CS 486/686
Introduction to Artificial Intelligence

Wenhu Chen

Lecture 1

Outline

Learning goals

- Get to know a bit about the instructor and one or more classmates.
- Understand the topics of the course
- Name an application of AI.

Outline

Who am I?

My name is Wenhu Chen. I am from GuangDong, China.

Work history (recent):

- Assistant Professor, University of Waterloo
- Canada CIFAR AI Chair

Previous:

- Research Scientist at Google Research/Brain

Education history:

- PhD, Computer Science, University of California, Santa Barbara

Research: natural language processing, deep learning, multi-modality

Who are the TAs?

The TAs will grade your homework, host TA office hours, and grade the final exam.

- Dhanraj, Varun (vdhanraj@)
- Hu, Theo (z97hu@)
- Jafari, Aref (a22jafar@)
- Luo, Yudong (y264luo@)
- Mokhtari, Sabrina (s4mokhta@)
- Xiong, Rui Ming (rmxiong@)
- Yang, Jheng-Hong (j587yang@)
- Zhang, Dake (d346zhang@)
- Zhu, Shuhui (s223zhu@)

Outline

Learning goals

Introductions

Topics in CS 486/686

Course Outline

Assessments

What is Artificial Intelligence?

Topics in CS 486/686 (1/3)

This part will be covered by Wenhu

Uncertainty Estimation:

1: Introduction to AI.

2: Uncertainty and Probability

3: Independence of Bayesian Networks

4: Bayesian Networks

Intro to Machine Learning:

5: Intro to Machine Learning

6: Intro to Neural Networks 1

7: Intro to Neural Networks 2

8: Intro to Neural Networks 3

Topics in CS 486/686 (2/3)

This part will be covered by Pascal

Decision Process and Reinforcement Learning:

9: Markov Decision Theory
10: Reinforcement Learning
11: Deep Reinforcement Learning
12: Model-based Reinforcement Learning
13: Multi-armed Bandits
14: Game Theory I
15: Game Theory II
16: Multi-agent RL

Topics in CS 486/686 (3/3)

This part will be covered by Wenhu

Search Algorithm:

17: Uninformed Search
18: Heuristic Search
19: Constraint Satisfaction Problems
20: Local Search

Deep Learning:

21: Unsupervised Learning
22: Self-supervised Learning
23: NLP & CV Applications
24: Recap

Part 1: Uncertainty Estimation

- Sum Rule, Product Rule, Chain Rule, Bayes Rule
- Understand the random variable independence
- D-separation Principle for Testing independence
- Understand how to construct Bayesian Networks
- Learn how to answer query about Bayesian Networks

Part 2:Machine Learning & Neural Networks

- Understanding Supervised/Unsupervised Learning, Bias-Variance Trade-off
- Understanding Empirical Risk Minimization
- Basics of Neural Networks
- Backward Propgation Algorithm in Neural Networks
- Optimization in Neural Networks

Part 3: Decision Process & Reinforcement Learning

- Understand Decision Network, Actions, and Utility
- Bandit & Reinforcement Learning
- Game Theory
- Multi-agent Reinforcement Learning

Part 4: Search

- Learn the generic search algorithm
- Discuss uninformed search algorithms like DFS, BFS
- Discuss heuristic search algorithms like the A* algorithm
- understand the complexity, and completeness of search algorithms
- Understand constraint satisfaction problem
- Understand how to solve CSP with the arc-consistency algorithm
- Learn about local search algorithms like annealing, genetic, etc

Part 5: Deep Learning

- Unsupervised Learning
- Self-supervised Learning
- LLMs like BERT, GPT-2, T5, etc
- Applications of Deep Learning in NLP & CV

Outline

Main Website

https://cs.uwaterloo.ca/~wenhuche/teaching/cs486-2024/

- Frequently updated with learning materials and announcements.
- Includes all the deadlines, release date, slides and notes.
- If there is any difference between main website and LEARN, use the information from the main website.

Piazza

Best way to reach me = office hour

Second-best way to reach me = private Piazza post

- Ask questions related to course
- You're encouraged to answer other students' questions
- We will try to respond in 2 business days

TA Support

If you have questions regarding assignments or learning materials, you can also consult the corresponding TAs.

