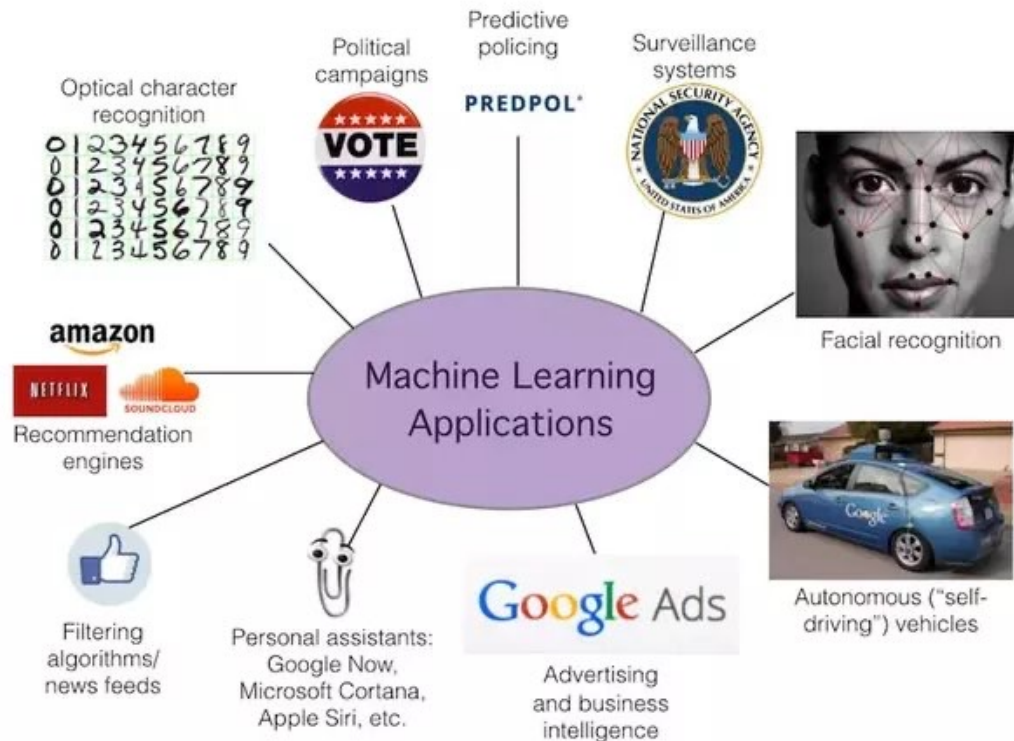
AI Is Transforming The Industries

☐ AI has transformed the IT industry
- ☐ Search Engine
- ☐ Speech Recognition
- ☐ Machine Translation
- ☐ Recommendation

☐ AI is transforming other industries
- ☐ Transportation
- ☐ Healthcare
- ☐ Finance
- ☐ Insurance, Law, HR, Travel, Media, …
- ☐ Semiconductor / Microprocessors

Why should you learn AI?

- Career opportunities
  - AI / ML jobs are highly paid
  - AI and ML are becoming the standard tools all software engineers are expected to know.

- Research opportunities and potential impact
  - "The golden age of physics is Newton's time and Einstein's time, the golden age of AI is right now!"

- Personal development
  - Consolidate your knowledge, connect the dots
  - Becoming better in solving problems

- "It is just my passion! "

Artificial Intelligence

□ Turing Test

The "standard interpretation" of the Turing Test, in which player C, the interrogator, is given the task of trying to determine which player – A or B – is a computer and which is a human. The interrogator is limited to using the responses to written questions to make the determination.

□ Artificial Intelligence is an ultimate dream of computer science

Supplementary reading: Turing (1950) "Computing Machinery and Intelligence" [here].

John McCarthy

Herb Simon

Allen Newell

Marvin Minsky

Seminal event for AI as a field, in 1956:
The **Dartmouth Summer Research Conference 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."

