

# Artificial Intelligence

CS 165A

Mar 29, 2022

**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 ['ju:'ʃi:lŋ],  
namely, y~eu~ee - sh~ih~ah~ng .

**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/>

# Teaching Assistant

- TA: Dan Qiao

[danqiao@ucsb.edu](mailto:danqiao@ucsb.edu)

- First Year PhD student in CS
- BS from Peking University
- Research focus:
  - Reinforcement learning






# 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



# AI / Machine Learning has revolutionized almost every aspect of our daily life



**Zillow** Edit Save Share More Close

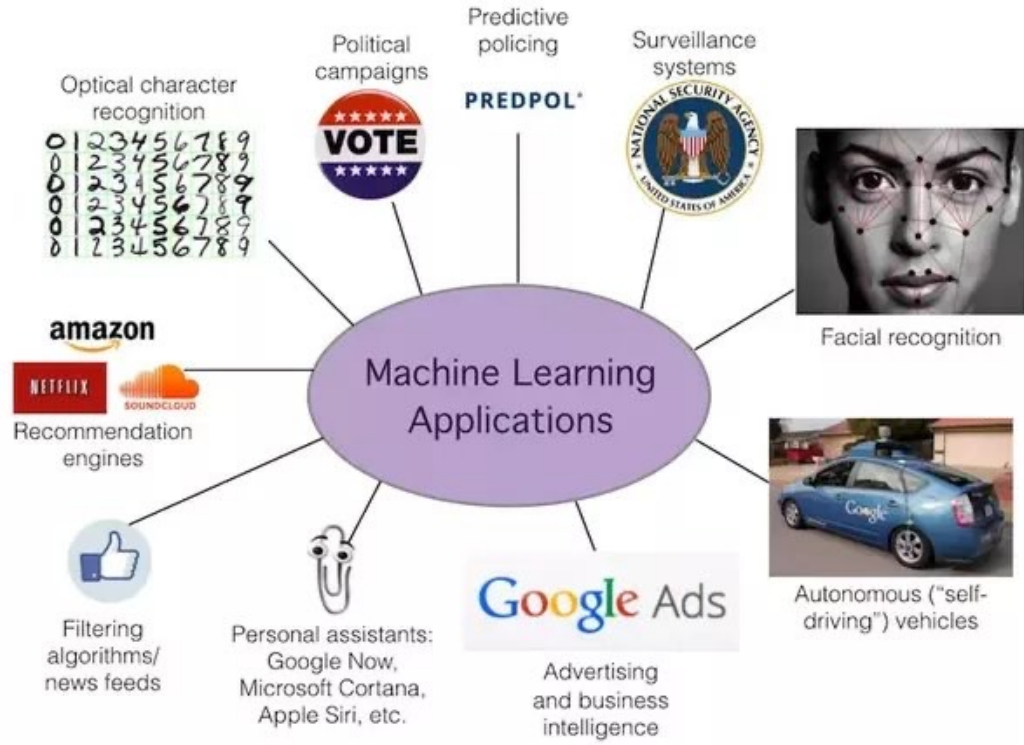
3 bd | 3 ba | 1,417 sqft  
 123 Main Street, Las Vegas, NV 89148  
 Off market | Zestimate®: \$266,436 | Rent Zestimate®: \$1,525/mo  
 Est. refi payment: \$1,277/mo See current rates

Home value Comparable homes Ways to sell Owner tools

**Sell to Zillow for your Zestimate**  
 Qualifying homes get a competitive cash offer.

**\$266,500**  
 Before taxes & fees

Get your offer



**Machine Learning Kaggle**

**Credit Card Risk Assessment**

**RISK**

# 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




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
- 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.

GOFAI:  
Good-Old-  
Fashioned  
Artificial  
Intelligence

(Statistical)  
Machine  
Learning

# 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.**

# State-of-the-art AI for Question and Answers

- GPT-3 from OpenAI
- See example here: [Giving GPT-3 a Turing Test](#)

Q: What is your favorite animal?

A: My favorite animal is a dog.

Q: Why?

A: Because dogs are loyal and friendly.

Q: What are two reasons that a dog might be in a bad mood?

A: Two reasons that a dog might be in a bad mood are if it is hungry or if it is hot.

Q: How many eyes does a giraffe have?

A: A giraffe has two eyes.

Q: How many legs does a frog have?

A: A frog has four legs.

Q: Are there any animals with three legs?


A: No, there are no animals with three legs.

Q: Why don't animals have three legs?

A: Animals don't have three legs because they would fall over.

