



Lecture 1

Introduction

Lin ZHANG, PhD
School of Software Engineering
Tongji University
Fall 2024



Course Info

Contact Information

Room 418L, Jishi Building

TA: Linfei Li, cslinfeili@tongji.edu.cn

All material can be found at

<https://cslinzhang.github.io/home/>



Materials

- Major materials
 - My slides
- References
 - 《计算机视觉：原理算法与实践》，张林等，2024年10月（拟）
 - 《机器学习》，周志华，2016
 - 《统计学习方法》（第2版），李航，2019
 - Some papers



Examination

- Homework 30%: 3 times, and each time 10%.
- Paper reading and presentation 20%
 - Read a paper related to machine learning and do a presentation
- Final report and presentation 50%
 - Select a problem related to your research direction, try to solve it with machine learning techniques, write an essay and finally do a presentation
- Being absent $\geq 1/3$ lectures, you will fail this course



Arrangement of Lectures (temporarily)

- Basic Concepts and Model Evaluation
- AdaBoost and Cascade Structure
- Principle Component Analysis
- Sparse Representation based Classification
- Linear Model
- Neural Network and CNN
- Fundamentals of Convex Optimization
- Support Vector Machines
- Transformer-based Object Detection
- Applications of CNN




A little history about AI

人工智能

1956年，**麦卡锡**召集哈佛大学、麻省理工学院、IBM公司、贝尔实验室的研究人员召开**达特茅斯会议**正式提出“**人工智能**”



2006年达特茅斯会议当事人重聚，左起：**摩尔**、**麦卡锡**、**明斯基**、**赛弗里奇**、**所罗门诺夫**



John McCarthy
人工智能之父


人工智能是指计算机系统具备的能力，该能力可以履行原本只有依靠人类智慧才能完成的复杂任务



A little history about AI

什么是人工智能?

- 指由人制造出来的机器所表现出来的智能
 - 通常指通过计算机程序来呈现人类智能的技术
- **遗憾的是，“智能”本身难以定义清楚!**
 - 行为定义的智能 Behavior defined intelligence
 - 即**图灵测试**定义的智能（不管内涵，只管外延）





A little history about AI

什么是人工智能?

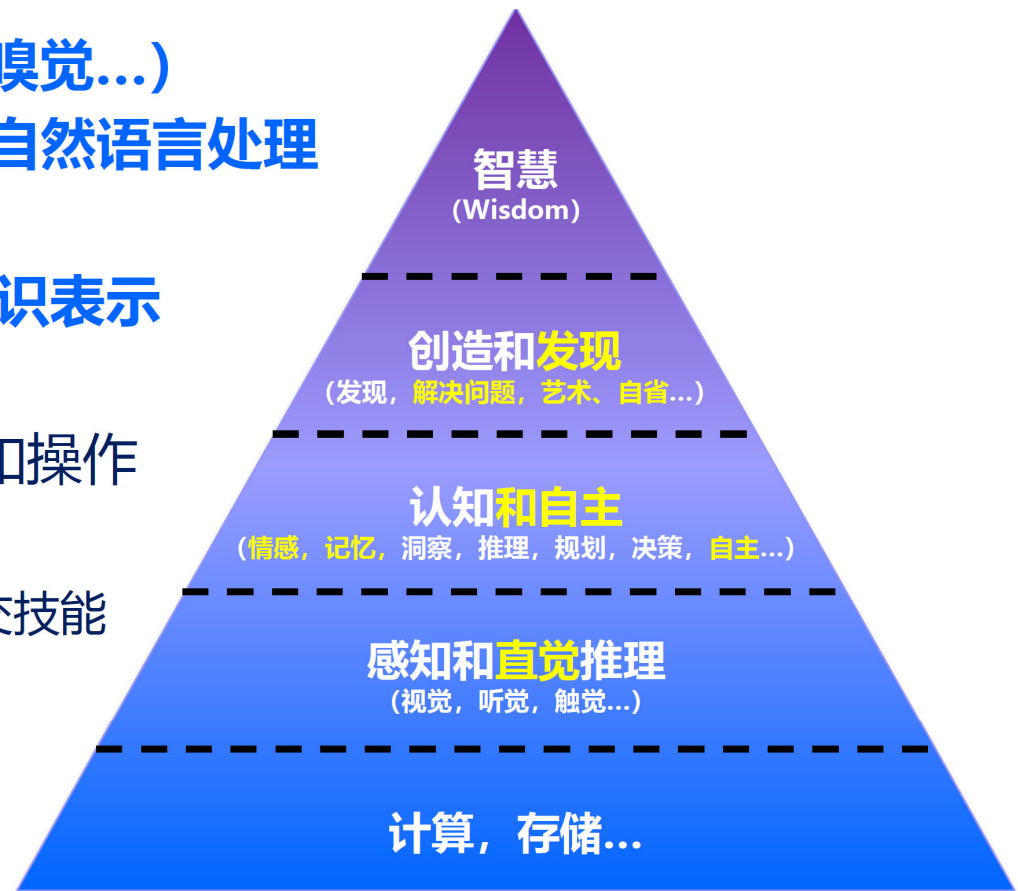
- 行为定义的智能 Behavior defined intelligence
 - 系统的表现是智能的
- 在计算机领域，人工智能是指对“智能代理”的研究
 - 任何可以**感知环境**并**采取行动**以最大可能达成其**特定目标**的任何设备都是智能代理【维基百科】