Over-Confidence

Herbert Simon (1957)

It is not my aim to surprise or shock you - but the simplest way I can summarize is to say that there are now in the world machines that think, that learn and that create, ... More precisely: within 10 years a computer would be **chess champion**, and an important new **mathematical theorem** would be proved by a computer.

Both of these milestones have now been achieved by computers, but each happened much later, more than **10** years:

1. Four color theorem proven in 1976 by Kenneth Appel and Wolfgang Haken

2. Deep Blue *chess* computer *defeated* Kasparov in 1997

History

☐ **1956**: The **Dartmouth Summer Research Conference on Artificial Intelligence**, Seminal event for AI as a field

☐ **1976**: **Four color theorem** was proven by Kenneth Appel and Wolfgang Haken

☐ **1997**: Deep Blue **chess** computer *defeated* Kasparov (10^{50})

  – "Saying Deep Blue doesn't really think about chess is like saying an airplane doesn't really fly because it doesn't flap its wings."   --Drew McDermott

GOFAI: Good-Old-Fashioned Artificial Intelligence

☐ Since 2012, we saw breakthroughs in **Speech Recognition, Image Recognition, Machine Translation, Autonomous Car,** etc.

☐ **2015**, **CNN+ReLU** outperforms humans in image recognition.

☐ **2016**, **AlphaGo** beat a 9-dan professional (Lee Sedol) (10^{170})

☐ **2020**, **GPT-3** is doing well simultaneously in many natural language tasks.

☐ **2023**, **Generative AI** such as StableDiffusion and ChatGPT provided first glimpse into what AGI looks like.

(Statistical) Machine Learning

12

Question Answering

- IBM Watson / Jeopardy!

- Apple Siri

  – The application uses a natural language user interface to answer questions, make recommendations, and perform actions by delegating requests to a set of Web services.

- Amazon Echo / Google Home

  – Associate the natural language user interface with a speaker.

Can Computers Understand?

- **They ate the pie with ice cream.**

- **They ate the pie with rhubarb.**

- **They ate the pie with paper plates.**

- **They ate the pie with cold milk.**

- **They ate the pie with friends.**

- **They ate the pie with dinner.**

- **They ate the pie with enthusiasm.**

- **They ate the pie with spoons.**

- **They ate the pie with napkins.**

from Dr. Douglas Lenatand Dr. Michael Witbrock   14

State-of-the-art AI for Question and Asnwers

LU How to interpret "with" in the following sentences? They ate the pie with ice cream.
They ate the pie with rhubarb.
They ate the pie with paper plates.
They ate the pie with cold milk.
They ate the pie with friends.
They ate the pie with dinner.
They ate the pie with enthusiasm.
They ate the pie with spoons.
They ate the pie with napkins.

In these sentences, "with" is used to indicate the accompaniment or manner in which the action of eating the pie occurred.

For example:

- "They ate the pie with ice cream" means that they consumed the pie while also having ice cream on the side.
- "They ate the pie with friends" means that they enjoyed the pie in the company of their friends.
- "They ate the pie with enthusiasm" means that they ate the pie with great interest and excitement.
- "They ate the pie with spoons" means that they used spoons to eat the pie.

In each of these cases, "with" is used to show the relationship between the pie and the accompanying item or manner in which it was eaten.

15

State-of-the-art AI for Question and Asnwers

LU how about the other sentences?

