



上海交通大学  
SHANGHAI JIAO TONG UNIVERSITY



上海交通大学  
约翰·霍普克罗夫特  
计算机科学中心

John Hopcroft Center for Computer Science

# CS410: Artificial Intelligence

Shuai Li

John Hopcroft Center, Shanghai Jiao Tong University

<https://shuaili8.github.io>

<https://shuaili8.github.io/Teaching/CS410/index.html>

Part of slide credits: CMU AI & <http://ai.berkeley.edu>





# Teaching Assistant

- Canzhe Zhao (赵灿哲)
  - Email: [canzhezhaos@sjtu.edu.cn](mailto:canzhezhaos@sjtu.edu.cn)
  - 1<sup>st</sup> year PhD student
  - Research interests on bandit algorithms and optimization
  - Office hour: Fri 7-9 PM
- Zhihui Xie (谢知晖)
  - Email: [fffffarmer@sjtu.edu.cn](mailto:fffffarmer@sjtu.edu.cn)
  - 1<sup>st</sup> year Master student
  - Research on causal machine learning and recommendation systems
  - Office hour: Wed 7-9 PM

# References (will add more during course)

- Artificial Intelligence: A Modern Approach by Stuart Russell and Peter Norvig (4<sup>th</sup> edition)
- Reinforcement Learning: An Introduction by Richard S. Sutton and Andrew G. Barto
- 周志华 《机器学习》 清华大学出版社，2016.



# Goal

- Know what is AI and what it usually covers
- Familiar and understand popular AI problems and algorithms
- Be able to build AI models in applications
  - Know which algorithms to adopt and when to adopt
- Get a touch of latest research

# Prerequisites

- Basic computer science principles
  - Big-O notation
  - Comfortably write non-trivial code in Python/numpy
- Probability
  - Random Variables
  - Expectations
  - Distributions
- Linear Algebra & Multivariate/Matrix Calculus
  - Gradients and Hessians
  - Eigenvalue/vector

# Grading

- Attendance and participation: 5%
- Homework (written & programming) 40%
- Project: 25%
- Final exam: 30%

# Honor code

- Discussions are encouraged
- **Independently** write-up homework and code
- Same reports and homework will be **reported**

# Course Outline

- Search
- Constraint satisfaction problems
- Game trees
- Markov decision processes (MDPs)
- Reinforcement learning
- Hidden Markov models (HMMs)
- Bayes nets
- Machine learning basics
- Neural networks




What is AI?


# What is AI?

The science of making machines that:


Think like people




Think rationally



Act like people


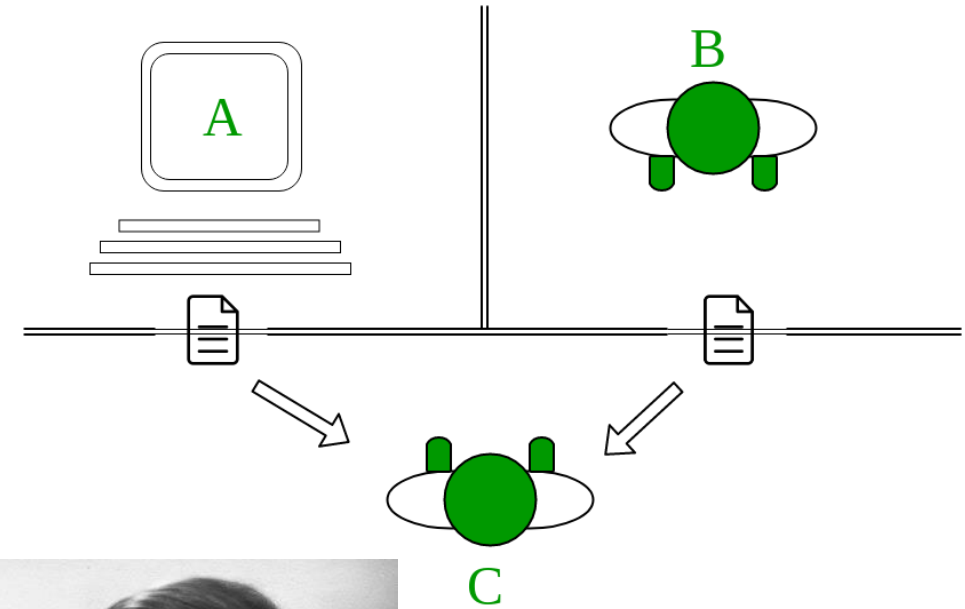


Act rationally



# Acting humanly: The Turing test approach

- In 1950, Turing defined a test of whether a machine could perform
- Practically though, it is a test of whether a machine can ‘act’ like a person
- “A human judge engages in a natural language conversation with one human and one machine, each of which tries to appear human. If judge can’t tell, machine passes the Turing test”



# Acting humanly: The Turing test approach 2

- The computer would need to possess the following capabilities
  - **Natural language processing** to enable it to communicate successfully in English/or other languages
  - **Knowledge representation** to store what it knows or hears
  - **Automated reasoning** to use the stored information to answer questions and to draw
  - **Machine learning** to adapt to new circumstances and to detect and extrapolate patterns
- Total Turing test includes a video signal, so the computer will need
  - **Computer vision** to perceive objects
  - **Robotics** to manipulate objects and move about

# Thinking humanly: The cognitive modeling approach

- The interdisciplinary field of **cognitive science** brings together computer models from AI and experimental techniques from psychology to construct precise and testable theories of the human mind
- Real cognitive science is necessarily based on **experimental investigation** of actual humans or animals
- In the early days of AI, people think that an algorithm performs well on a task  $\Leftrightarrow$  it is a good model of human performance

# 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!
- “Brains are to intelligence as wings are to flight”
- Lessons learned from the brain: memory and simulation are key to decision making



# Thinking rationally: The “laws of thought” approach

- The Greek philosopher Aristotle, syllogisms (三段论)
- The logicians hope to build on logic systems to create intelligent systems
- The emphasis was on **correct inferences**


# Acting rationally: The rational agent approach

- Making **correct inferences** is sometimes *part* of being a rational agent, but not *all*
- An **agent** is just something that acts (agent comes from the Latin *agere*, to do)
- A **rational agent** is one that acts so as to achieve the best outcome or, when there is uncertainty, the best expected outcome
- This approach has two advantages:
  - It is more general than the “laws of thought” approach because correct inference is just one of several possible mechanisms for achieving rationality
  - It is more amenable to scientific development than are approaches based on human behavior or human thought



# AI Definition by John McCarthy

- What is artificial intelligence
  - It is the science and engineering of making intelligent machines, especially intelligent computer programs
- What is intelligence
  - Intelligence is the computational part of the ability to achieve goals in the world
- John McCarthy (1927-2011)
  - co-authored the document that coined the term "artificial intelligence" (AI), developed the Lisp programming language family



# AI and this course

- Describe machines (or computers) that mimic "cognitive" functions that humans associate with the human mind, such as "learning" and "problem solving".


--Russell, S. J., & Norvig, P. (2016). *Artificial intelligence: a modern approach*. Malaysia; Pearson Education Limited.

- 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
  - **Computational Rationality**


# What is AI?

The science of making machines that:


A: Think like people




C: Think rationally




B: Act like people



D: Act rationally



# Maximize Your Expected Utility




# What is Machine Learning?

- Term “Machine Learning” coined by Arthur Samuel in 1959
  - Samuel Checkers-playing Program
- Common definition (by Tom Mitchell):
  - ***Machine Learning is the study of computer algorithms that improve automatically through experience***
- Subfield of Artificial Intelligence (AI)
  - The hottest subfield - reinvigorated interest in AI due to deep learning!

# Difference between AI and ML

- AI is a bigger concept to create intelligent machines that can simulate human thinking capability and behavior, whereas, machine learning is an application or subset of AI that allows machines to learn from **data** without being programmed explicitly.




# An example of AI but is not machine learning

- A\* search algorithm
  - Objective: Find the shortest path between two nodes of a weighted graph
  - Use heuristic information
  
- Compare with Breadth First Searching and Greedy Searching

# Breadth First Searching

- Pink: start point, Purple: end point;
- Blue: visited points, the darker the earlier



Each time it visits, or expand the point with least  $g(n)$  value

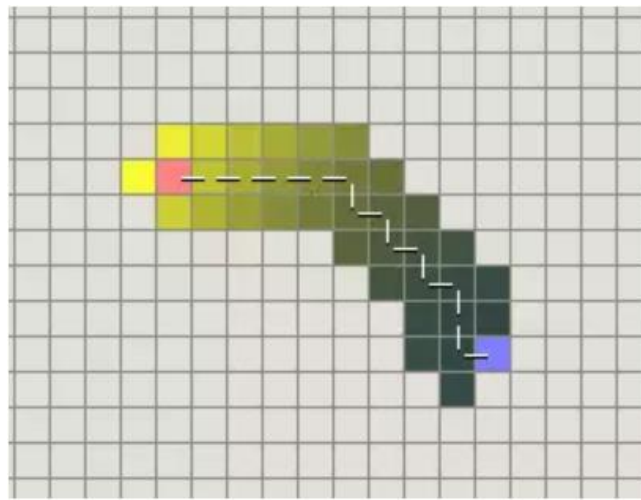
- $g(n)$  is the distance from start point to point  $n$

Short comings: computing burden is too high, it visited too many points before getting the end point.




# Greedy Searching

- Each time it visit or expand the point with least  $h(n)$  value
  - $h(n)$  is the distance from point  $n$  to end point. It works fine when there is no obstacles.