A little history about AI

人工智能的内涵和任务

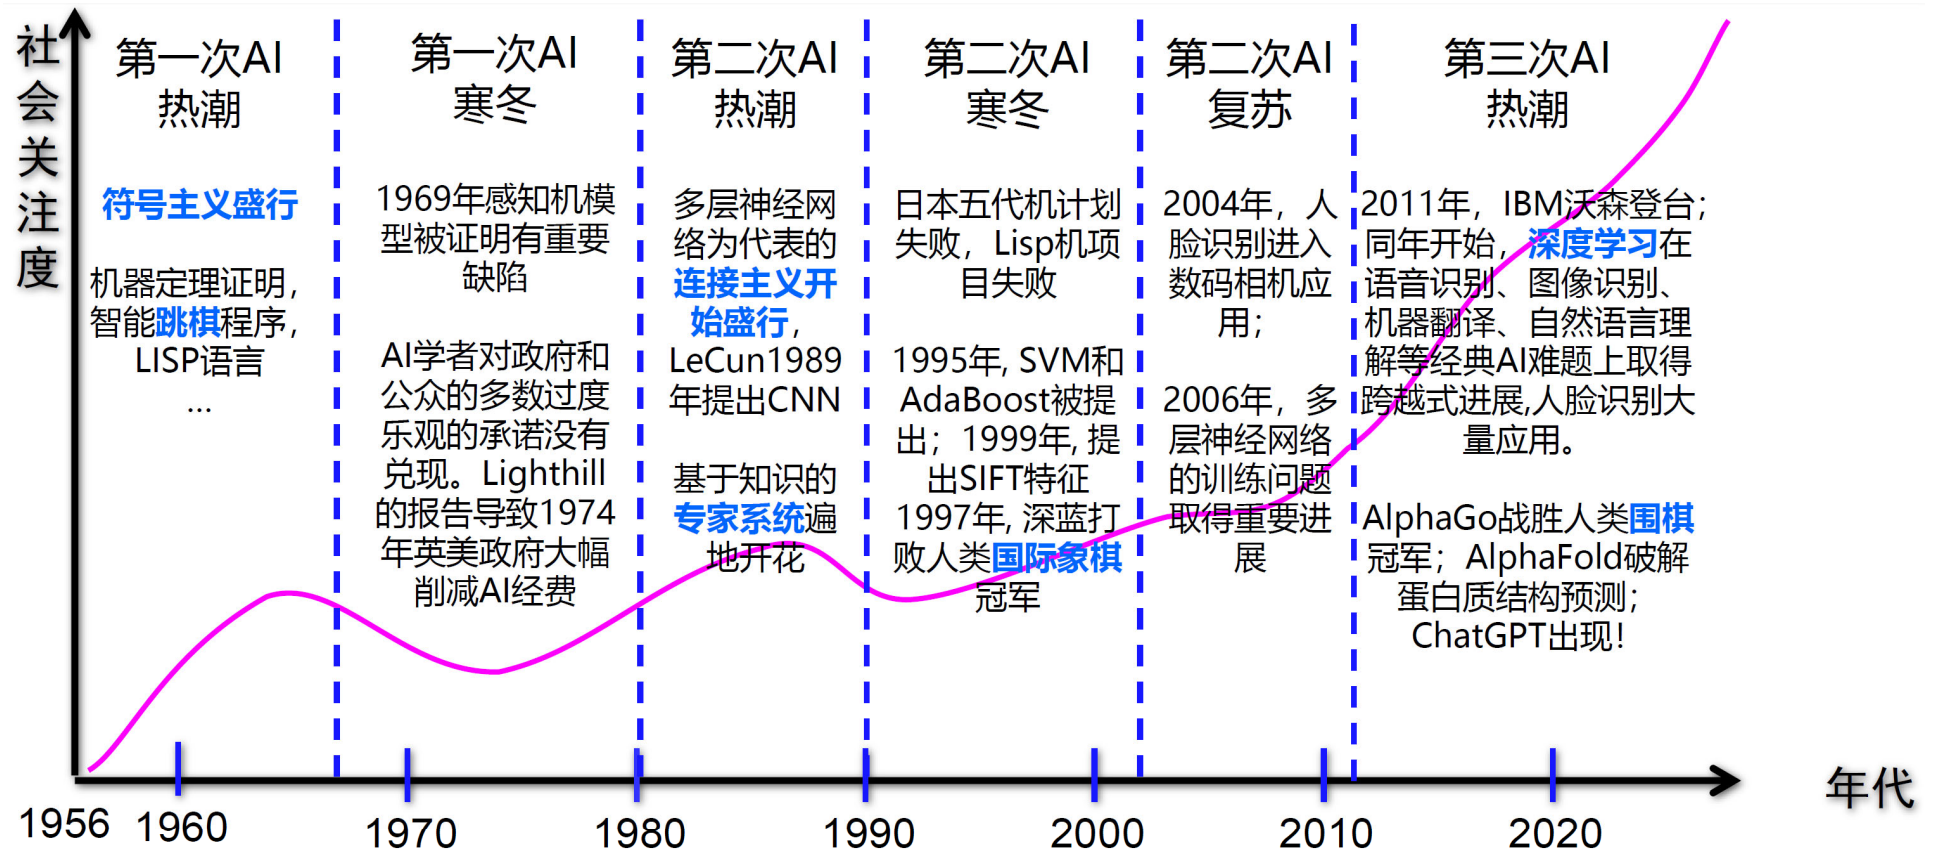
- Perception感知 (视觉, 听觉, 嗅觉...)
- Natural language processing自然语言处理
- Learning学习
- Knowledge representation 知识表示
- Planning规划
- Motion and manipulation运动和操作
- Social intelligence社会智能
 - Affective Computing情感计算/社交技能
- Reasoning, problem solving
- Creativity创造力
- General intelligence通用智能





A little history about AI

人工智能发展的历史沿革



符号主义、知识驱动的AI方法论

连接主义、大规模数据驱动的AI方法论



A little history about AI

三次人工智能浪潮背后的方法论

■ 第一次浪潮：基于符号的推理与搜索

- 模拟人的符号推理方式
- 搜索树：解决迷宫问题

■ 符号主义

- 符号主义的实现基础是纽威尔和西蒙提出的物理符号系统假设
 - 人类**认知和思维的基本单元是符号**，而**认知过程就是在符号表示上的一种运算**。人是一个物理符号系统，计算机也是一个物理符号系统，故可用计算机来模拟人的智能行为，即用计算机的符号操作来模拟人的认知过程
 - 实质就是模拟人的左脑抽象逻辑思维，通过**研究人类认知系统的功能机理，用符号之间的逻辑关系来描述人类的认知过程**，并把这种符号输入到能处理符号的计算机中，就可以模拟人类的认知过程，从而实现人工智能



A little history about AI

三次人工智能浪潮背后的方法论

■ 第一次浪潮：基于符号的推理与搜索

- 模拟人的符号推理方式
- 搜索树：解决迷宫问题

■ 符号主义

- 太乐观了【承诺太多，最终做不到】！

- 解决不了更复杂的现实问题
搜索空间太大了：如围棋

- **大量问题难以转化为符号推理问题**

例如：人脸识别、语音识别等模式识别问题（非结构化数据的结构化）



A little history about AI

三次人工智能浪潮背后的方法论

■ 第二次浪潮：依赖人类符号化知识的专家系统

□ 以依赖符号化知识库和符号推理的专家系统为主

□ 知识表示是人工智能的核心难题

□ 人工智能研究早期主流的知识表示方法——**符号主义的知识观**

- 从符号主义的观点来看，**认知就是符号的处理过程**，是智能的基础
- **符号化的知识表示、推理、运用**是人工智能的核心
- 知识表示：采用符号表示（实体、关系等）所有知识
- 知识推理：推理是采用启发式知识及启发式搜索对**问题求解的过程**，其过程可以用某种**形式化的语言**来描述，因而有可能建立起基于符号化知识的人类智能和机器智能的**同一理论体系**



A little history about AI

三次人工智能浪潮背后的方法论

■ 第二次浪潮：依赖人类符号化知识的专家系统

□ 多数AI系统建立在符号基础上的知识表示（知识库、知识图谱）

•例：医疗辅诊系统，症状->疾病

□ 巅峰之作：IBM 的沃森自动问答系统

• 2011年，IBM 沃森在问答竞赛 《危险边缘》（Jeopardy）上击败人类

• 类似问题：地球上最北端的机场是哪个？

• 背后的技术

自然语言处理、消息检索、知识表示、自动推理、机器学习等开放式问答技术





A little history about AI

三次人工智能浪潮背后的方法论

■ 第二次浪潮：依赖人类符号化知识的专家系统

□ 还是太乐观了

- 雄心勃勃的日本**第五代计算机**计划失败

- 美国Cyc常识知识库项目陷入困境（至少不算成功）

1984年启动，Douglas Lenat教授领衔，以手工建立知识库为主，包含了320万条人类定义的断言，涉及30万个概念，15000个谓词

- **挑战性问题：常识是否可穷尽枚举？**

- 难以解决复杂的现实问题

知识表示困境：文字识别尚可，人脸识别用什么知识表示？

□ 方法论层面

- 方法论上的悄然变迁，基于**专家知识人工设计特征**，采用**统计模式识别**和**机器学习**（包括神经网络）工具，学习较小规模数据之间统计关系，成为主流方法