Similarly, in the other sentences:
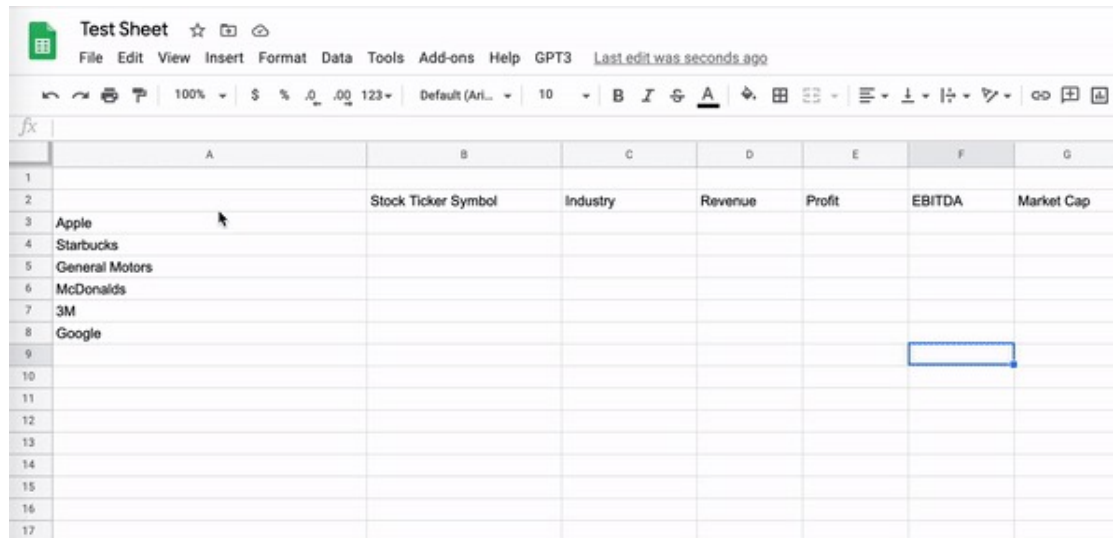
- "They ate the pie with rhubarb" means that they ate the pie with rhubarb as an ingredient or a topping.
- "They ate the pie with paper plates" means that they used paper plates to serve and eat the pie.
- "They ate the pie with cold milk" means that they consumed the pie with cold milk as a beverage.
- "They ate the pie with dinner" means that they had the pie as a part of their meal.
- "They ate the pie with napkins" means that they used napkins to wipe their mouths while eating the pie.

Again, in each of these cases, "with" is used to show the relationship between the pie and the accompanying item or manner in which it was consumed.

A lot more that GPTs are able to do:

- Generate well-written text

- Generate code for website layout, format charts / figures.

- Parse and complete spreadsheets

- Translations: English ⇔ Chinese, Legal ⇔ Plain English

- Check out this video: https://youtu.be/_x9AwxfjxvE?t=219

It can even be creative

LU Could you write a rap song in Eminem style about Prof. Wang's course on Artificial Intelligence at UCSB?

Yeah, let me give it a shot,
Talking 'bout Prof. Wang's course,
On AI, it's time to learn a lot.

Verse 1:
Yo, listen up, class is in session,
Artificial Intelligence, that's the lesson.
Prof. Wang's the teacher, the AI mastermind,
UCSB's the place, the classroom's the grind.

Chorus:
AI, AI, Prof. Wang's course on AI,
Gonna learn about machine learning, and the AI supply.
Gonna code, gonna train, gonna build a bot or two,
It's gonna be lit, that's just what we do.

Verse 2:
Starting with the basics, the fundamentals of AI,
Data structures, algorithms, and the models we rely.
We'll dive into neural nets, deep learning, and more,
Regression, classification, it's time to explore.

18

It can even be creative

Verse 3:

We'll tackle NLP, natural language processing,
Sentiment analysis, and topic modeling, no messing.
Computer vision, object recognition, and more,
We'll train our models, make predictions, and score.

Chorus:

AI, AI, Prof. Wang's course on AI,
Gonna learn about machine learning, and the AI supply.
Gonna code, gonna train, gonna build a bot or two,
It's gonna be lit, that's just what we do.

Verse 4:

At the end of the course, we'll have the skills we need,
To build our own AI systems, no matter the feed.
Big data, small data, we can handle it all,
Thanks to Prof. Wang, UCSB, and this AI haul.

Chorus:

AI, AI, Prof. Wang's course on AI,
Gonna learn about machine learning, and the AI supply.
Gonna code, gonna train, gonna build a bot or two,
It's gonna be lit, that's just what we do.

Outro:

So that's it, class, time to hit the books,
Study hard, code harder, and don't give dirty looks.
Thanks for listening to my rap, it's been a pleasure,
Good luck in the course, and let's all strive for treasure.