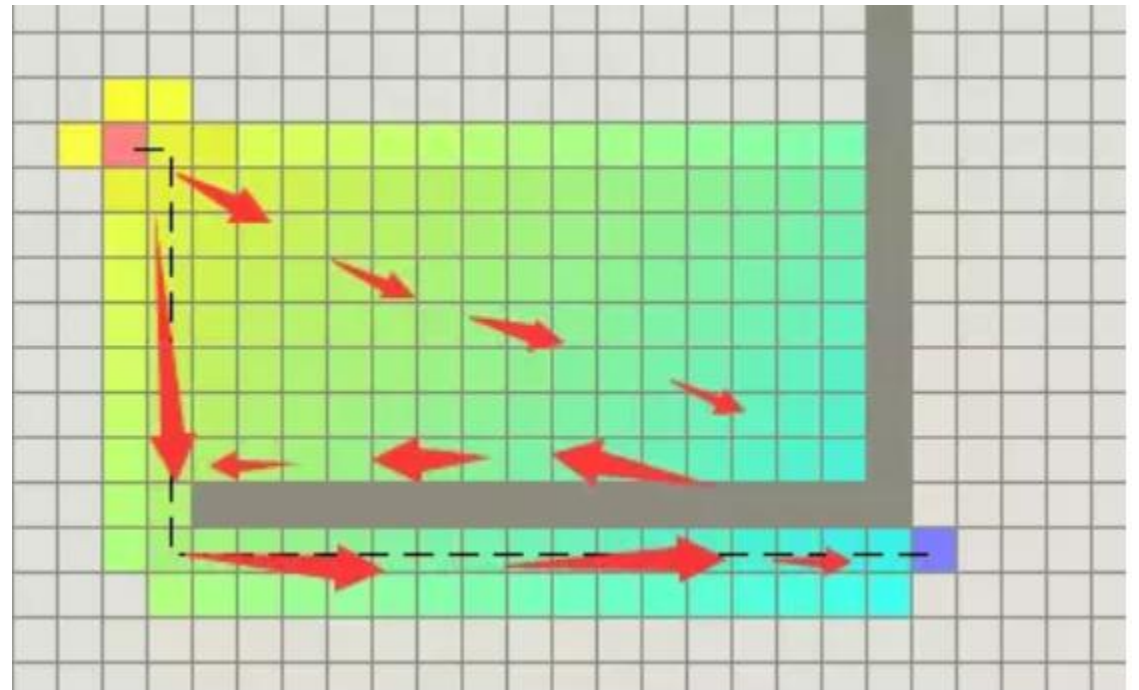



- The cost doubles when there is obstacles



# A\* algorithm

- It combines the stability of BFS and the heuristics in greedy searching.
- Each time it visits point with the least  $f(n) = g(n) + h(n)$  value.





# The Foundations of AI

The disciplines that contributed ideas, viewpoints, and techniques to AI

# Philosophy

- Can formal rules be used to draw valid conclusions?
- How does the mind arise from a physical brain?
- Where does knowledge come from?
- How does knowledge lead to action?
  
- Rationalism (理性主义)/materialism (唯物主义)/empiricism (经验主义)

# Mathematics

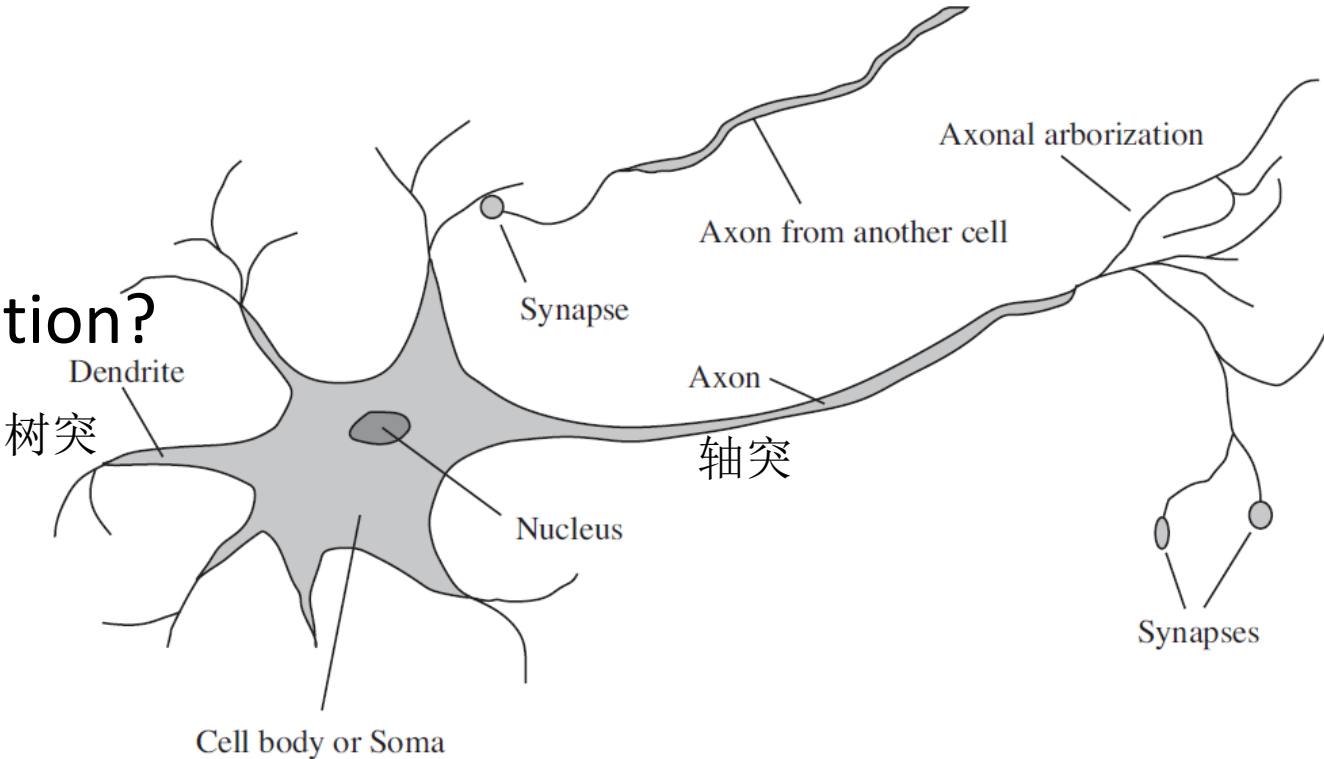
- What are the formal rules to draw valid conclusions?
- What can be computed?
- How do we reason with uncertain information?
  
- The first nontrivial **algorithm** is thought to be Euclid's algorithm for computing greatest common divisors
- The word **algorithm** (and the idea of studying them) comes from al-Khowarazmi, a Persian mathematician of the 9th century
- NP-completeness/probability/entropy

# Economics

- How should we make decisions so as to maximize payoff?
- How should we do this when others may not go along?
- How should we do this when the payoff may be far in the future?
  
- The pioneering AI researcher Herbert Simon (1916–2001) won the Nobel Prize in economics in 1978 for his early work showing that models based on **satisficing**—making decisions that are “good enough,” rather than laboriously calculating an optimal decision—gave a better description of actual human behavior (Simon, 1947).

# Neuroscience

- How do brains process information?



	Supercomputer	Personal Computer	Human Brain
Computational units	$10^4$ CPUs, $10^{12}$ transistors	4 CPUs, $10^9$ transistors	$10^{11}$ neurons
Storage units	$10^{14}$ bits RAM $10^{15}$ bits disk	$10^{11}$ bits RAM $10^{13}$ bits disk	$10^{11}$ neurons $10^{14}$ synapses
Cycle time	$10^{-9}$ sec	$10^{-9}$ sec	$10^{-3}$ sec
Operations/sec	$10^{15}$	$10^{10}$	$10^{17}$
Memory updates/sec	$10^{14}$	$10^{10}$	$10^{14}$

# Psychology

- How do humans and animals think and act?
- **Cognitive psychology** views the brain as an information-processing device
- **Developmental psychology** is the scientific study of how and why human beings change over the course of their life, especially concerned with infants and children



# Computer engineering

- How can we build an efficient computer?
- Designing algorithms is not enough
- Hardware
  - modern digital electronic computer
- Software
  - operating systems, programming languages, and tools needed to write modern programs (and papers about them)
- Work in AI has also pioneered many ideas that have made their way back to mainstream computer science
  - time sharing, interactive interpreters, personal computers with windows and mice

# Control theory and cybernetics

- How can artifacts operate under their own control?
- Goal of modern control theory, especially the branch known as stochastic optimal control, is to
  - Design systems that maximize an **objective function** over time
  - Roughly match our view of AI: designing systems that behave optimally
- Differences of control theory and AI:
  - Control theory more care about continuous variables with calculus and matrix algebra as tools
  - AI uses logical inference and computation to escape these limitations and consider problems such as language, vision, and planning that fell completely outside the control theorist's purview


# Linguistics

- How does language relate to thought?
- Understanding language requires an understanding of the subject matter and context, not just an understanding of the structure of sentences
- Knowledge representation (the study of how to put knowledge into a form that a computer can reason with)
  - decades of work on the philosophical analysis of language

# The History of AI

# 1950s

- Turing's test
- Dartmouth Conference 1956: the birth of AI




## 1956 Dartmouth Conference: The Founding Fathers of AI



John MacCarthy



Marvin Minsky



Claude Shannon



Ray Solomonoff




Alan Newell



Herbert Simon



Arthur Samuel



Oliver Selfridge



Nathaniel Rochester



Trenchard More

# Chess as the First Killer App for AI

- Claude Shannon proposed the first chess playing program in 1950
  - It included adversarial search and minimax (later lecture)
  - It also included many heuristics for faster searching