# A lot more that GPT-3 is able to do:

- Generate well-written text
- Generate code for website layout, format charts / figures.
- Parse and complete spreadsheets



The screenshot shows a Google Sheets spreadsheet with the following data:


	A	B	C	D	E	F	G
1							
2		Stock Ticker Symbol	Industry	Revenue	Profit	EBITDA	Market Cap
3	Apple						
4	Starbucks						
5	General Motors						
6	McDonalds						
7	3M						
8	Google						
9							
10							
11							
12							
13							
14							
15							
16							
17							

- Translations: English ⇔ Chinese, Legal ⇔ Plain English
- Check out this video: [https://youtu.be/\\_x9AwxvjxvE?t=219](https://youtu.be/_x9AwxvjxvE?t=219)




# 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/Lu56xVIZ40M?t=50>

# “We are in the middle of a new industrial revolution.” --- Bernhard Schölkopf

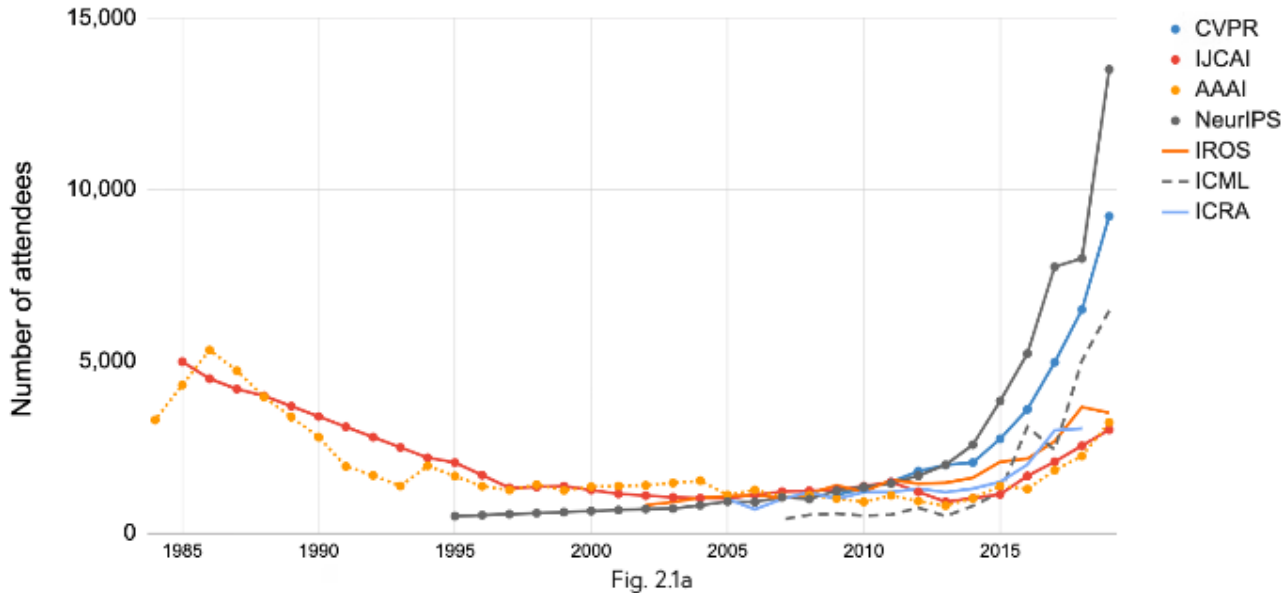
- First two revolutions: our ability to harness energies
  - Coal / Fossil fuel
  - Electricity
- Third revolution: our ability to harness information
  - Digital revolution from 1950s.
  - Big Data revolution, most recently, AI revolution
- Fundamental societal changes
  - Job loss
  - Addiction, polarization
  - Privacy, fairness
  - Synthesize voice, face of any one
  - AI Weapons?



# The exponential growth of AI as a field

Attendance at large conferences (1984-2019)

Source: Conference provided data.