Other cool demos of AI

- Learning to walk / run with Reinforcement Learning
  - https://youtu.be/gn4nRCC9TwQ

- Learning to generate a realistic face
  - https://www.thispersondoesnotexist.com/

EBGAN (Zhao et al, 2016)

BEGAN (Berthelot et al, 2017)

- Learning to play "hide and seek" in creative ways:
  https://youtu.be/Lu56xVlZ40M?t=50

Art generation

Tell DALL E 2:

"Students eager to learn machine learning on a Californian beach in Monet style"

"We are in the middle of a new industrial revolution." --- Bernhard Schölkopf

- First two revolutions: our ability to harness energies
  - Coal / Fossil fuel
  - Electricity

- Third revolution: our ability to harness information
  - Digital revolution from 1950s.
  - Big Data revolution, most recently, AI revolution

- Fundamental societal changes
  - Job loss
  - Addiction, polarization
  - Privacy, fairness
  - Synthesize voice, face of any one
  - AI Weapons?

Further reading: https://arxiv.org/pdf/1911.10500.pdf

The exponential growth of AI as a field

Attendance at large conferences (1984-2019)
Source: Conference provided data.

Fig. 2.1a

NIPS @NipsConference · 4m
#NIPS2018 The main conference sold
out in 11 minutes 38 seconds

💬 3   🔁 21   ♡ 25

"For those debating ML hype, getting a ticket to a ML conference is now more challenging than a Taylor Swift conference or a Hamilton showing."

– Kaggle CTO Ben Hamner

Is it all just hype?

- AI research is supported by rigorous mathematical foundation.
  - Although… there's caveat.
  - Theory is always a bit behind.

"Machine learning has become alchemy"
~ Ali Rahimi

- AI Researchers know the limitations of their algorithms
  - (although.... sometimes it's easy to go over the top)

Andrew Ng
@AndrewYNg

Following

Should radiologists be worried about their jobs? Breaking news: We can now diagnose pneumonia from chest X-rays better than radiologists.
stanfordmlgroup.github.io/projects/chexn…

3:20 PM - 15 Nov 2017 from Mountain View, CA

1,440 Retweets 2,401 Likes

114   1.4K   2.4K

Regardless, it is an exciting time!

- The field of AI is gathering talents from CS, Stats, Math, Physics, Engineering, Business like never before!

  - 70% PhD applicants to UCSB CS want to specialize in AI / ML!

- We are training the first generation of Da Vinci-like engineers / researchers in large scale
  - Who can swiftly leverage multiple branches of mathematics
  - Understand computing / computer systems / write efficient codes
  - Understand statistics / causality / handles big data
  - Inspired by ideas from neural science / biological world.

Goals of this course

- To teach you some ideas of AI

- To introduce you to a set of key techniques and algorithms from AI

- To get you thinking about how AI can be applied to a variety of real problems

- To help you understand what's hard in AI

- To know where to find additional materials when needed in the future

- **It is not about vision, natural language processing, machine learning, AlphaGo, ChatGPT…it is an entry level course**

The schedule of the course

Week	Topic
1	Course Overview & Intelligent Agents
2	Machine Learning
	Machine Learning
2	Machine Learning
	Probabilistic Graphical Models
3	Probabilistic Graphical Models
	Search: Problem solving with search
4	Search: Search algorithms
	Search: Minimax search and game playing
5	Midterm Review
	Midterm
6	RL: Intro / Markov Decision Processes
	RL: Solving MDPs
7	RL: Bandits and Exploration
	RL: Reinforcement Learning Algorithms
8	RL: Reinforcement Learning Algorithms
	Logic: Propositional Logic
9	Logic: First order Logic
	Responsible AI
10	Final Review
	Final Exam

Machine Learning

Probabilistic Reasoning

Search

Reinforcement Learning

Logic

Responsible AI

27

There will be some steep learning curves

- You need to use math:
  - Calculus, Linear Algebra, Probabilities.
  - Know how to do complexity analysis in Big O notation.