TABLE 1 Comparison of Current Chess Programs

	Turing	Kister, Stein, Ulam, Walden, Wells (Los Alamos)	Bernstein, Roberts, Arbuckle, Belsky (Bernstein)	Newell, Shaw, Simon (NSS)
<b>Vital statistics</b>				
Date	1951	1956	1957	1958
Board	8 × 8	6 × 6	8 × 8	8 × 8
Computer	Hand simulation	MANIAC-I 11,000 ops./sec	IBM 704 42,000 ops./sec	RAND JOHNNIAC 20,000 ops./sec
<b>Chess program</b>				
Alternatives	All moves	All moves	7 plausible moves Sequence of move generators	Variable Sequence of move generators
Depth of analysis	Until dead (exchanges only)	All moves 2 moves deep	7 plausible moves 2 moves deep	Until dead Each goal generates moves
Static evaluation	Numerical Many factors	Numerical Material, mobility	Numerical Material, mobility Area control King defense	Nonnumerical Vector of values Acceptance by goals
Integration of values	Minimax	Minimax (modified)	Minimax	Minimax
Final choice	Material dominates Otherwise, best value	Best value	Best value	1. First acceptable 2. Double function
<b>Programming</b>				
Language		Machine code	Machine code	IPL-IV, interpretive
Data scheme		Single board No records	Single board Centralized tables Recompute	Single board Decentralized List structure Recompute
Time	Minutes	12 min/move	8 min/move	1-10 hr/move (est.)
Space		600 words	7000 words	Now 6000 words, est. 16,000
<b>Results</b>				
Experience	1 game	3 games (no longer exists)	2 games	0 games
Description	Loses to weak player Aimless Subtleties of evaluation lost	Beats weak player Equivalent to human with 20 games experience	Passable amateur Blind spots Positional	Some hand simulation Good in spots (opening) No aggressive goals yet

45

**Allen Newell**  
**J. C. Shaw**  
**H. A. Simon**

# Chess-Playing Programs and the Problem of Complexity

# The Promise of AI

- In 1965, Herbert Simon predicted that “machines will be capable, within 20 years, of doing any work a man can do”
- In 1967, AI pioneer Marvin Minsky predicted “in from three to eight years we will have a machine with the general intelligence of an average human being.”
- In 1967, John McCarthy told the U.S. government that it would be possible to build “a fully intelligent machine” in the space of a decade




# 1970s - first AI winter



- Limited computer power
- Intractability and the combinatorial explosion
- Commonsense knowledge and reasoning
  - Hard to encode so many concepts and rules
  - Didn't know how to teach computers to learn these



# Computer Vision, Blocks World, Natural Language



Larry Roberts 1963 Thesis



Terry Winograd's 1971 Thesis on SHRDLU  
for natural language understanding

# Early Robots




Fig. 8. The CMU Rover.




Fig. 1. The Stanford Cart.

1983 – mobile robots by Hans Moravec



Dean Pomerleau (CMU) 1986  
NAVLAB controlled by NNs

<https://www.nytimes.com/video/science/1247468063802/stanford-cart.html>  
<https://www.youtube.com/watch?v=ntIczNQKfjQ>

# Deep Blue

- Started in the mid-1980s at CMU, didn't win until 1997
- Project moved to IBM
- “Good Old-Fashioned” Brute Force Search using custom hardware



- Win Garry Kasparov by 3.5:2.5 on Chess
- Search over 12 following steps

- <https://www.youtube.com/watch?v=KF6sLCeBj0s>

# Rise of Statistical Approaches: 1990s – 2000s

- Knowledge-based:
  - Search for a solution using general search algorithms
  - Encode knowledge that humans have and use logic to solve
- Statistical:
  - Learning patterns and choosing solutions based on observed likelihood







# Evolution of AI Research: 2010s



# 2010s-now

- Deep learning
  - The return of neural networks
- Big data
  - Large datasets, like ImageNet
- Computational power
- Artificial general intelligence (AGI)

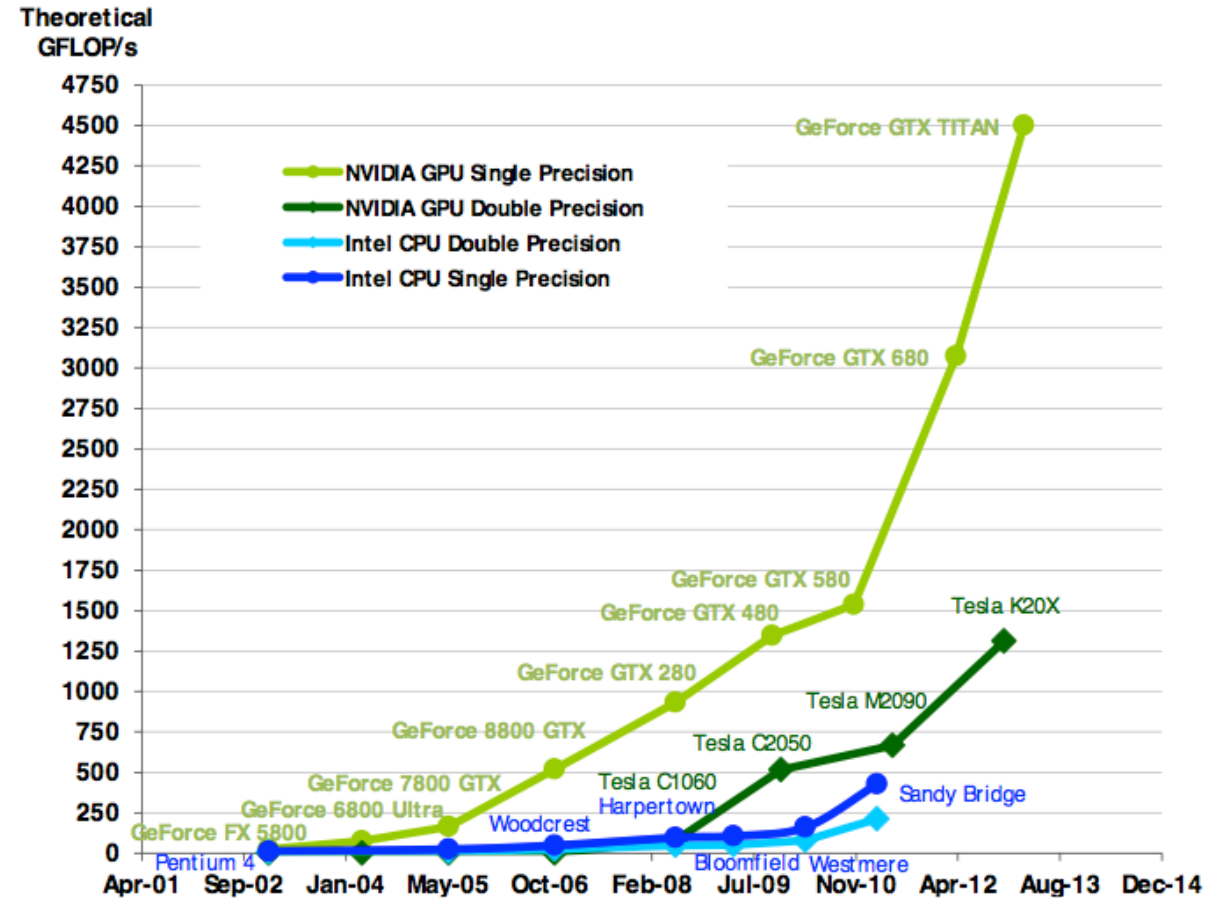



Figure 1 Floating-Point Operations per Second for the CPU and GPU

# Computer Vision (CV) -- ImageNet, AlexNet




**IMAGENET** [www.image-net.org](http://www.image-net.org)

**22K** categories and **15M** images

- Animals
  - Bird
  - Fish
  - Mammal
  - Invertebrate
- Plants
  - Tree
  - Flower
  - Food
  - Materials
- Structures
  - Artifact
  - Tools
  - Appliances
  - Structures
- Person
  - Scenes
  - Indoor
  - Geological Formations
  - Sport Activities


Deng, Dong, Socher, Li, Li, & Fei-Fei, 2009



Deng, J., Dong, W., Socher, R., Li, L. J., Li, K., & Fei-Fei, L. (2009, June). Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition* (pp. 248-255). IEEE.

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems* (pp. 1097-1105).


# CV -- GAN




a)



b)




c)



d)

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... & Bengio, Y. (2014). Generative adversarial nets. In *Advances in neural information processing systems* (pp. 2672-2680).

# CV (Detection) -- R-CNN, Fast R-CNN, Faster R-CNN






1. Input image

2. Extract region proposals (~2k)

3. Compute CNN features

4. Classify regions



1. Girshick, R., Donahue, J., Darrell, T., & Malik, J. (2014). Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 580-587).

2. Girshick, R. (2015). Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision* (pp. 1440-1448).

3. Ren, S., He, K., Girshick, R., & Sun, J. (2015). Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems* (pp. 91-99).