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



# 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,...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

# 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 Friday, Office hours, Piazza

# 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-2022Spring/>
    - ◆ Announcements, syllabus, schedule, lecture notes, assignments, related links
- Discussion on Piazza:
  - <https://piazza.com/ucsb/spring2022/cs165a/home>
- 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: Thursday 2:00-3:00pm
- TA Office Hour: To be announced on Piazza.
- **Ask your questions on Piazza!**

# 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
  - The Model-Inference-Learning Paradigm
  - Task environment
  - Model vs. reality



# What's an Agent?

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

- Possible properties of agents:
  - Agents are **autonomous** – they act on behalf of the user
  - Agents can **adapt** to changes in the environment
  - Agents don't only act **reactively**, but sometimes also **proactively**
  - Agents have **social ability** – they communicate with the user, the system, and other agents as required
  - Agents also **cooperate** with other agents to carry out more complex tasks than they themselves can handle
  - Agents **migrate** from one system to another to access remote resources or even to meet other agents



# 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 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....

# Our view of AI (cont.)

- So this course is about designing rational agents
  - 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 problem-solving techniques (ways of constructing  $f$ ), not agents 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  $\neq$  Omniscient  
Rational  $\neq$  Clairvoyant  
Rational  $\neq$  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 explicit or implicit (built into performance measure)
- Not limited to physical agents (robots)
  - Any AI program

# The Vacuum World: 3 min discussion



Performance measure, Environment, Actuators, Sensors

# 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

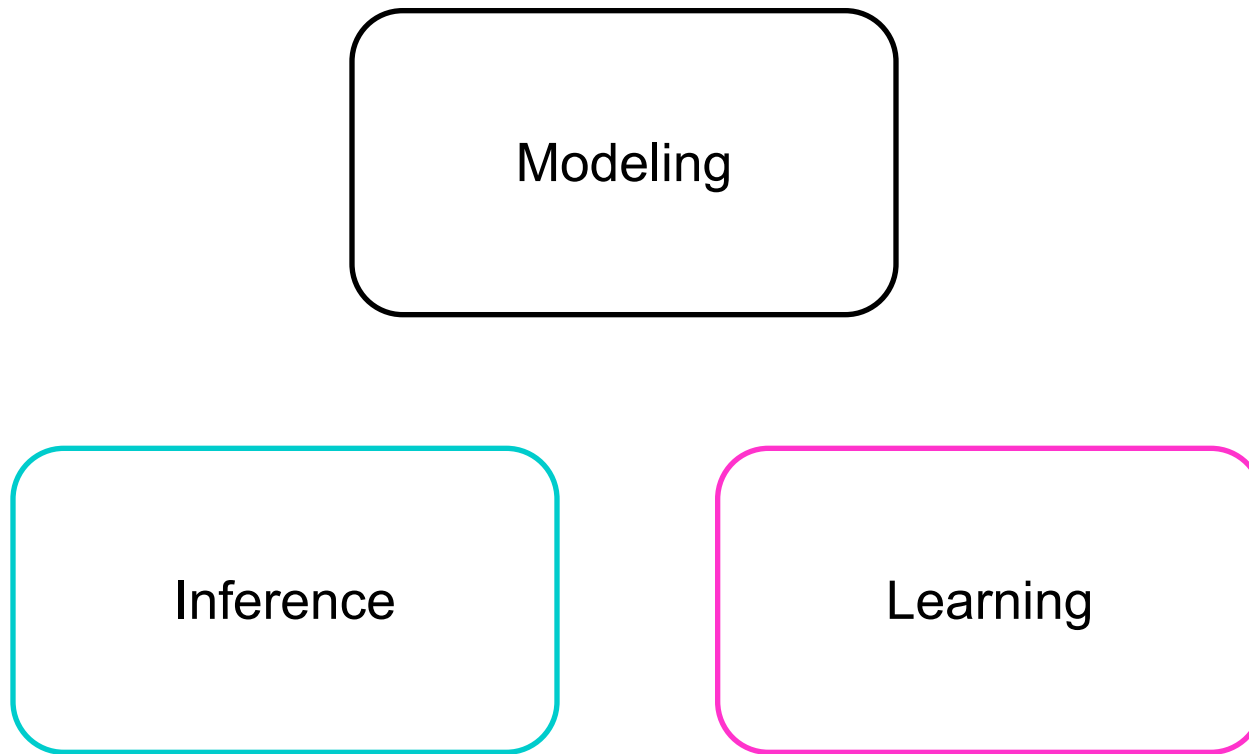
**Figure 2.5** Examples of agent types and their PEAS descriptions.