- You need to be able to code:
  - Mainly in Python (numpy, scipy)
  - Data structures: trees, graphs, hashtables, sparse matrices, etc…

- Coding projects are required ($15\% \times 3$)

- **Theoretical homework are optional**, but recommended
  - they will help you gain deeper understanding of the topics.

- Formal prerequisite:
  - CS20,CS50, CS130A, PSTATS 120A.

It will be "worth it".

- Like Jedi training.

  - The information can be presented to you, but you can only learn if you are actively looking and thinking.
  - I learn a lot more from solving coding projects and homeworks on my own than just from the lectures.

- We (your instructor, TA and peers) are ready to help you!
  - Discussion sessions on Wednesdays, Office hours, Ed Stem

Expectation of you

- Academic Integrity
    - Coding projects are individual work
    - Discussion with your peers is fine but you need to disclose it.
    - Finding solutions to the exact questions on the internet is not OK.
    - Sharing equations, sharing code / pseudo-code are not OK.
    - If not sure if anything is a violation of the course policy, check with the instruction team on Piazza.

- Creating a respectful learning environment for all
    - Be respectful to your peers and instructors
    - Harassment and discrimination of any form will not be tolerated.

* More details on these in the syllabus

Course information

- Web sites
  - https://sites.cs.ucsb.edu/~yuxiangw/classes/AICourse-2023Spring/
    - Announcements, syllabus, schedule, lecture notes, assignments, related links

- Discussion on Ed Discussion:
  - https://edstem.org/us/courses/38544/discussion/
  - Coding project submission on Gradescope
  - You will be added directly to it.

- Textbook
  - Stuart Russell and Peter Norvig, Artificial Intelligence: A Modern Approach, **Fourth Edition**, 2010
  - Other reference books will be referred to. (e.g., Sutton and Barto will be our main text for the RL part of the class)

Workload and Grade

- 3 Coding Projects, 2 Exams

- Evaluation:
  - 45% Coding Projects, 25% Midterm exam, 30% Final exam
  - x% bonus points: For doing extras in coding projects; for class participation

- Letter grades:
  - A: 90%. A-: 85%, B+: 80%, B: 70%, C: 65%, D:60%
  - I do curve it sometimes. Also, if you are near the top of the class I am happy to write you **strong letters**.

- Late Policy:
  - Each student has **a total of 4 late days** without penalties.
  - When your late days are used up, you may still submit the project. When submitted within a week of the due date you will receive 75% of the credits. Beyond that only 50%.

Office Hours

- Instructor Office Hour: Tuesday 2:00-3:00pm

- TA Office Hour:

    Zihan: Monday 1 pm

    Vihaan: Wednesday 4 pm

    Esha: TBD

- Venue (Tentative): Space in front of Henley Hall 2118

- **Ask your questions on Ed Discussion!**

Useful resources

- Appendix A in the AIMA book for a mathematical review.

- Standard mathematical notations used in ML / AI:
  http://www.deeplearningbook.org/contents/notation.html

- Chapter 2, Chapter 17-18 of the D2L book: http://d2l.ai/

- Probability / Linear Algebra review: https://tinyurl.com/ydanmuk8

Remainder of the lecture: Intelligent agents

- Formally setting up the problem
    - Intelligent Agents
    - Task environment
    - The Model-Inference-Learning Paradigm
    - Model vs. reality

The concept of an "intelligent agent"

"An intelligent agent is an entity capable of combining cognition, perception and action in behaving autonomously, purposively and flexibly in some environment."

Our model of an agent

- An agent **perceives** its environment, **reasons** about its goals, and **acts** upon the environment
  - Abstractly, a function from percept histories to actions

$$f : P^* \rightarrow A$$

- Main components of an agent
  - Perception (sensors)
  - Reasoning/cognition
  - Action (actuators)

- Supported by
  - knowledge representation, search, inference, planning, uncertainty, learning, communication….

Discussion: Examples of Intelligent Agents in our daily life

- We talked about Chess AI, Chatbot, Mars rovers… anything else?