- Uncertain Estimation (L2 - L4): Yudong Luo
- Intro to ML (L5 - L8): Aref Jafari and RuiMing Xiong
- Decision Making and Reinforcement Learning (L8 - L16): Zeou Hu, Shuhui Zhu, and Varun Dhanraj
- Search Algorithm (L17 - L20): Dake Zhang and Sabrina Mokhtari
- Deep Learning (L21 - L23): Jheng-Hong Yang

Office Hours

- Instructor Office Hour: weekly on Thursday 3 - 4PM
- TA office Hour: before the assignment due date (TBA)

Outline

Assignments

There are 3 assignments in total (10% for each one).

- Assignment 1: Part 1 and Part 2
- Assignment 2: Part 3
- Assignment 3: Part 4 and Part 5

Quiz (Undergrads)

There are 10 quiz in total (2% for each one).

- The quiz date is on the website
- Each quiz contains 5-10 multi-choice questions
- The quiz is only open for a few hours after the class
- No late submission is allowed

Project (Grads)

Solve a problem with AI (or for AI) and write a report

- Topics are not restricted, it could be reinforcement learning, image classification, natural language processing, human-computer interaction, biomedical, etc.
- It does not need to reach publication-level novelty.
- It cannot be trivial, like cloning some code from existing public GitHub.
- Final report due after the final exam.

Exam

Final: 2.5 hours, date/time (TBD by the registar office)

- Final exam designed by Wenhu and Pascal
- Every student must attend the final exam, no makeup exam
- All students must pass the final exam (50%)
- Final exam will cover all the materials, roughly uniformly distributed across the 5 parts

Score Breakdown

- For CS486: 30% from Assignment (10% $\times$ 3), 20% from Quiz, 50% from Final Exam
- For CS686: 30% from Assignment (10% $\times$ 3), 30% from Project, 40% from Final Exam

Outline

The Goal of AI

Artificial Intelligence:

- Enable machines to sense, comprehend, act and learn with human-like levels of intelligence
- Combined with computer science to solve challenging problems
- Not confined to methods that are biologically observable

The History of AI: Alan Turing

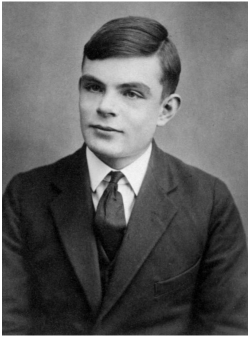

Alan Turing: 1912-1954 was known as the father of computer.
Also asked 'can machines think?', proposed the imitation game.

The History of AI: Alan Turing

- Turing illustrated his ideas on machine intelligence by reference to chess.
- Turing believes that machines can learn from experience as well as solve new problems.
- He needs to content himself with theory in absence of computer to run chess program.
- 50 Years later, Deep Blue from IBM was invented to play good chess, which is precisely what Turing predicted.

The History of AI: Dartmouth Conference

- During Dartmouth conference in mid-1950s, Jon MacCarthy proposed the term "Artificial Intelligence"
- The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it
- McCarthy also created Lisp, the standard programming language used in robotics

The History of AI: the Turing Test

The interrogator is connected to one person and one machine via a terminal, therefore can't see her counterparts. Her task is to find out which of the two candidates is the machine, and which is the human only by asking them questions. If the machine can 'fool' the interrogator, it is considered intelligent.

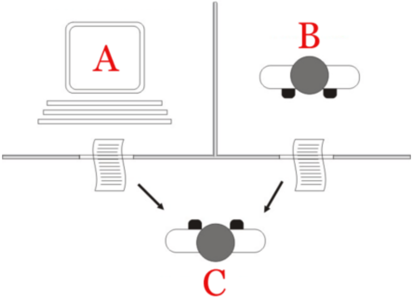

The History of AI: the Turing Test

ELIZA - (1966 Joseph Weizembaum, Lisp)

ME: Hi
ELIZA: How do you do? What brings you to see me?
ME: I'm thinking of taking a course in artificial intelligence.
ELIZA: Why do you say that?
ME: I'm worried about it
ELIZA: Are you worried often?
ME: when I take AI courses
ELIZA: What happened when you take ai courses?
ME: I usually spend long nights doing assignments
ELIZA: Earlier you said you felt worried?
ME: about AI courses
ELIZA: Maybe your plans have something to do with this.
ME: no
ELIZA: Why not?
ME: because I think you are a computer
ELIZA: Is it because you think I am a computer that you came to me?