A little history about AI

三次人工智能浪潮背后的方法论

■ 第二次浪潮：依赖人类符号化知识的专家系统

- 第二次浪潮末期：数据驱动的机器学习方法的崛起
- 基本原理—基于函数拟合的预测问题
 - 用[较大量的]成对的 (x_i, y_i) 数据，拟合一个带有 θ 参数的函数 f
 - 本质：学习 x 和 y 的相关性；类比：学生学习过程， x_i 是考题， y_i 是答案
 - 函数 f 经常是人工设计的，例如：线性函数 $y = f(x) = Ax$
 - 参数 θ 量相对较少（但也经常在数十万甚至数百万量级）





A little history about AI

三次人工智能浪潮背后的方法论

■ 第二次浪潮：依赖人类符号化知识的专家系统

■ 方法论层面——基于知识的特征设计

□ 两篇文章的相似度计算

The Chinese Type 520 guided missile destroyer Guanying participates in a naval parade on April 23, 2019.

William Choong, senior fellow at the Shangri-La Dialogues, said in a tweet Tuesday that the presence of both Wei and Shanahan would set up "a clash of two visions—the US Japan-led free and open Indo-Pacific and China's Asia for Asians."

Analyst Carl Schutte, a former director of operations at the US Pacific Command's Joint Intelligence Center, told CNN: "Chinese leaders now recognize the value of multilateral defense venues and want to deny the US a monopoly of great power influence."

US intentions for the region have already been interpreted wrongly.

The Pentagon has stepped up freedom-of-navigation operations to as often as weekly. And the commander of the US Pacific Air Forces said this month that Air Force jets were flying in and around the South China Sea almost daily.

Washington has also sent warships through the Taiwan Strait separating China from what it calls to renegade province several times this year.

One of Washington's Taiwan Strait operations included a US Coast Guard cutter, which later sailed into the South China Sea—sending the fifth arm of its military and its main maritime law enforcement agency into the Pacific fray.



Cooperation with the probe
In a tweet, Trump said Mueller had "unlimited access, people, resources and cooperation."

Facts First: The White House largely cooperated with the investigation, but it's wrong to say there was "unlimited" cooperation. Trump repeatedly refused a sit-down interview with Mueller's team. Some Trump campaign associates "deleted relevant communications" or gave conflicting information. Others lied to investigators and were charged with obstruction offenses.

Trump submitted written testimony about Russian meddling but refused to answer any questions about obstruction. Mueller made it clear that Trump's responses were "incomplete" and insufficient. The President's son, Donald Trump Jr., also declined an in-person interview.

At least three Trump associates were charged with lying to investigators, which is an obstructive act, and two others were charged with lying to congressional inquiries about Russian meddling.

Mueller's conflicts of interest
In a tweet, Trump said Mueller was "highly conflicted."

Facts First: Mueller did not have conflicts of interest, and Trump knows it. The Justice Department cleared Mueller of any conflicts when he took the job in 2017. Trump's top aides told him that these perceived conflicts were "ridiculous" and were not considered true conflicts.

Trump has long claimed that Mueller was conflicted for a few reasons: Because he once sought a refund from a Trump-owned golf course, because he interviewed to be FBI director after Trump fired James Comey in 2017, and because his old law firm represented key figures in the investigation.



A little history about AI

三次人工智能浪潮背后的方法论


■ 第二次浪潮：依赖人类符号化知识的专家系统

■ 方法论层面——基于知识的特征设计

□ 两张人脸的相似度计算

■ 步骤1：图像中若干个**点**形成的**微模式类型**

■ 步骤2：统计人脸上**不同微模式的出现频次**作为不同人脸的**特征表示**





A little history about AI


三次人工智能浪潮背后的方法论

■ 第三次浪潮：依赖大量数据的深度学习的方法

- 用神经网络作为映射函数 **直接学习** 从输入 x 预测输出 y
 - 较少依赖人工设计
- **题海战术** (动辄百万, 千万量级)



x







A little history about AI

三次人工智能浪潮背后的方法论

- 第三次浪潮：依赖大量数据的深度学习方法

深度学习(深度卷积神经网络)的缘起

- 生物脑中的神经网络，单个神经元的**功能**
 - 接收前面神经元的输入，汇总→决策→传递





A little history about AI

三次人工智能浪潮背后的方法论


■ 第三次浪潮：依赖大量数据的深度学习方法

深度学习(深度卷积神经网络)的缘起

■ 生物脑中的神经网络，单个神经元的**功能**

- 接收前面神经元的输入，汇总→决策→传递
- 初级视觉皮层(V1)区简单细胞

■ **功能是检测不同朝向的线段** Hubel & Wiesel, 1959, 1962, ...