# 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



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


# Paradigm: Modelling

Real world




**Modeling**

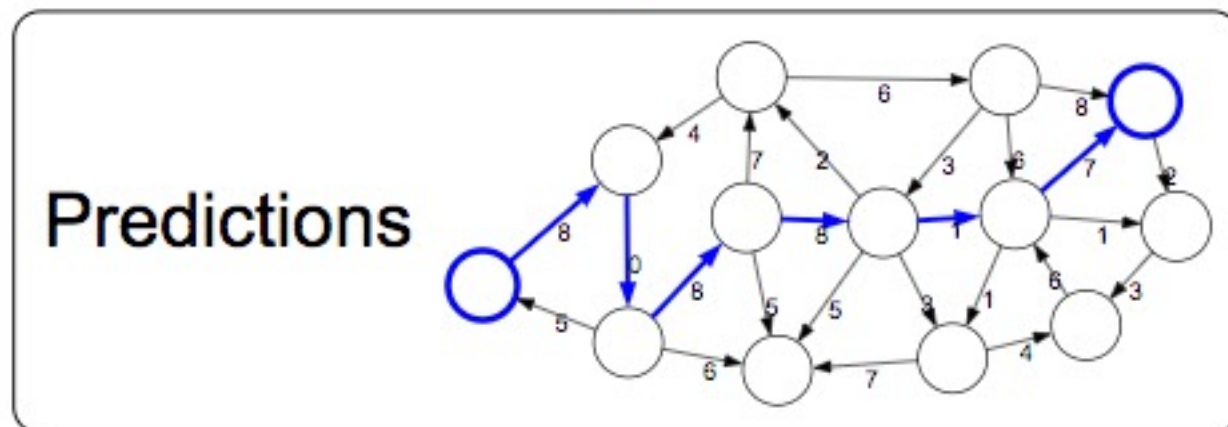
Model



# Paradigm: Inference




**Inference**



# Paradigm: Learning


Model without parameters



+data

**Learning**

Model with parameters



# Structure of the course

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*

# 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.,

Table-Driven-Agent

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

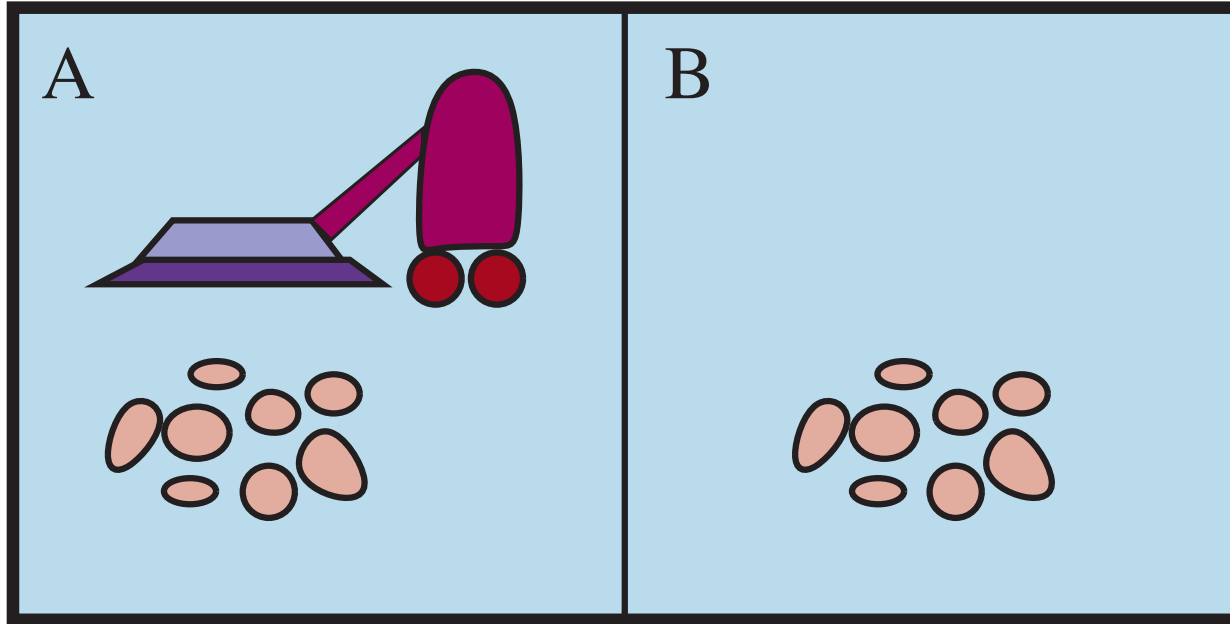
(Read more in Section 2.4 of the AIMA book.)

Potential mid-term questions:

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



# Quiz: 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 (4<sup>th</sup> Edition)
- Optional readings:
  - Chapter 1 in D2L: Intro to ML  
[http://d2l.ai/chapter\\_introduction/index.html](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/Pyoe7d7ePqmDKfnaA>
  - 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. Please try it before coming to the discussion class on Friday.