Blocks World - SHRDLU - Winograd 1968

- first autonomous physical interaction
- blocks world only

Pick up a big red block
OK
grasp the pyramid
I DON'T UNDERSTAND WHICH PYRAMID YOU MEAN
What does the box contain?
THE BLUE PYRAMID AND THE BLUE BLOCK.
What is the pyramid supported by?
THE BOX.
How many blocks are not in the box?
FOUR OF THEM

Game Playing

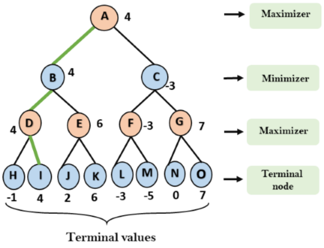

- Game playing is also a very important metric to evaluate intelligence
- Min-Max Game of two players, one wants to maximize the chance of wining, the other aims to minimize it
- Game playing can be formulated as a search problem with nodes representing the state

Chess

- game tree has more than 10^{100} nodes
- IBM - Deep Blue uses lookahead search and complex evaluation function to predict the outcome.
- In 1997, IBM Deep Blue defeated the world champion Gary Kasparov in a 6-match game.
- Deep Blue AI 2022 (https://www.scientificamerican.com/article/20-years-after-deep-blue-how-ai-has-advanced-since-conquering-chess/)

Go (Weiqi)

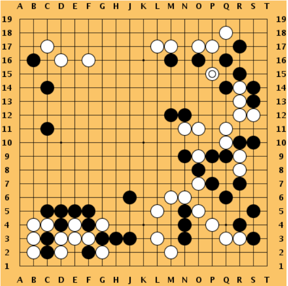

- game tree has more than 10^{360} nodes, which takes longer for AI to master
- AlphaGo (Deepmind): Silver et al. (2016). Mastering the game of Go with deep neural networks and tree search. Nature, 529(7587), 484-489.
- Deep neural networks; supervised learning; reinforcement learning.

AlphaGo

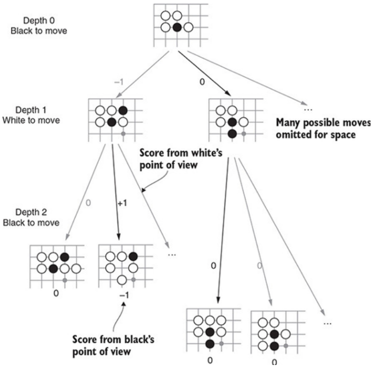

- AlphaGo combines machine learning and tree search techniques
- Alphago uses Monte-Carlo tree search algorithm, guided by a "value network" and "policy network" (covered in our lectures)
- AlphaGo learns by playing against both humans and itself

Poker

- Play with uncertainty. Must model opponent(s).
  Care about long-term payoff.
- Bowling, M., Burch, N., Johanson, M., & Tammelin, O. (2015). Heads-up limit hold'em poker is solved. Science, 347(6218), 145-149.
- Brown, N., & Sandholm, T. (2019). Superhuman AI for multiplayer poker. Science, 365(6456), 885-890.

Atari Games

- Mnih et al. (2013). Playing atari with deep reinforcement learning. arXiv preprint arXiv:1312.5602.
- Reinforcement learning; Convolutional neural network; High-dimensional sensory input. Previous approaches use hand-crafted visual features.
- Outperforms previous approaches; Surpasses a human expert on 3/7 Atari 2600 games.

AlphaFold

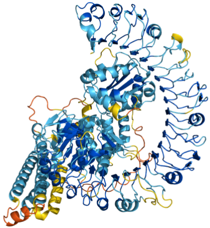

- Synthesize structures of unknown protein (previously requires painstaking months to years human labor)
- Jumper et al. (2021) Highly accurate protein structure prediction with AlphaFold
- Benefit greatly from the growth of the Protein Data Bank
- Uses a model called EvoFormer to make structure prediction

The surge Deep Neural Network

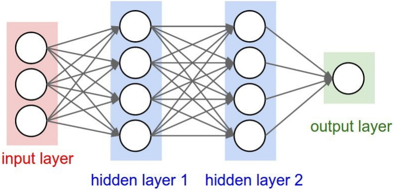

- Old-fashioned ML used decision tree, SVM, linear regression, boosting algorithms.
- Neural Network exist for quite a while. It really took off with AlexNet (Alex et al. ImageNet Classification with Deep Convolutional Neural Networks)
- Increasing the depth of neural networks to hundreds of layers.