# Speech recognition (Unsupervised, ICA)

Mixed



Separated



# Speech recognition (Unsupervised, ICA, cont.)

Mixed



Separated



# Speech recognition

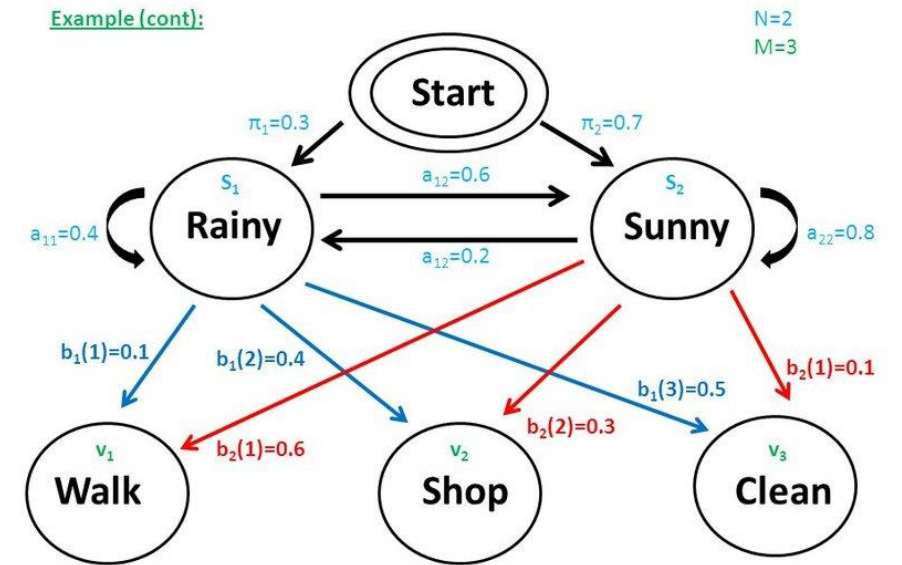
- Previous works use

- Hidden Markov models (HMMs)
  - Deal with the temporal variability of speech
- Gaussian mixture models (GMMs)
  - Determine how well each state of each HMM fits a frame or a short window of frames of coefficients that represents the acoustic input

- New

- Feed-forward neural network
  - Takes several frames of coefficients as input and produces posterior probabilities over HMM states as output



## Hidden Markov Model





# Speech recognition


## ■ Deep Learning: From **GMM-HMM** to **DNN-HMM**



# Natural Language Processing (NLP) -- Word2Vec

Image and audio processing systems work with rich, high-dimensional datasets encoded as vectors.

AUDIO



Audio Spectrogram

DENSE

IMAGES




Image pixels

DENSE

TEXT



Word, context, or document vectors

SPARSE

Pennington, J., Socher, R., & Manning, C. (2014, October). Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)* (pp. 1532-1543).

# Natural Language Processing (NLP) -- Word2Vec (cont.)

## Word Analogies

Test for linear relationships, examined by Mikolov et al. (2014)



a:b :: c:?




$$d = \arg \max_x \frac{(w_b - w_a + w_c)^T w_x}{\|w_b - w_a + w_c\|}$$

man:woman :: king:?

+	king	[ 0.30 0.70 ]
-	man	[ 0.20 0.20 ]
+	woman	[ 0.60 0.30 ]
<hr/>		
	queen	[ 0.70 0.80 ]



# NLP -- CNN



Kim, Y. (2014, October). Convolutional Neural Networks for Sentence Classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (pp. 1746-1751).

# NLP -- BERT

- BERT
  - Bidirectional Encoder Representations from [Transformers](#)
  - The pre-train deep bidirectional representations from unlabeled text by jointly conditioning on both left and right context in all layers
  - The pre-trained BERT model can be finetuned with just one additional output layer to create state-of-the-art models for a wide range of tasks, such as question answering and language inference, without substantial task-specific architecture modifications
  - It obtains new state-of-the-art results on eleven natural language processing tasks

Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

# NLP -- BERT

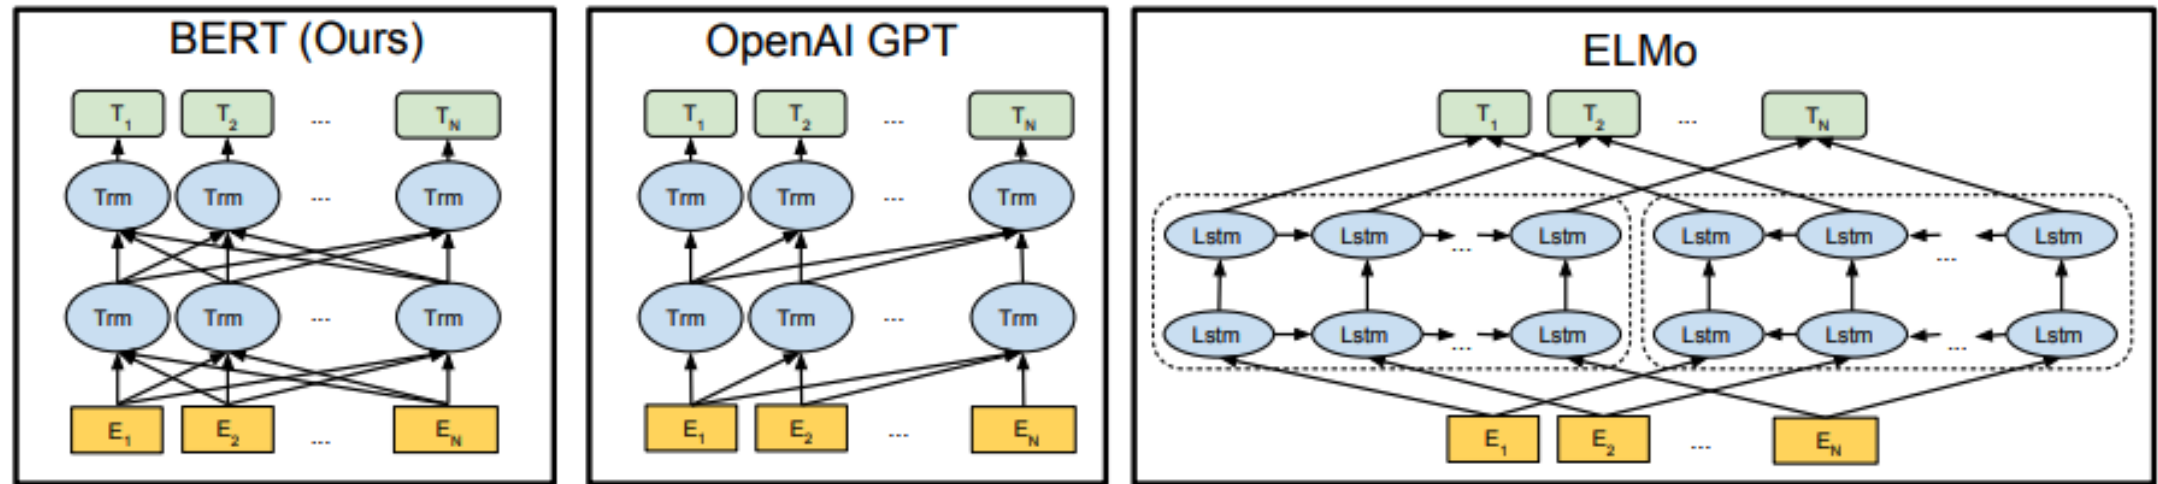



Figure 1: Differences in pre-training model architectures. BERT uses a bidirectional Transformer. OpenAI GPT uses a left-to-right Transformer. ELMo uses the concatenation of independently trained left-to-right and right-to-left LSTM to generate features for downstream tasks. Among three, only BERT representations are jointly conditioned on both left and right context in all layers.


# AlphaGo 2016



- Win Lee Sedol by 4:1 on Go
- Efficient search on large solution space

Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., ... & Dieleman, S. (2016). Mastering the game of Go with deep neural networks and tree search. *nature*, 529(7587), 484.

# Texas hold'em 2017



## DeepStack

- In a study involving 44,000 hands of poker, DeepStack defeated with statistical significance professional poker players in heads-up no-limit Texas hold'em
- Imperfect information setting

Moravčík, M., Schmid, M., Burch, N., Lisý, V., Morrill, D., Bard, N., ... & Bowling, M. (2017). Deepstack: Expert-level artificial intelligence in heads-up no-limit poker. *Science*, 356(6337), 508-513.



# History of Game AI

1956 checkers

1992 backgammon


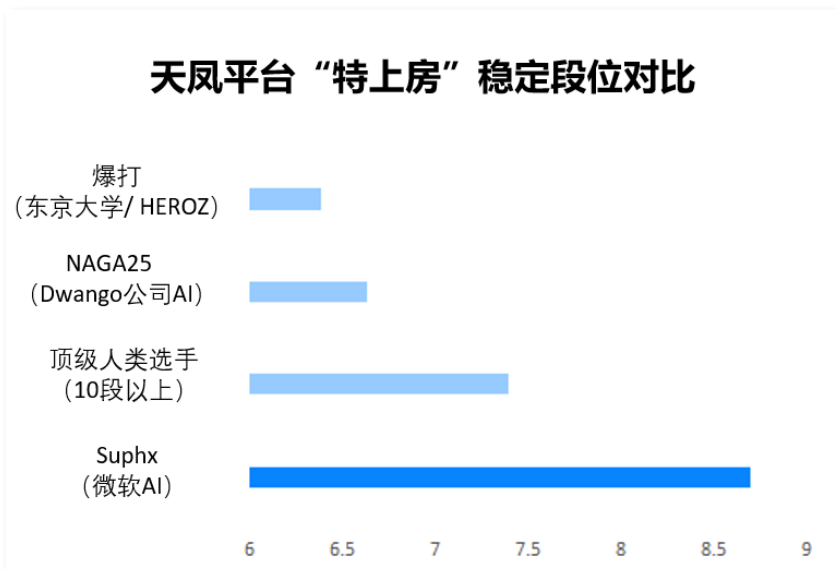
1994 checkers

1997 chess

2016 Go

2017 Texas hold'em

2019 Majiang



1956年

Arthur Samuel开发的AI程序首先应用于国际跳棋

1992年

Hans Berliner开发的AI程序在双陆棋中取得突破

1994年

Jonathan Schaeffer团队开发的AI程序战胜国际跳棋世界冠军

Chinook



许峰雄团队开发的AI程序击败国际象棋世界冠军

1997年

DeepMind团队开发的AI程序战胜围棋世界冠军

2016年

Deep Blue



2017年

卡耐基梅隆大学、Facebook AI以及阿尔伯特大学开发的AI程序在德州扑克取得突破

下一个里程碑在哪里?