A little history about AI

三次人工智能浪潮背后的方法论


- 第三次浪潮：依赖大量数据的深度学习方法

深度学习(深度卷积神经网络)的缘起

- 生物脑中的神经网络，大量神经元互联

- 视觉通路神经细胞层级感受野假设

- 响应越来越复杂的模式 → 祖母细胞理论
- 可见越来越大的(视网膜)感受野：类比从普通士兵到总司令





A little history about AI

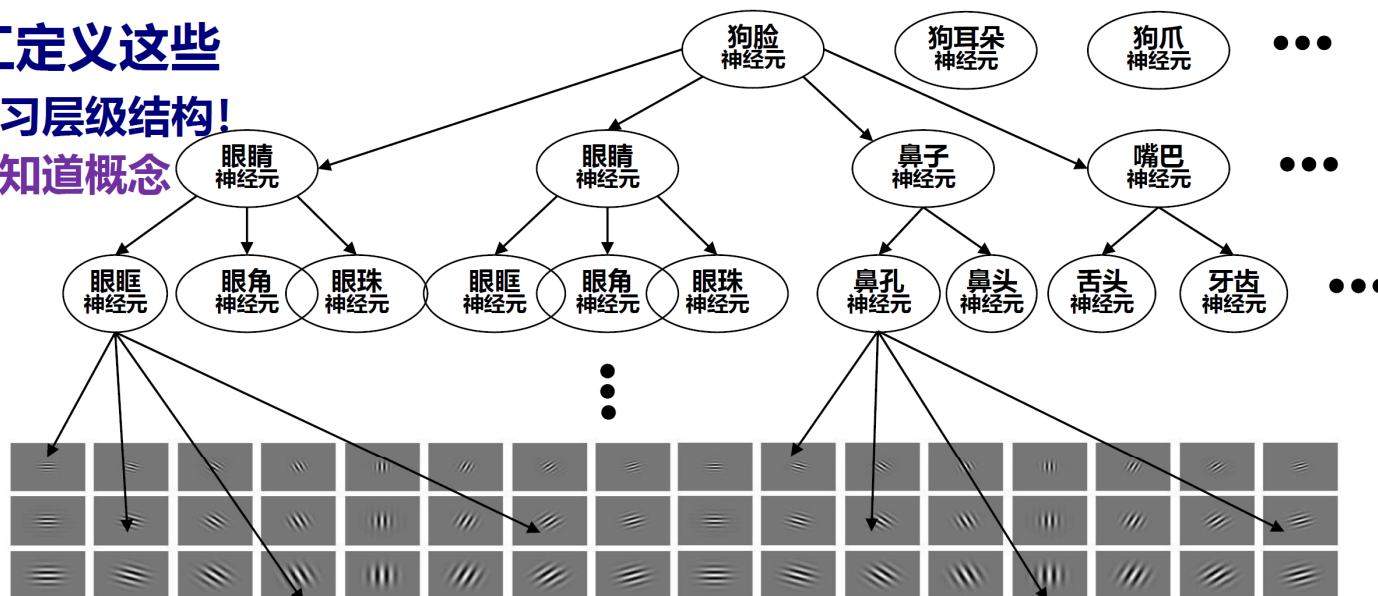
三次人工智能浪潮背后的方法论

■ 第三次浪潮：依赖大量数据的深度学习的方法

□ 所谓深度学习主要是指多层神经网络

- 一个例子：多层神经网络怎么找到狗？
- 但不需要人工定义这些

- 算法自动学习层级结构！
- 算法不需要知道概念



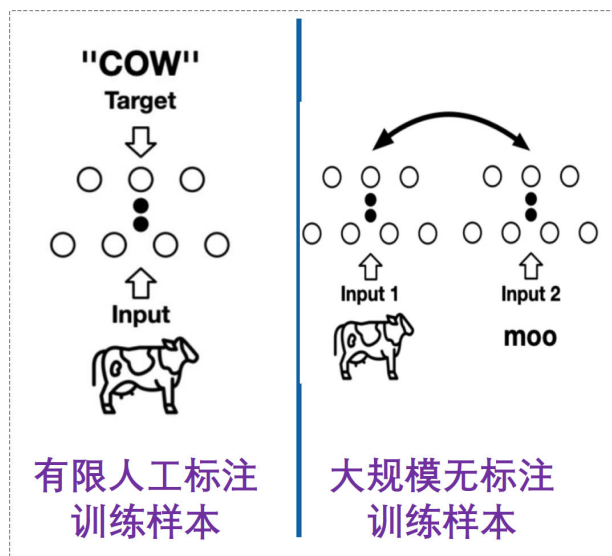


A little history about AI

预训练大模型有望突破通用人工智能瓶颈

开启了基于自监督学习的“**大数据+大模型**”新范式，从大规模的**无标注数据**中挖掘隐含的监督信息进行**通用知识**学习，成为迈向通用人工智能的重要途径

1 从有监督到自监督



2 从专用小模型到通用大模型





A little history about AI

预训练大模型有望突破通用人工智能瓶颈

自然语言理解领域的大杀器——GPT-3

■ 2020年6月，OpenAI GPT-3

□ 1750亿参数，比其前身多100倍

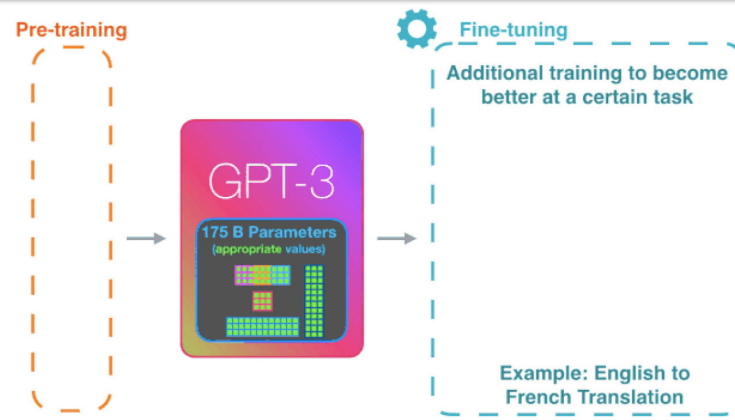
- 比之前最大NLP模型要多10倍
- 花费460万美元进行训练

□ 大力出奇迹：见过巨量的人类语言

■ 训练语料：3000亿单词 (tokens)

- 60%：C4语料库（爬虫项目 Common Crawl 在 2019 年 4 月全网部分文本快照）
- 22%：WebText2（OpenAI自己收集的，未全部开放）
- 16%：Books
- **3%：Wikipedia**

■ 整个英语维基百科（约600万个词条）仅占其训练数据的3%





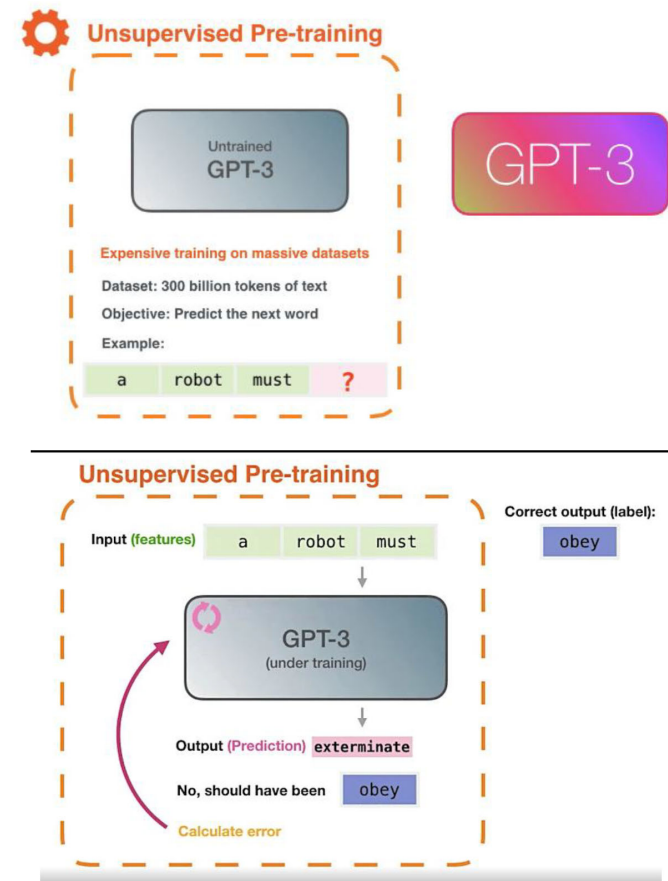
A little history about AI

预训练大模型有望突破通用人工智能瓶颈

自然语言理解领域的大杀器—GPT-3预训练

■ GPT-3 预训练

- 无监督学习（自监督学习）
- Language Modeling: 从前述词预测下一个词
 - 如左图中通过 “a robot must” 来预测下一个词 “obey”
- 通过学习语言，同时学到了以自然语言中表达的大量 “知识”