AlexNet (2012)

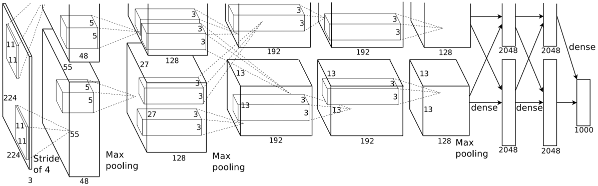

Figure 2: An illustration of the architecture of our CNN, explicitly showing the delineation of responsibilities between the two GPUs. One GPU runs the layer-parts at the top of the figure while the other runs the layer-parts at the bottom. The GPUs communicate only at certain layers. The network's input is 150,528-dimensional, and the number of neurons in the network's remaining layers is given by 253,440–186,624–64,896–64,896–43,264–4096–4096–1000.

- 60 million parameters, 650K neurons with five conv layers
- Reaching the top-1 error rate of 37.5% on on ILSVRC-2012 competition
- AlexNet was trained on 1M labled images
- The network trains between five and six days to train on two GTX 580 3GB GPUs

Distbelief (2012)

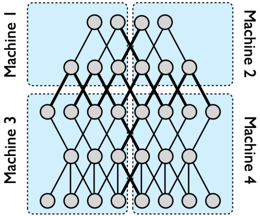

Figure 1: An example of model parallelism in DistBelief. A five layer deep neural network with local connectivity is shown here, partitioned across four machines (blue rectangles). Only those nodes with edges that cross partition boundaries (thick lines) will need to have their state transmitted between machines. Even in cases where a node has multiple edges crossing a partition boundary, its state is only sent to the machine on the other side of that boundary once. Within each partition, computation for individual nodes will the parallelized across all available CPU cores.

- Distbelief proposes to shard model and data across Google clusters with Downpour SGD and Sandblaster L-BFGS
- Utilizing 512 CPU cores to train single large neural networks

ImageNet Challenge

IMAGENET

- 1,000 object classes (categories).
- Images:
  - 1.2 M train
  - 100k test.

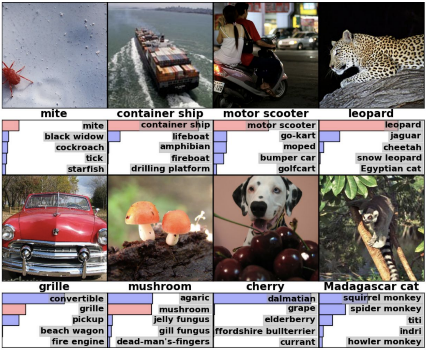

Figure: ImageNet Challenge

A standard classification problem to categorize an image into 1 of 1000 object classes.

Image Classification

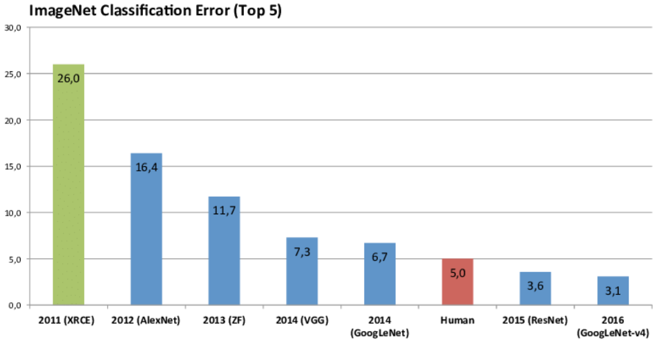

Figure: ImageNet winner over time

Before 2011, the curve was relatively flat. In 2012, AlexNet was proposed as the first deep neural network to address the long-standing problem. Already solved.