AlphaGo




Libratus

Pluribus


DeepStack

# Game playing – state of the art



# Recent popularity of AI and ML

Google Trends Keywords 2009 - 2017 for Germany



# AI and Machine Learning Together: 2010s and 2020s

We've spent years feeding neural nets vast amounts of data, teaching them to think like human brains. They're crazy-smart, but they have absolutely no common sense. What if we've been doing it all wrong? BETH HOLZER

CLIVE THOMPSON

BUSINESS 11.13.2018 06:00 AM

## How to Teach Artificial Intelligence Some Common Sense

We've spent years feeding neural nets vast amounts of data, teaching them to think like human brains. They're crazy-smart, but they have absolutely no common sense. What if we've been doing it all wrong?

2,456 views | Oct 16, 2018, 08:30am

## AI Requires More Than Machine Learning



Jans Aasman Forbes Councils Member  
Forbes Technology Council  
COUNCIL POST | Paid Program  
Innovation


## Researchers: Are we on the cusp of an 'AI winter'?

By Sam Shoad  
Technology reporter

12 January 2020 | Technology

# What Can AI Do?

# Sci-Fi AI



# Face recognition, real-time detection




<https://bitrefine.group/home/transportation/face-recognition-support-system>

[https://cdn-images-1.medium.com/max/1600/1\\*q1uVc-MU-tC-WwFp2yXJow.gif](https://cdn-images-1.medium.com/max/1600/1*q1uVc-MU-tC-WwFp2yXJow.gif)

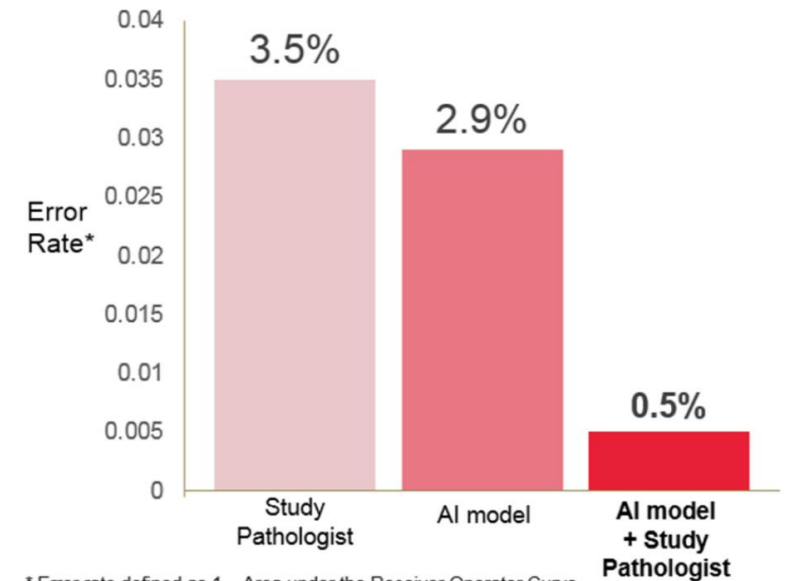
# Medical image analysis

- Segmentation results
- Breast Cancer Diagnoses

a



(AI + Pathologist) > Pathologist



\* Error rate defined as 1 - Area under the Receiver Operator Curve

\*\* A study pathologist, blinded to the ground truth diagnoses, independently scored all evaluation slides.

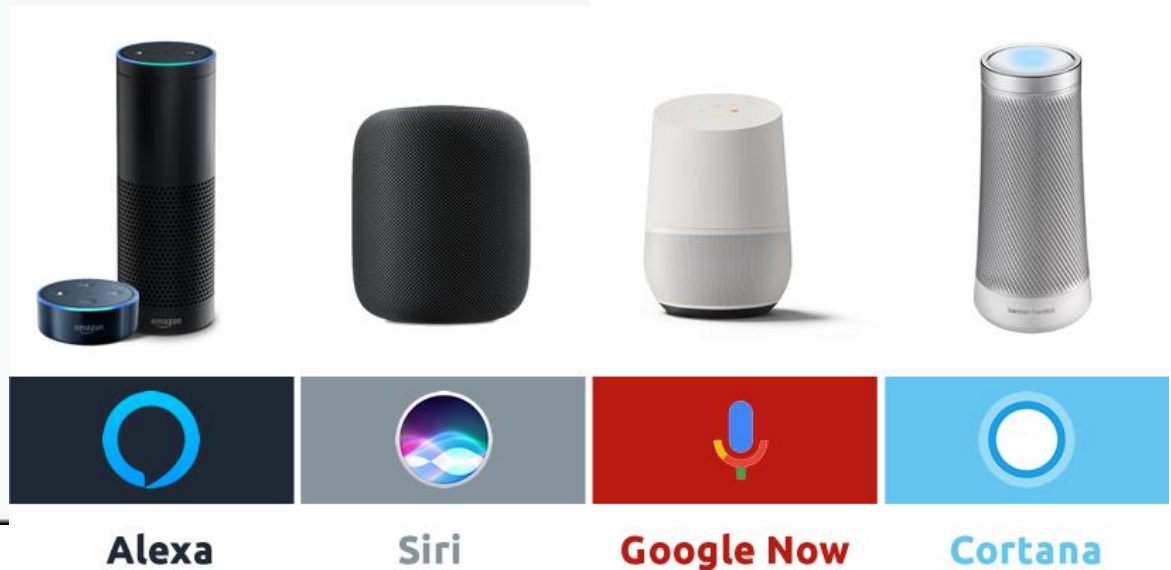
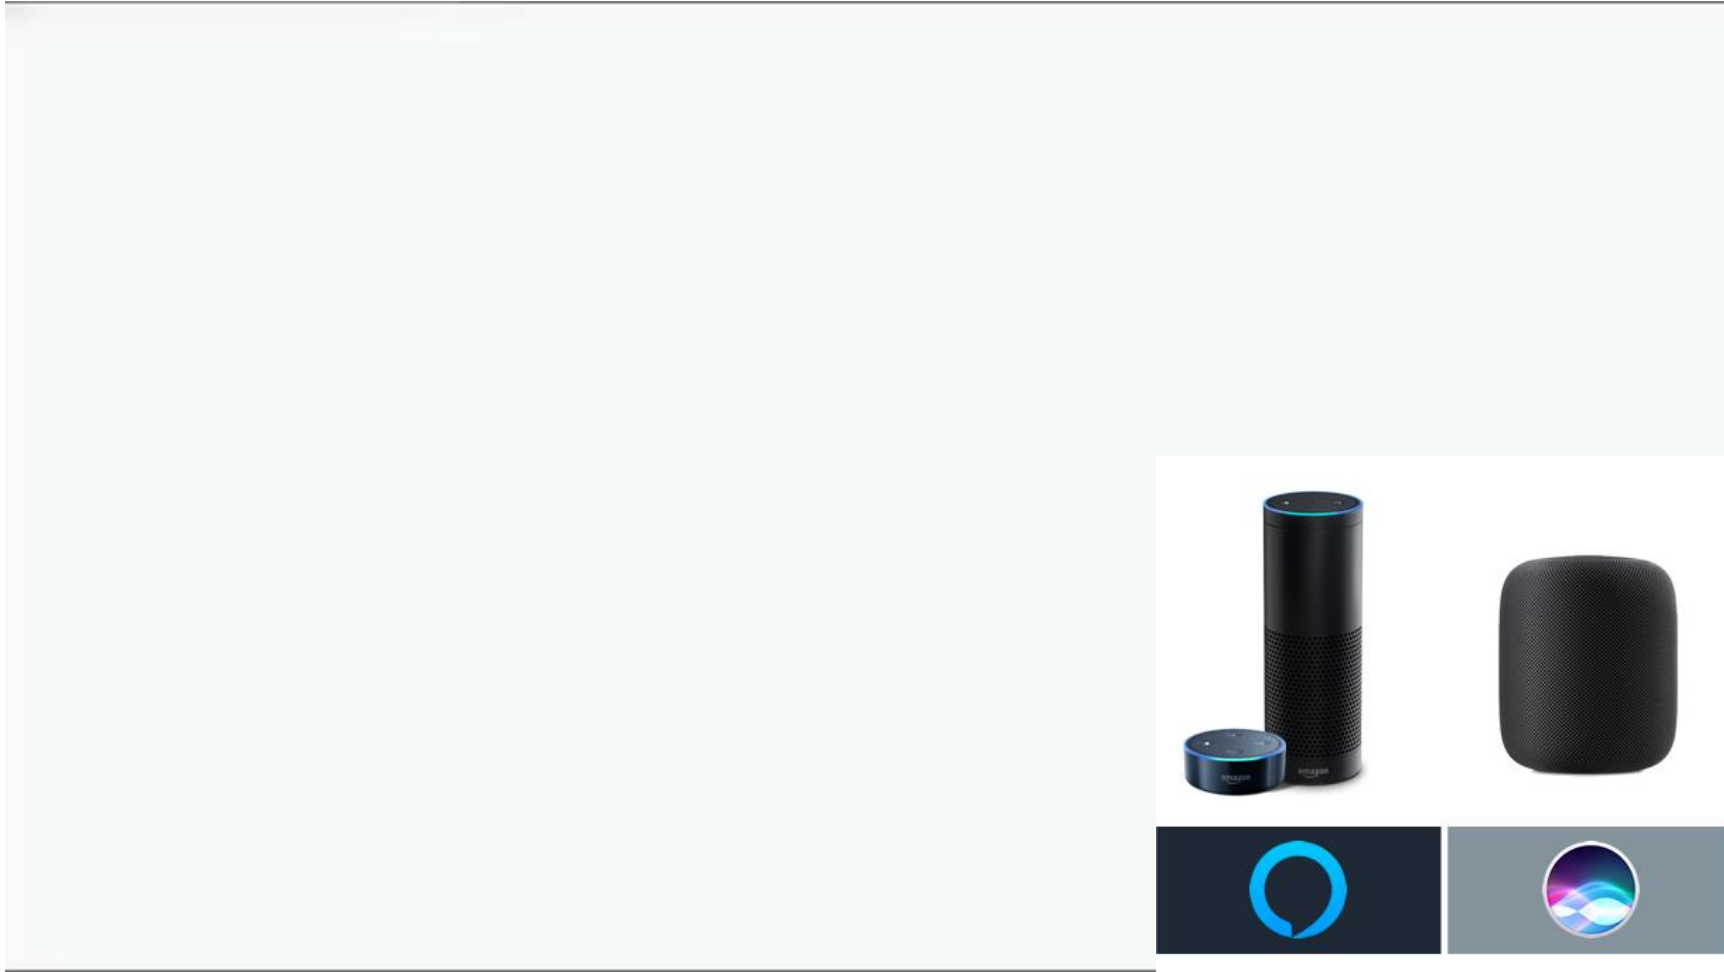
© 2016 PathAI

Ronneberger, O., Fischer, P., and S. S. Segmentation. In *Inter* 234-241). Springer, Cham.