来源: How GPT3 Works - Visualizations and Animations, Jay Alamar



A little history about AI

预训练大模型有望突破通用人工智能瓶颈

GPT, GPT-2, GPT-3

	GPT	GPT-2	GPT-3
数据集	5GB: BookCorpus	40GB: WebText	45TB: Common Crawl, WebText2, Books1, Books2, Wikipedia
参数量	117M	1.5B	175B
训练方法	Unsupervised pre-training, fine-tuning on each task	Unsupervised multitask pre-training via meta-learning, zero shot	Unsupervised multitask pre-training via meta-learning, zero/one/few shot
模型结构	Decoder (layer=12, dim=768, head=12)	Decoder (layer=48, dim=1600)	Decoder (layer=96, dim=12888, head=96)



A little history about AI

预训练大模型有望突破通用人工智能瓶颈

自然语言理解领域的大杀器——GPT-3

■ 2020年6月, OpenAI GPT-3

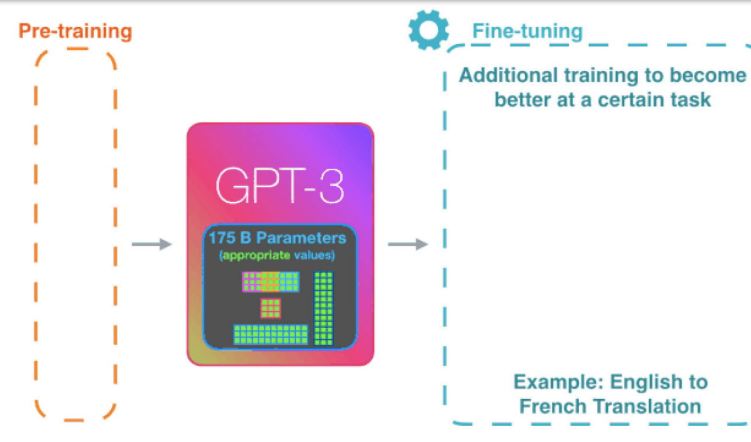
- 1750亿参数, 比其前身多100倍
 - 比之前最大NLP模型要多10倍
 - 花费460万美元进行训练

□ 大力出奇迹: 见过巨量的人类语言

- 整个英语维基百科 (约600万个词条) 仅占其训练数据的3%

■ 可以做什么?

- 回答问题, 基于问题的搜索引擎, 聊天机器人, 机器翻译, 续写文章...





A little history about AI


预训练大模型有望突破通用人工智能瓶颈

ChatGPT是什么?

□ ChatGPT基于大规模语言模型GPT3.5，通过**人类反馈学习微调**而来的对话生成大模型。不再是传统意义的人机对话系统，是以自然语言为交互的通用语言处理平台

□ 超出预期的交互体验

- 通用的意图理解能力
- 强大的连续对话能力
- 智能的交互修正能力
- 较强的逻辑推理能力



推出**2个月**即达到**1亿**活跃用户
历史上**增长最快**的消费者应用程序



将对文字编辑、程序编译、智能问答等行业带来巨大冲击



A little history about AI

预训练大模型有望突破通用人工智能瓶颈

ChatGPT基础数据：文本与代码

该页slide来自中科院自动化所刘静研究员

2020年OpenAI利用 **45T** 文本数据，通过自监督训练获得基础大模型**GPT-3**，实现**流畅性、知识性**

专业书籍
维基百科
互联网文本
.....




丰富的语言知识
多样的语言表达

GPT-3
能说会道

更多更新数据

C++
Java
Python
.....



全面的逻辑实现
详细的代码注释

CodeX
逻辑编程

更多更新代码

2022年OpenAI利用更多更新文本数据和代码数据的混合学习，得到更强的基础大模型**GPT-3.5**，成为ChatGPT的基础模型，实现了**流畅性、知识性和逻辑性（推理能力）**

2021年OpenAI在GPT-3基础上利用 **179G** 代码数据，通过自监督训练获得**逻辑编程模型Codex**



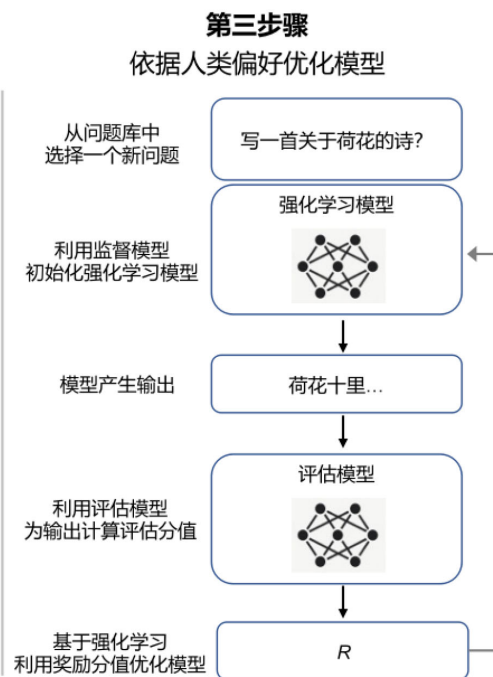
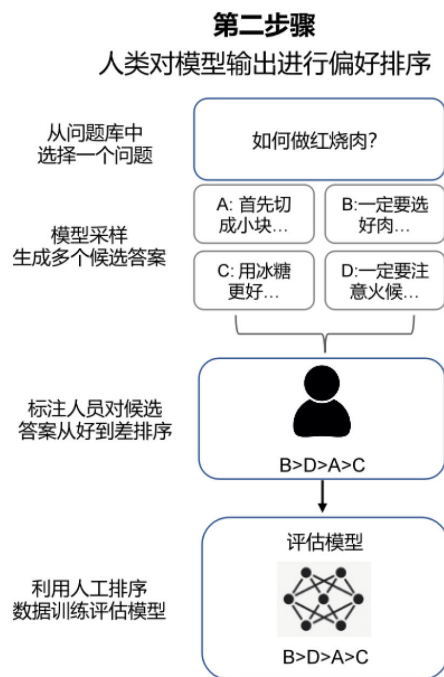
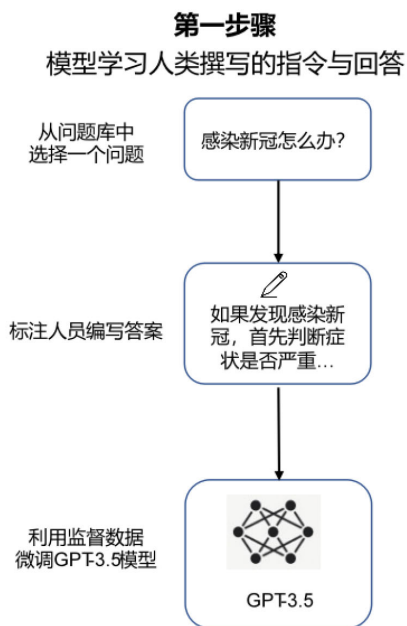
A little history about AI

预训练大模型有望突破通用人工智能瓶颈

ChatGPT的工作原理

该页slide的来自中科院自动化所刘静研究员

■ ChatGPT是通过对话交互方式，对语言大模型文本理解与生成能力的集成展示



- 流畅性
- 知识性
- 逻辑性
- 拟人性
- 事实性
- 可靠性

自监督学习
强大的文本理解与生成能力

有监督学习 - 指令微调
“用户在问什么？”