Image Generation

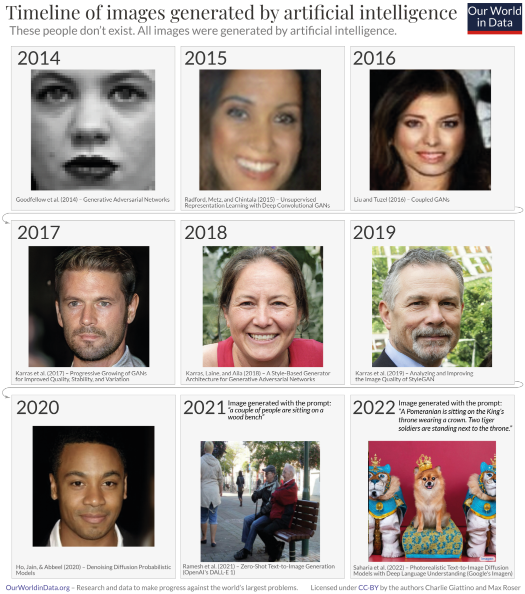

Supervised Learning (Human Annotation)

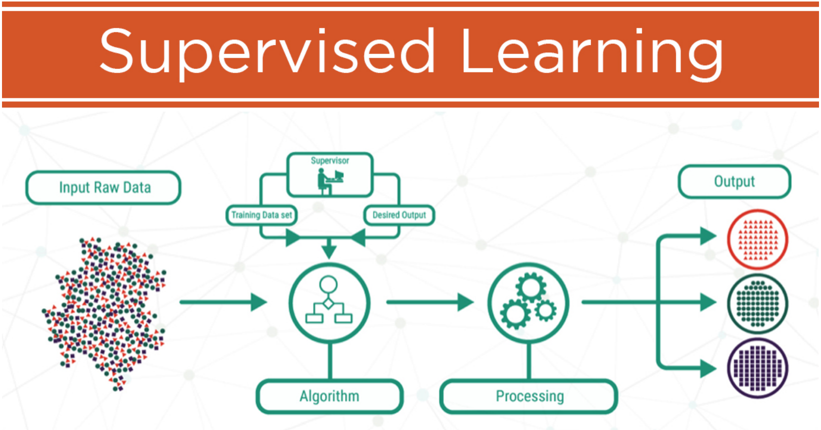

Self-Supervised Pre-training

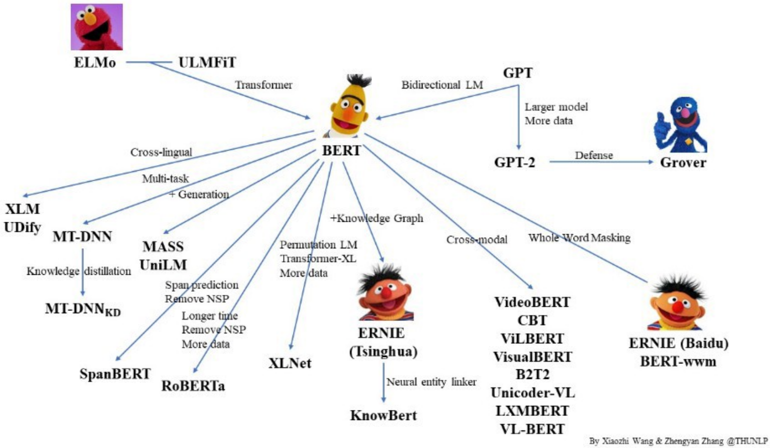

Instead of training on annotated data, the model trains on the crawled web with self-supervised objective

The Rise of Computation

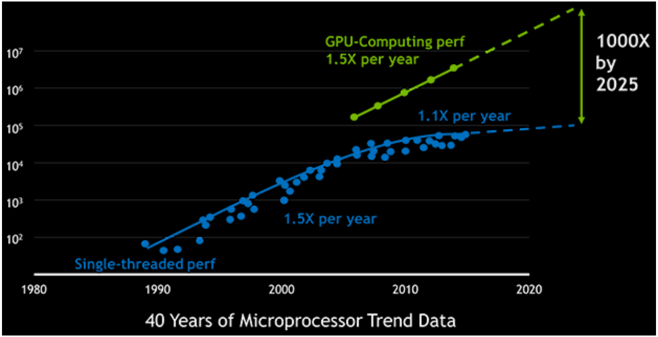

New Moore's Law: every year, the compute capability increases by 1.5x. The underlying computation power drives advancement of AI.