Wang, Dayong, et al. "Deep learning for identifying metastatic breast cancer." arXiv preprint arXiv:1606.05718 (2016).  
<https://blogs.nvidia.com/blog/2016/09/19/deep-learning-breast-cancer-diagnosis/>



# Voice assistants: Google AI 2018



# Web app: search, recommendation, ad

The image shows a Google search interface for 'shanghai jiao tong university'. The search bar contains the text 'shanghai jiao tong university'. Below the search bar, a dropdown menu lists search suggestions: 'shanghai jiao tong university ranking', 'shanghai jiao tong university international students', 'shanghai jiao tong university school of medicine', and 'shanghai jiao tong university admission'. The 'admission' suggestion is highlighted with a red box. To the right of the search bar, a red banner reads 'Other Movies You Might Enjoy'. Below this banner, four movie posters are displayed: 'Amélie', 'Tu wama Iambien', 'Guys and Balls', and 'Mostly Martha'. Each poster has an 'Add' button and a star rating. A notification box on the right states 'Eiken has been added to your Queue at position 2.' Below the search results, there are several search results for 'Shanghai Jiao Tong University', including a Wikipedia entry which is also highlighted with a red box. At the bottom of the page, there is a large Tmall advertisement for 'New New Power Week' (新势力周) featuring a woman in a purple sweater and a 'BUY' button. The Tmall logo and slogan '理想生活上天猫' are visible at the bottom right.

Google shanghai jiao tong university

shanghai jiao tong university ranking  
shanghai jiao tong university international students  
shanghai jiao tong university school of medicine  
shanghai jiao tong university admission

About 1,000,000 results (0.34 seconds)

Scholarly articles for shanghai jiao tong univer  
Shanghai Jiao Tong University - Wang - Cited by 21  
Shanghai Jiao-Tong University - Xue - Cited by 14  
Nanosheet-constructed porous TiO<sub>2</sub>-B for advanced ... - Liu - Cite

Shanghai Jiao Tong University  
en.sjtu.edu.cn/ ▾  
Site Search. Home; About SJTU; Admission; Academics; Resear  
Rose to No.7 in 2016 Financial Times EMBA Ranking ...  
Programs in English · Schools · Fall 2016 SJTU Graduate ... · Scl

上海交通大学  
www.sjtu.edu.cn/ ▾ Translate this page  
全面介绍上海交通大学新闻的网站。

Shanghai Jiao Tong University - Wikipedia  
https://en.wikipedia.org/wiki/Shanghai\_Jiao\_Tong\_Universi  
Shanghai Jiao Tong University is a public research university loc  
Established in 1896 by an imperial edict issued by the Guangxu ...  
Name · History · Academics, enrollment, and staff · Organization

Amélie  
Tu wama Iambien

Add  
Add

★★★★☆  
Not Interested

★★★★☆  
Not Interested

Guys and Balls  
Mostly Martha

Add  
Add

★★★★☆  
Not Interested

★★★★☆  
Not Interested

Eiken has been added to your Queue at position 2.  
This movie is available now.  
Move To Top Of My Queue

天猫 聚划算 天猫超市 | 淘抢购 电器城 司法拍卖 淘宝心选 兴农

新势力周  
哎哟!辣么好看  
时髦新品, 满减200-15

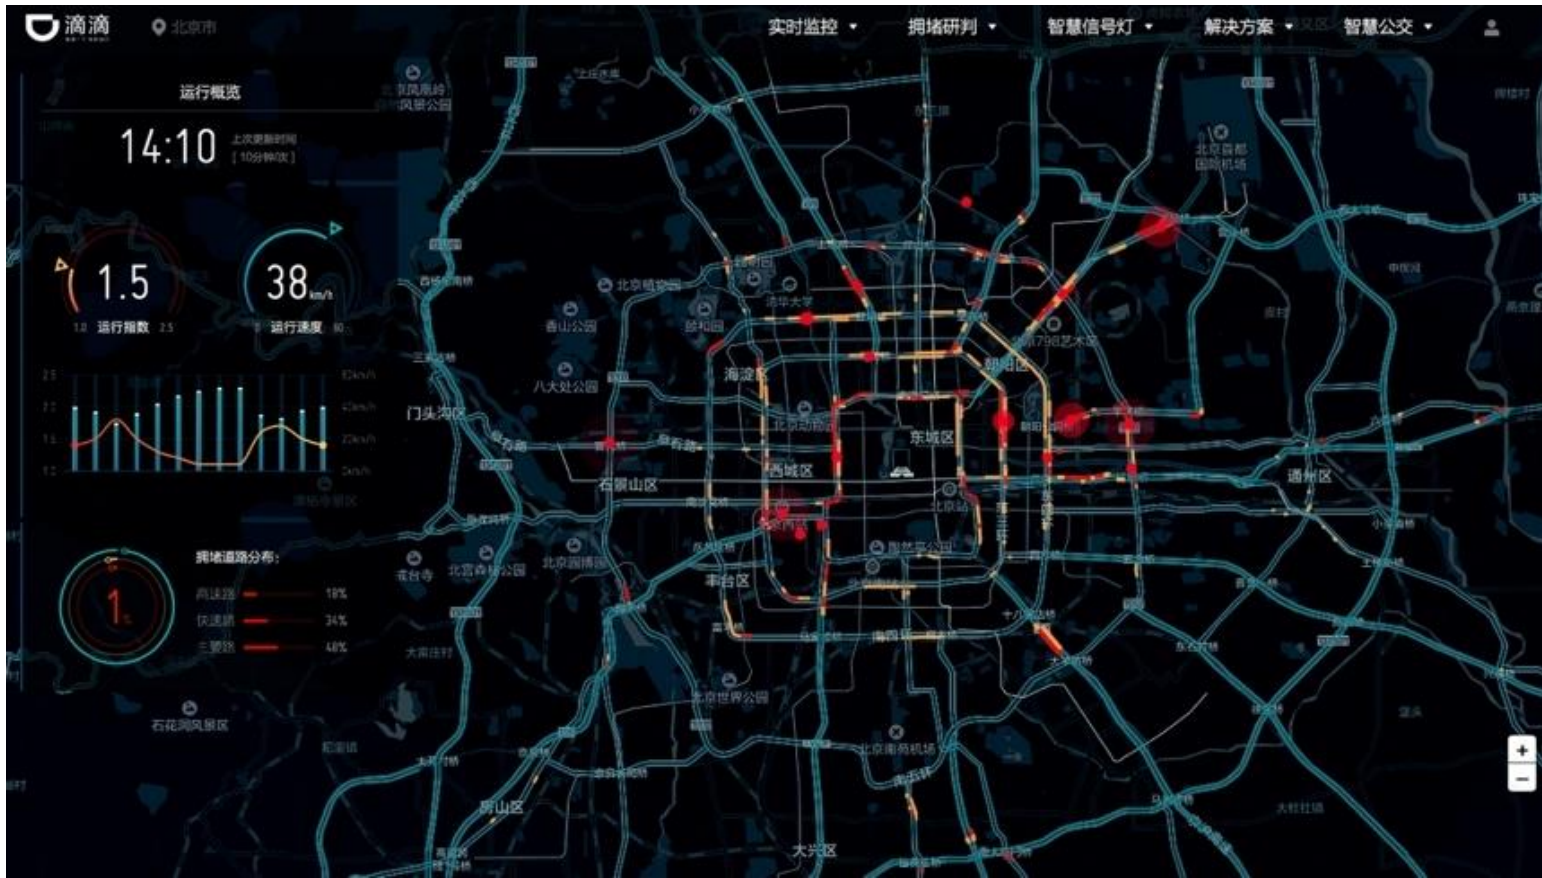
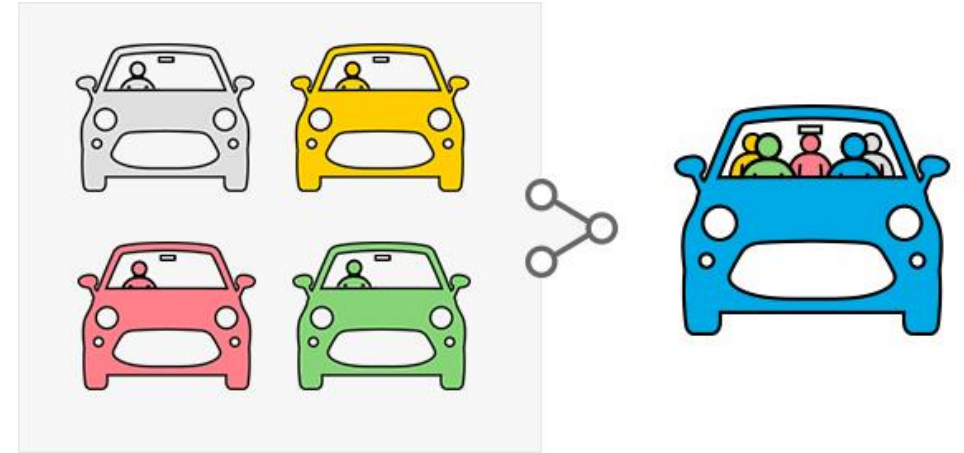
BUY

TMALL天猫 理想生活上天猫

1/6

# Alleviate traffic congestion

- Ride sharing
- Disperse traffic



# Exoskeletons



# Agriculture: Crop-dusting

- DJI drones (unmanned aerial vehicles)



# Transportation: Sorting parcels




# Boston Dynamics: Atlas | Partners in Parkour




## EXPLORER

\$74,500.00

The Spot Explorer kit puts the power of robotics into your hands and makes robotics easy, so you can focus on building your application.