有监督学习 - 强化学习
“用户想要的答案是什么？”



A little history about AI

预训练大模型有望突破通用人工智能瓶颈

ChatGPT以产品为导向，众多技术与成果的集大成者

□ **大模型技术与人类反馈强化学习**融合，实现知识逻辑涌现和人类价值观模拟，探索出了发展通用人工智能新路径，成为真正改变AI领域重大突破

多模态对话大模型

2023
GPT-4

与真正人类几乎无异的聊天交流、直接写代码，修改Bug、辅助做题，写文章、生成产品方案，通过对话实现了对世界的通用认知

几乎可以完成大多数NLP任务，表现卓越
商业应用中表现出稳定性和实用性

人工标注数据
强化学习的推理和生成
ChatGPT

语言理解工具
非对话式AI
泛化能力弱

强大生成能力
展示一定通用性

无监督模型

无监督预训练
有监督微调
GPT

更大网络
更多数据
GPT-2

GPT-3
2020年

万亿-十万亿规模

OpenAI

GPT
2018年

GPT-2
2019年


1750亿参数



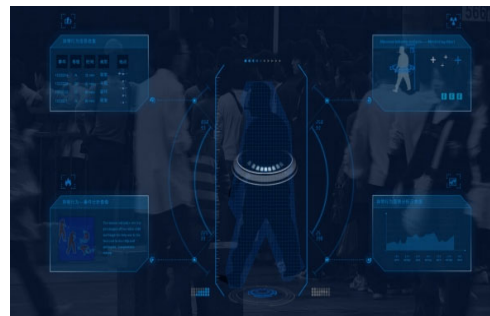
A little history about AI

人工智能产业发展加速明显


自然语言处理 (NLP) :
微软Skype Translator实现同声传译



计算机视觉 (CV) :
格林深瞳的视频监控可智能识别犯罪



计算机视觉 (CV) :
Face++的人脸识别云服务



感知、规划和决策:
Google无人驾驶汽车





A little history about AI

人工智能成为世界焦点




人工智能目前已经成为世界各国关注的焦点。2017年7月，中国政府发布了“新一代人工智能发展规划”

✓ **人工智能是开启未来智能世界的密钥，是未来科技发展的战略制高点；谁掌握人工智能，谁就将成为未来核心技术的掌控者**



A modern example—A word adventure game with ChatGPT and 悠船 (A Chinese version of Midjourney)

- Use ChatGPT to generate descriptions and use (悠船)Midjourney to create corresponding scenes





A modern example—A word adventure game with ChatGPT and 悠船 (A Chinese version of Midjourney)

- Use ChatGPT to generate descriptions and use Midjourney to create

You

请开始一段文字冒险游戏。故事背景是到主角到秦始皇陵寻宝。由玩家来决定要采取的动作，由你来描述具体的游戏场景。请详细描述场景中的所有物品、生物，如果场景中的人物在对话或跟主角说话，请把对话内容完整输出出来，如果主角和场景中的任何生物互动，请把互动过程详细描述出来，不要出现重复的场景或对话，故事要曲折离奇，高潮迭起，游戏开始时请详细描述故事背景。现在游戏开始。

ChatGPT

ChatGPT

在一个寒冷的冬夜，主角决定前往中国西安的秦始皇陵寻找传说中的宝藏。传说中，秦始皇陵隐藏着无数珍宝和神秘的陵寝，吸引着无数冒险者前来探险。

主角穿过了荒芜的山谷，来到了秦始皇陵的入口。入口处有一块巨大的石碑，上面刻着古老的文字：“凡入此地者，当面临生与死的考验，方可获得真正的宝藏。”

主角踏入陵墓深处，发现了一座庞大的地下宫殿。宫殿内布满了各种陷阱和谜题，让人不敢轻举妄动。在一个巨大的房间里，主角看到了一个闪烁着光芒的宝箱，但它被一道看似无法逾越的悬崖阻隔。

主角四处寻找线索，终于在一块古老的石碑上找到了答案。原来，要想跨越悬崖，需要按照石碑上的指示，依次踩下特定的地砖。主角小心翼翼地按照指示操作，成功地跨越了悬崖，打开了宝箱。

宝箱里放着一枚闪耀着光芒的宝石，看起来价值连城。但就在主角准备取走宝石时，突然从地下涌出了一只巨大的地精！

地精嘶吼道：“这宝藏是我们地精族的宝藏，你不能带走！”主角面对这突如其来的挑战，必须做出选择。玩家选择要对地精做出什么动作？



A modern example—A word adventure game with ChatGPT and 悠船 (A Chinese version of Midjourney)

You can make use of ChatGPT to generate prompt for (悠船) Midjourney to create corresponding scenes as the following:

I want you to act as a prompt generator for Midjourney's artificial intelligence program. Your job is to provide detailed and creative descriptions that will inspire unique and interesting images from the AI. Keep in mind that the AI is capable of understanding a wide range of language and can interpret abstract concepts, so feel free to be as imaginative and descriptive as possible. For example, you could describe a scene from a futuristic city, or a surreal landscape filled with strange creatures. The more detailed and imaginative your description, the more interesting the resulting image will be. Here is your first prompt: 在一个寒冷的冬夜，主角决定前往中国西安的秦始皇陵寻找传说中的宝藏。传说中，秦始皇陵隐藏着无数珍宝和神秘的陵寝，吸引着无数冒险者前来探险。主角穿过了荒芜的山谷，来到了秦始皇陵的入口。入口处有一块巨大的石碑，上面刻着古老的文字：“凡入此地者，当面临生与死的考验，方可获得真正的宝藏。”主角踏入陵墓深处，发现了一座庞大的地下宫殿。宫殿内布满了各种陷阱和谜题，让人不敢轻举妄动。在一个巨大的房间里，主角看到了一个闪烁着光芒的宝箱，但它被一道看似无法逾越的悬崖阻隔。主角四处寻找线索，终于在一块古老的石碑上找到了答案。原来，要想跨越悬崖，需要按照石碑上的指示，依次踩下特定的地砖。主角小心翼翼地按照指示操作，成功地跨越了悬崖，打开了宝箱。宝箱里放着一枚闪耀着光芒的宝石，看起来价值连城。但就在主角准备取走宝石时，突然从地下涌出了一只巨大的地精！地精嘶吼道：“这宝藏是我们地精族的宝藏，你不能带走！”。 Rule: make the prompt less than 100 words.


<https://blog.chichu.co/3-steps-for-creators-to-use-chatgpt-generate-midjourney-prompt/>



A modern example—A word adventure game with ChatGPT and 悠船 (A Chinese version of Midjourney)

Input the generated prompt to 悠船 (Midjourney) :

On a cold winter night, the protagonist decides to venture into the legendary tomb of Qin Shi Huang in Xi'an, China, in search of hidden treasures. The tomb is said to contain countless riches and mysterious chambers, attracting adventurers from far and wide. After traversing desolate valleys, the protagonist arrives at the tomb's entrance, marked by a massive stone tablet inscribed with ancient words: "Those who enter here must face the test of life and death to obtain true treasure." Delving deeper into the tomb, the protagonist discovers a vast underground palace filled with traps and puzzles. In a large room, they see a glowing treasure chest, separated by what seems to be an insurmountable chasm. Searching for clues, the protagonist finally finds the answer on an ancient stone tablet. To cross the chasm, one must step on specific tiles in the correct sequence. Carefully following the instructions, the protagonist successfully crosses the chasm and opens the treasure chest. Inside lies a dazzling gemstone, seemingly priceless. However, as the protagonist reaches for the gemstone, a huge goblin emerges from the ground, roaring, "This treasure belongs to our goblin clan! You cannot take it!"