Large Language Models

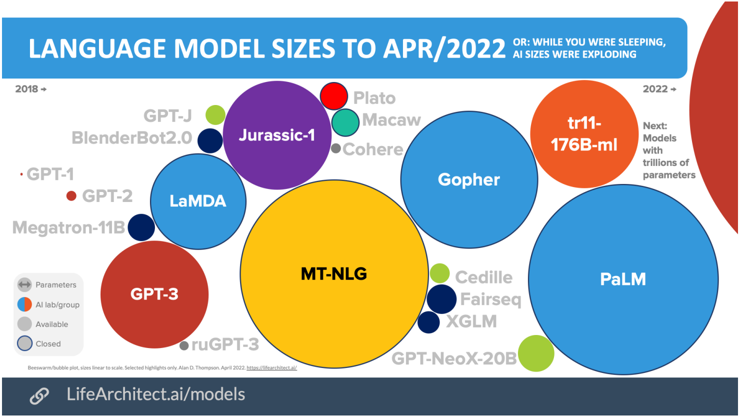

- the current text/language models grow larger and larger
- Google's PaLM model is already achieving 540B parameters

Large Language Models

The current large language models have shown striking performance in many aspects. The most recent ChatGPT model from OpenAI can give astonishingly accurate responses.

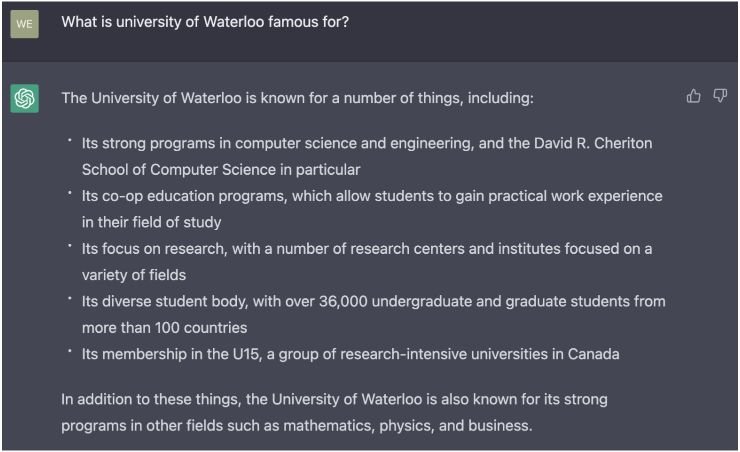

Large Language Models

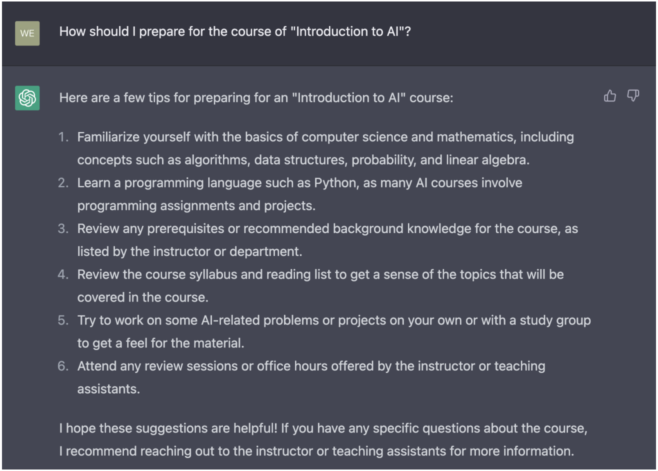
How should I prepare for the course of "Introduction to AI"?

Here are a few tips for preparing for an "Introduction to AI" course:

1. Familiarize yourself with the basics of computer science and mathematics, including concepts such as algorithms, data structures, probability, and linear algebra.
2. Learn a programming language such as Python, as many AI courses involve programming assignments and projects.
3. Review any prerequisites or recommended background knowledge for the course, as listed by the instructor or department.
4. Review the course syllabus and reading list to get a sense of the topics that will be covered in the course.
5. Try to work on some AI-related problems or projects on your own or with a study group to get a feel for the material.
6. Attend any review sessions or office hours offered by the instructor or teaching assistants.

I hope these suggestions are helpful! If you have any specific questions about the course, I recommend reaching out to the instructor or teaching assistants for more information.

Learning goals

- Get to know a bit about the instructor and one or more classmates.
- Understand the topics of the course
- Learn the history of AI
- Learn the history of deep learning