Our view of AI

- AIMA view – AI is building intelligent (rational) agents
  - Principles of rational agents, and
  - Models/components for constructing them

- Rational = "Does the right thing" in a particular situation
  - Maximize *expected* performance (not *actual* performance)

- So a rational agent does the "right" thing (at least tries to)
  - Maximizes the likelihood of success, given its information
  - How is "the right thing" chosen?
    - Possible actions (from which to choose)
    - Percept sequence (current and past)
    - Knowledge (static or modifiable)
    - Performance measure (*wrt* goals – defines success)

Our view of AI (cont.)

- So this course is about <u>designing rational agents</u>

  - Constructing f

  - For a given class of environments and tasks, we seek the agent (or class of agents) with the "best" performance

  - Note: Computational limitations make complete rationality unachievable in most cases

- In practice, we will focus on <u>problem-solving techniques</u> (ways of constructing f), not <u>agents</u> per se

Ideal Rational Agent

- In other words…

   "For each possible percept sequence, an ideal rational agent should do whatever action is expected to maximize its performance measure, on the basis of the evidence provided by the percept sequence and whatever built-in knowledge the agent has."

Note that:

| Rational ≠ Omniscient |
| --- |
| Rational ≠ Clairvoyant |
| Rational ≠ Successful |

Describing the Task Environment

Figure 2.1 Agents interact with environments through sensors and actuators.

- **PEAS** – Performance measure, Environment, Actuators, Sensors
  - Goals may be   into performance measure)
- Not limited to p   ts)
  - Any AI program

The Vacuum World: 3 min discussion

<u>P</u>erformance measure, <u>E</u>nvironment, <u>A</u>ctuators, <u>S</u>ensors

The Vacuum World

- Performance (P)
  - Keep world clean
  - Possible performance measures

- Environment (E)
  - Location
  - Cleanliness

- Three actions (A)
  - Move right
  - Move left
  - Remove dirt

- Sensed information (percepts) of environment (S)
  - Two locations
    - Left
    - Right
  - Two states
    - Dirty
    - Clean

PEAS Descriptions

Agent Type	Performance Measure	Environment	Actuators	Sensors
Medical diagnosis system	Healthy patient, reduced costs	Patient, hospital, staff	Display of questions, tests, diagnoses, treatments	Touchscreen/voice entry of symptoms and findings
Satellite image analysis system	Correct categorization of objects, terrain	Orbiting satellite, downlink, weather	Display of scene categorization	High-resolution digital camera
Part-picking robot	Percentage of parts in correct bins	Conveyor belt with parts; bins	Jointed arm and hand	Camera, tactile and joint angle sensors
Refinery controller	Purity, yield, safety	Refinery, raw materials, operators	Valves, pumps, heaters, stirrers, displays	Temperature, pressure, flow, chemical sensors
Interactive English tutor	Student's score on test	Set of students, testing agency	Display of exercises, feedback, speech	Keyboard entry, voice

Figure 2.5 Examples of agent types and their PEAS descriptions.

45

Environments

- Properties of environments
  - Fully vs. partially observable
  - Deterministic vs. stochastic
  - Episodic vs. sequential
  - Static vs. dynamic
  - Discrete vs. continuous
  - Single agent vs. multiagent

- The environment types largely determine the agent design

- The real world is partially observable, stochastic, sequential, hostile, dynamic, and continuous
  - Bummer…

New paradigm: Modeling-Inference-Learning

Modeling

Inference

Learning

(Idea and example taken from Percy Liang's teachings)