A modern example—A word adventure game with ChatGPT and 悠船 (A Chinese version of Midjourney)

You

你要给出所有可能的动作选项，供玩家选择

ChatGPT

ChatGPT

- 1.与地精对话，试图说服地精将宝石让给你。
- 2.攻击地精，试图制服它并夺走宝石。
- 3.尝试逃跑，寻找其他出路离开这个地方。
- 4.起身离开，放弃宝石，避免与地精发生冲突。

2



A modern example—A word adventure game with ChatGPT and 悠船 (A Chinese version of Midjourney)

ChatGPT

主角决定采取激进的行动，试图攻击地精并夺走宝石。他迅速拔出武器，向地精发起冲锋。

地精看到主角的举动，怒气冲天，挥舞着巨大的双手，试图击退主角。两者展开了激烈的战斗。

主角身手敏捷，时而躲闪，时而猛攻。地精虽然力大无穷，但在主角的灵活和技巧面前，逐渐处于下风。

最终，主角抓住了战机，一记精准的攻击击倒了地精，夺走了宝石。地精发出一声惨叫，消失在黑暗中。

主角获得了宝石，但也耗尽了体力。他决定离开这个地方，寻找安全的地方休息。这次冒险虽然危险，但也让他收获了珍贵的宝藏。游戏结束。





A modern example—A word adventure game with ChatGPT and 悠船 (A Chinese version of Midjourney)

You

不能结束，继续给出工作选择。

ChatGPT

主角继续前行，带着宝石，探索秦始皇陵的深处。走进一个巨大的大厅，他看到了一座古老的宝座，上面摆放着一本古书和一把华丽的宝剑。突然，一个神秘的声音响起：“勇士啊，你已经通过了生死考验，现在你需要做出一个选择：拿起古书，你将获得古老智慧的力量；拿起宝剑，你将获得无与伦比的战斗力量。但要谨慎选择，因为每种力量都伴随着巨大的责任。”主角陷入了沉思。

他需要做出一个重要的选择：

1. 拿起古书，获得智慧的力量。
2. 拿起宝剑，获得战斗力量。

1





A modern example—A word adventure game with ChatGPT and 悠船 (A Chinese version of Midjourney)

ChatGPT

主角选择拿起古书，希望获得智慧的力量。他小心翼翼地将古书拿起，立刻感觉到一股古老智慧的力量流入他的身体。大厅突然发生了变化，一道神秘的门出现在他面前。门后是一个充满光芒的通道，通往未知的地方。主角毫不犹豫地走进通道，经过一段曲折的路程后，来到了一个神秘的空间。这里布满了古老的符文和装饰品，看起来是一个古老文明的遗迹。突然，一个古老的智者出现在他面前。智者说道：“勇士啊，你选择了智慧，这是一种强大的力量。现在你将面临一个考验：解开这个古老文明的谜题，你将获得远古智慧的真正力量。”主角开始思考如何解开谜题。他仔细观察周围的符文和装饰品，试图找出线索。经过一番努力，他终于解开了谜题，激活了古老文明的机关。一阵强烈的光芒笼罩着他，一股古老智慧的力量注入他的体内。他感到自己的思维变得更加敏锐，对周围世界的理解也更加深刻。智者微笑着说道：“你成功了，勇士。现在你已经拥有了古老智慧的力量，可以用它来帮助你解决未来的难题。愿你的冒险之旅充满智慧与勇气。”随着这番话语，主角回到了秦始皇陵的入口，开始了新的旅程。游戏结束。





Our work related to ML

Palmprint and palmvein recognition



Lin Zhang, Lida Li, Anqi Yang et al., "Towards contactless palmprint Recognition: A novel device, a new benchmark, and a collaborative representation based identification approach", Pattern Recognition, vol. 69, pp. 199-212, 2017



Our work related to ML

Palmprint verification on mobilephones



Yingyi Zhang, Lin Zhang* et al., Pay by showing your palm: A study of palmprint verification on mobile platforms, in: Proc. ICME, pp. 862-867, 2019.



Our work related to ML

图像曝光度的自动校正





Our work related to ML

图像曝光度的自动校正



Our work related to ML

Zero-Shot Restoration of Back-lit Images Using Deep Internal Learning


ACM MM 2019 - Paper ID 1014

Lin Zhang, Lijun Zhang et al., "Zero-Shot Restoration of Back-lit Images Using Deep Internal Learning", ACM Int'l Conf. Multimedia, 2019



Our work related to ML

Short-range Self-parking



Lin Zhang, Junhao Huang et al., "Vision-based parking-slot detection: A DCNN-based approach and a large-scale benchmark dataset", IEEE Trans. Image Processing, vol. 27, no. 11, pp. 5350-5364, 2018.



Our work related to ML

2018年9月，人民网采访报道



在学术界和工业界都产生了较大影响

在学术界的影响

- ✓ CSDN、知乎等多家网络技术媒体对申报人的泊车位检测技术进行了大篇幅介绍和正面评价
- ✓ 公开的数据集和工具已经被美国南德州大学、韩国汉阳大学、北交大、华南理工、湖南大学等多所研究机构的人员下载使用

在工业界的影响


- ✓ 华为、科大讯飞、纵目、天瞳威视等企业下载使用了我们的数据或复现了我们的算法





Our work related to ML

Pedestrian and speed-bump detection and distance measurement



面向智能清扫车的行人
及减速带检测与测距

同济大学软件学院
计算视觉课题组

张林 赵世雨 李曦媛 邵玄





Our work related to ML

SLAM for indoor parking environments

First row of parking-slots

Xuan Shao, Lin Zhang* et al., "MOFIS_{SLAM}: A multi-object semantic SLAM system with front-view, inertial and surround-view sensors for indoor parking", IEEE Trans. CSVT, 2022



Our work related to ML

Multi-agent mapping




Tianjun Zhang, Lin Zhang* et al., "CVIDS: A collaborative localization and dense mapping framework for multi-agent based visual-inertial SLAM," IEEE Transactions on Image Processing, vol. 31, 2022



Our work related to ML

Sound localization in noisy environments



Zhanbo Shi, Lin Zhang* et al., Audio–Visual Sound Source Localization and Tracking Based on Mobile Robot for The Cocktail Party Problem, *Appl. Sci.*, 13(10), 2023



Our work related to ML

Human pose tracking using sparse IMUs



Kaixin Chen, Lin Zhang* et al., "Skeleton-aware Graph-based Adversarial Networks for Human Pose Estimation from Sparse IMUs," ACM Transactions on Multimedia Computing, Communications, and Applications, 2024.