# SMART CITY COMPONENTS



# What Can AI Do?

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



- ✓ ■ Play a decent game of table tennis?
- ✓ ■ Play a decent game of Jeopardy?
- ✓ ■ Drive safely along a curving mountain road?
- ? ■ Drive safely across Pittsburgh?
- ✓ ■ Buy a week's worth of groceries on the web?
- ✗ ■ Buy a week's worth of groceries at a local market?
- ? ■ Discover and prove a new mathematical theorem?
- ✗ ■ Converse successfully with another person for an hour?
- ? ■ Perform a surgical operation?
- ✓ ■ Put away the dishes and fold the laundry?
- ✓ ■ Translate spoken Chinese into spoken English in real time?
- ? ■ Write an intentionally funny story?





# Intelligent Agents


# Agents and environments

- Agents interact with environments through sensors and actuators
- 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




# Pac-Man as an Agent





# Environment 1: Pac-Man

- Performance measure
  - -1 per step; +10 food; +500 win; -500 die; +200 hit scared ghost
- Environment
  - Pacman dynamics (incl ghost behavior)
- Actuators
  - North, South, East, West, (Stop)
- Sensors
  - Entire state is visible



# Environment 2: Automated taxi

- Performance measure
  - Income, happy customer, vehicle costs, fines, insurance premiums
- Environment
  - streets, other drivers, customers
- Actuators
  - Steering, brake, gas, display/speaker
- Sensors
  - Camera, radar, accelerometer, engine sensors, microphone



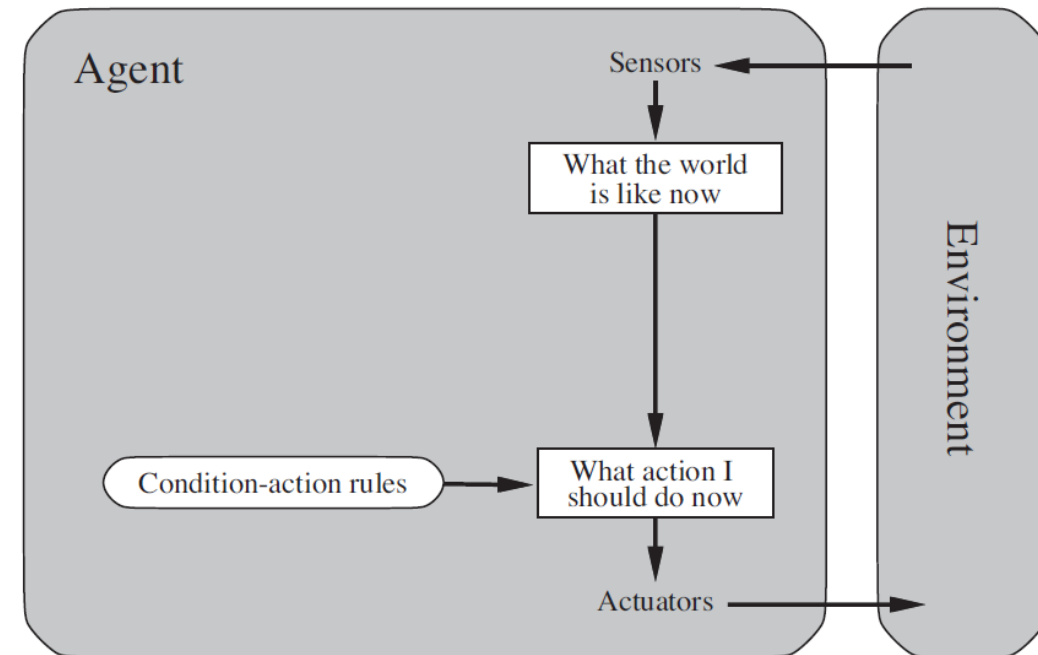

# Environment Types

	Pacman	Taxi
Fully or partially observable		
Single agent or multi-agent		
Deterministic or stochastic		
Static or dynamic		
Discrete or continuous		



# Simple reflex agents

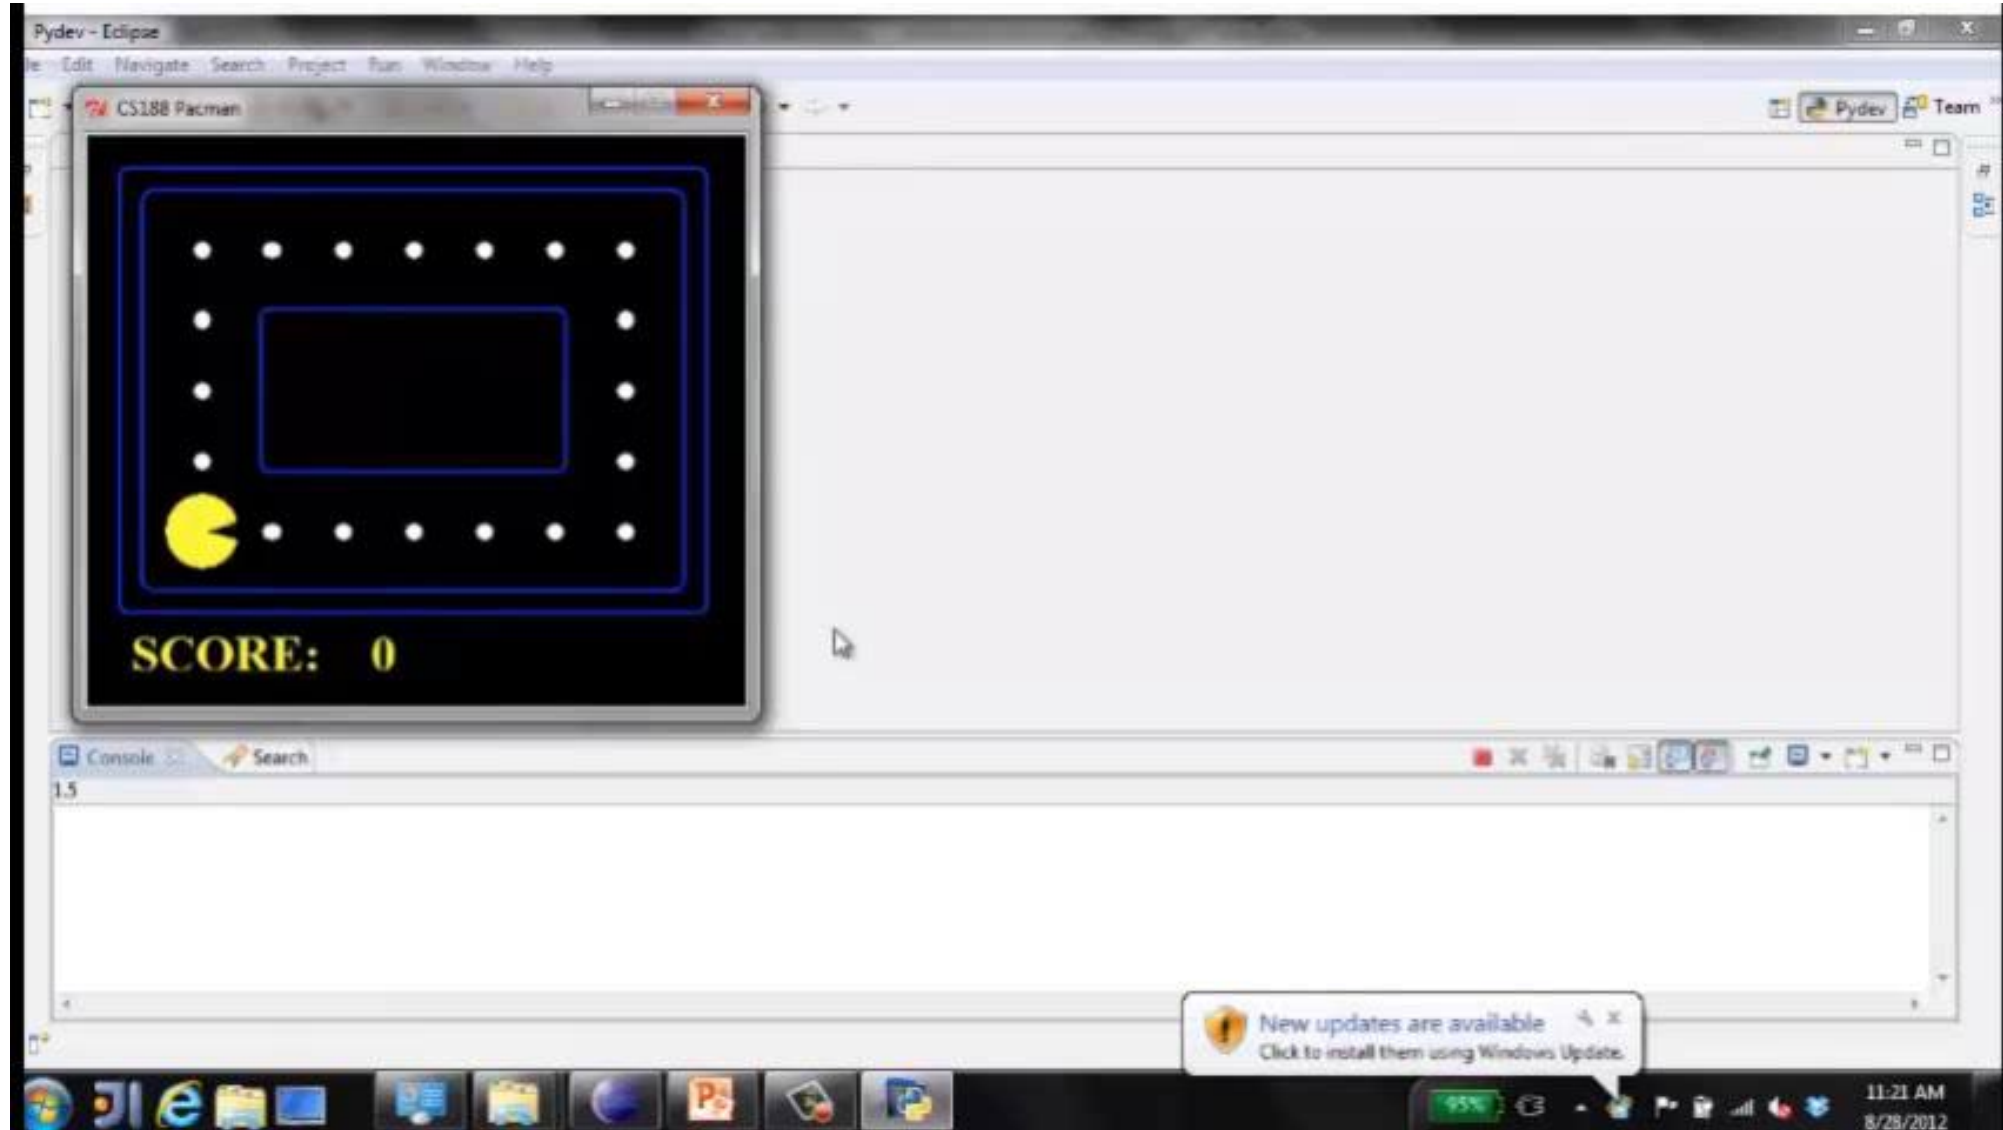
- Reflex agents:
  - Choose action based on current percept (and maybe memory)
  - May have memory or a model of the world's current state
  - Do not consider the future consequences of their actions
  - **Consider how the world IS**
- Can a reflex agent be rational?