# Paradigm: Modelling

Real world

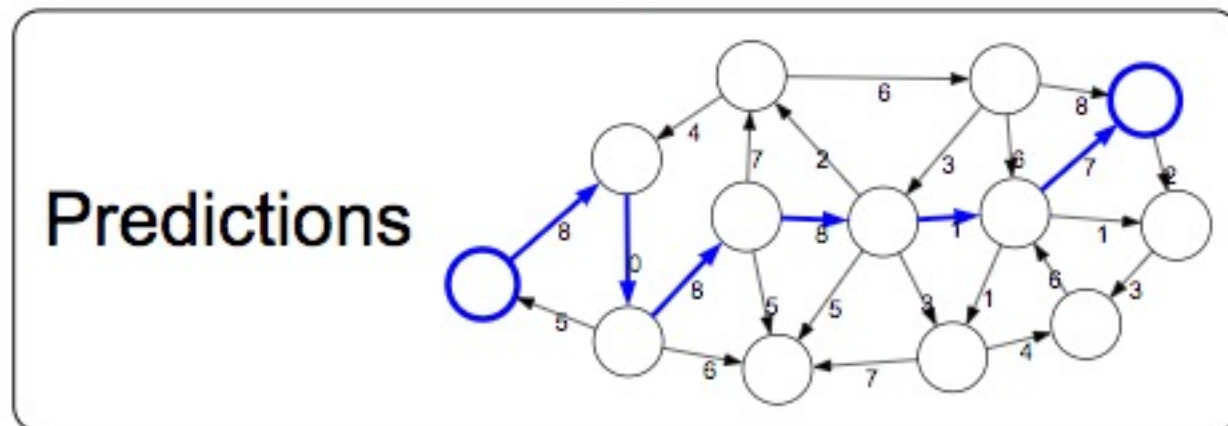**Modeling**

Model

Paradigm: Inference

Paradigm: Learning

Model without parameters

Model with parameters

+data

Learning

Structure of the course

Modelling Languages

Probabilistic Graphical Models / Deep Neural Networks

| Classification / Regression Bandits | Search game playing | Markov Decision Processes Reinforcement Learning | Logic, knowledge base Probabilistic inference |

Reflex Agents   **Planning Agents**   **Reasoning agents**

Low-level intelligence   High-level intelligence

Machine Learning

Algorithms that help AI to improve with data

(Again this idea is adapted from Percy Liang's teachings)   51

Generic Agent Program

- Implementing $f : P^* \rightarrow A$ ...or... $f(P^*) = A$
  - Lookup table?
  - Learning?

Knowledge, past percepts, past actions

function SKELETON-AGENT(*percept*) **returns** action
static: *memory*, the agent's memory of the world

memory ← UPDATE-MEMORY(*memory*, *percept*)
action ← CHOOSE-BEST-ACTION(*memory*)
memory ← UPDATE-MEMORY(*memory*, *action*)
return *action*

e.g.,

<u>Table-Driven-Agent</u>

Add *percept* to percepts

LUT [percepts, table]

NOP

AIMA's categorization of agent programs

- Simple reflex agent

- Model-based reflex agent

- Goal-based agent

- Utility-based agent

- Learning agent

(Read more in Section 2.4 of the AIMA book.)

Potential mid-term questions:
1. Where do these agent fall under our new categorization?
2. What are these agent's "Modelling-Inference-Learning" components?

What kind of agent it is in the Vacuum world?

- Reflex, planning, reasoning?

- What is the model? Are there any learning components?

When to use which type of agent?

- Depends on the problem (task environment)
  - Stochastic/deterministic/stateful/adversarial …

- Depends the amount of data available
  - Often we need to learn how the world behaves

- Depends on the dimensionality of your observations

Solving the right problem approximately

vs

Solving an approximation of the problem exactly

"All models are wrong, but some are useful."

George Box
(1919 - 2013)

Next three lectures: machine learning

- Chapter 19 in AIMA (4th Edition)

- Optional readings:
  - Chapter 1 in D2L: Intro to ML
    http://d2l.ai/chapter_introduction/index.html
  - Chapter 3,4,11 in D2L:
    - Chapter 3,4: From linear classification to neural networks
    - Chapter 11: Optimization in ML.

Final notes

- Submit (anonymous) feedbacks here:
  - https://forms.gle/rfPb5KwRwaaX38Zu6

  - What do you like / dislike about the course so far?
  - Are you following the lectures?
  - Am I going too fast or too slow?
  - Which AI topic from the course-schedule that you are most excited about?

- Homework 0 is on the website. It's optional but may help you. Please try it before coming to the discussion class on Wednesday.