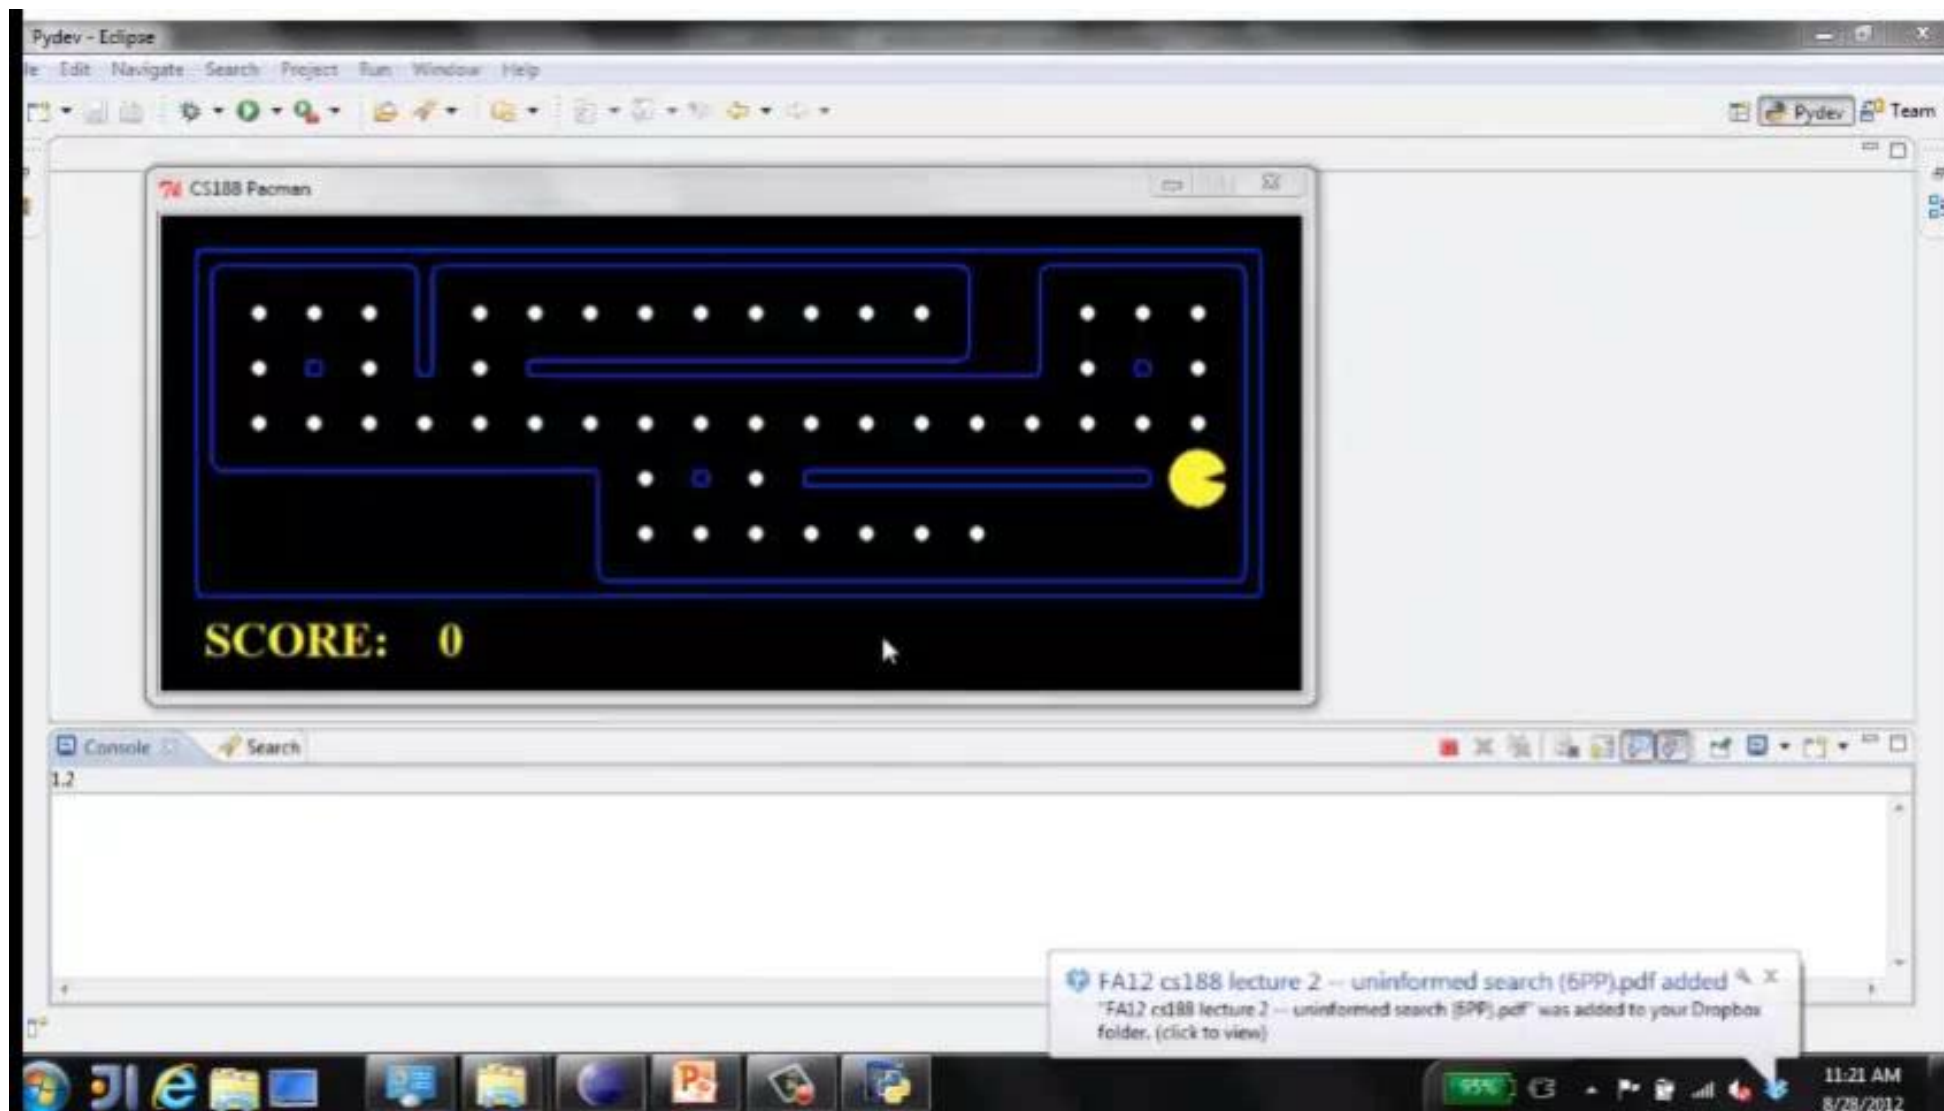
[Demo: reflex optimal (L2D1)]

[Demo: reflex optimal (L2D2)]

# Video of Demo Reflex Optimal



# Video of Demo Reflex Odd



# Summary

- What is AI and ML
- An example of AI but not ML
  - A\* algorithm
- Foundation of AI
- History of AI
- What can AI do
  - Many applications in different industries/many aspects of life
- Intelligent agents
  - reflex agents

**Shuai Li**

<https://shuaili8.github.io>

## Questions?