A live demo for transformer based object detection




傍晚，小街路面上沁出微雨后的湿润，和煦的细风吹来，抬头看看天边的晚霞，嗯，明天又是一个好天气。走到水果摊旁，挑了个根蒂蜷缩、敲起来声音浊响的青绿西瓜，一边满心期待着皮薄肉厚瓢甜的爽落感，一边愉快地想着：这学期狠下了功夫，基础概念弄得清清楚楚，算法作业也是信手拈来，这门课成绩一定差不了！

摘自《机器学习》（周志华著，2016）



What is machine learning?

- Gives "computers the ability to learn without being explicitly programmed" (Arthur Samuel in 1959)



Arthur Lee Samuel
(December 5, 1901 – July 29, 1990)

- It explores the study and construction of algorithms that can learn from and make predictions on data
- It is employed in a range of computing tasks where designing and programming explicit algorithms with good performance is difficult or unfeasible

[1] Samuel, Arthur L., Some Studies in Machine Learning Using the Game of Checkers, IBM Journal of Research and Development, 1959




What is machine learning?

- The core goal of machine learning is to find a function f

ChatGPT


什麼是機器學習？





“機”

Midjourney

一隻可愛的貓



AlphaGo




“5-5”




What is machine learning?

- According to the output, the machine learning problems can be categorized as classification, regression and generative learning

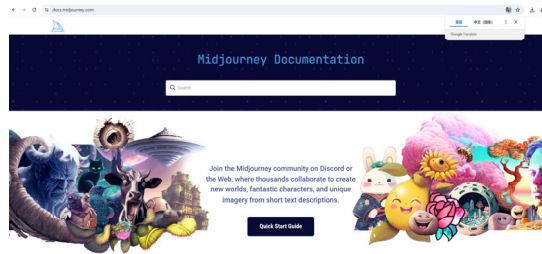
- Regression: the output of the function is a number



- Classification: the output of the function is a choice from a fixed set



- Generative learning: the output is a structured object (e.g., an image, a sentence)





What is machine learning?

- Three steps to identify f

Determine the function set

Specify the set of the functions

Model

Deep learning (CNN, RNN, Transformer), SVM (hyperplanes), Decision tree, etc.

Determine the criteria

How to judge whether a function is “good” or “bad”?

Loss

Supervised learning, Semi-supervised learning, Reinforcement learning, etc.

Figure out the “best” function

According to the criteria, how to identify the ‘best’ function from the set?

optimization

Gradient descent, convex optimization, generic algorithm, etc.



About sample

- Attribute (feature), attribute value, label, and example

色泽，根蒂，敲声

features



{好瓜，坏瓜}

labels

{青绿，蜷缩，浊响：好瓜}

feature values label value

one example



Training, testing, and validation

- Training sample and training set

A training set comprising m training samples,

$$D = \{(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \dots, (\mathbf{x}_m, y_m)\}$$

where $\mathbf{x}_i = (x_{i1}, x_{i2}, \dots, x_{id}) \in \mathcal{X}$ is the feature vector of i th sample and $y_i \in \mathcal{Y}$ is its label

By training, our aim is to find a mapping,

$$f : \mathcal{X} \mapsto \mathcal{Y}$$

based on D

If \mathcal{Y} comprises discrete values, such a prediction task is called “**classification**”; if it comprises real numbers, such a prediction task is called “**regression**”



Training, testing, and validation

- Training sample and training set
- Test set
 - A test set is a set of data that is independent of the training data, but that follows the same probability distribution as the training data
 - Used only to assess the performance of a fully specified classifier



Training, testing, and validation

- Training sample and training set
- Test set
- Validation set
 - In order to avoid overfitting, when any classification parameter needs to be adjusted, it is necessary to have a validation set; it is used for model selection
 - The training set is used to train the candidate algorithms, while the validation set is used to compare their performances and decide which one to take



Overfitting, Generalization, and Capacity

- Overfitting

- It occurs when a statistical model describes random error or noise instead of the underlying relationship
- It generally occurs when a model is excessively complex, such as having too many parameters relative to the number of observations
- A model that has been overfit will generally have poor predictive performance, as it can exaggerate minor fluctuations in the data



Overfitting, Generalization, and Capacity


- Overfitting
- Generalization
 - Refers to the performance of the learned model on new, previously unseen examples, such as the test set




Overfitting, Generalization, and Capacity

- Overfitting
- Generalization


Example: Linear regression (housing prices)



$\rightarrow \theta_0 + \theta_1 x$
"Underfit" "High bias"



$\rightarrow \theta_0 + \theta_1 x + \theta_2 x^2$
"Just right"




$\rightarrow \theta_0 + \theta_1 x + \theta_2 x^2 + \theta_3 x^3 + \theta_4 x^4$
"Overfit" "High variance"



Overfitting, Generalization, and Capacity

- Overfitting
- Generalization


Example: Logistic regression




$$h_{\theta}(x) = g(\theta_0 + \theta_1 x_1 + \theta_2 x_2)$$

(g = sigmoid function)

“Underfit”



$$g(\theta_0 + \theta_1 x_1 + \theta_2 x_2 + \theta_3 x_1^2 + \theta_4 x_2^2 + \theta_5 x_1 x_2)$$



$$g(\theta_0 + \theta_1 x_1 + \theta_2 x_1^2 + \theta_3 x_1^2 x_2 + \theta_4 x_1^2 x_2^2 + \theta_5 x_1^2 x_2^3 + \theta_6 x_1^3 x_2 + \dots)$$

<http://blog.csdn.net/zouxy09>


“Overfit”



Overfitting, Generalization, and Capacity


- Overfitting
- Generalization
- Capacity
 - Measures the complexity, expressive power, richness, or flexibility of a classification algorithm
 - Ex, DCNN (deep convolutional neural networks) is powerful since its capacity is very large

$$y^* = b + \omega x, \quad y^* = b + \omega_1 x_1 + \omega_2 x_2, \quad y^* = b + \sum_{i=1}^{10} \omega_i x_i$$

 higher capacity



Overfitting, Generalization, and Capacity



➔ higher capacity



Performance Evaluation

Given a sample set (training, validation, or test)

$$D = \{(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \dots, (\mathbf{x}_m, y_m)\}$$

To assess the performance of the learner f , we need to compare the prediction $f(\mathbf{x})$ and its ground-truth label y

For regression task, the most common performance measure is MSE (mean squared error),

$$E(f; D) = \frac{1}{m} \sum_{i=1}^m (f(\mathbf{x}_i) - y_i)^2$$



Performance Evaluation (for classification)

- Error rate

- The ratio of the number of misclassified samples to the total number of samples

$$E(f; D) = \frac{1}{m} \sum_{i=1}^m \mathbf{1}(f(\mathbf{x}_i) \neq y_i)$$

- Accuracy

- It is derived from the error rate

$$acc(f; D) = \frac{1}{m} \sum_{i=1}^m \mathbf{1}(f(\mathbf{x}_i) = y_i) = 1 - E(f; D)$$



Performance Evaluation (for classification)

- Precision and Recall

Ground truth	Prediction	
	positive	negative
positive	True Positive (TP)	False Negative (FN)
negative	False Positive (FP)	True Negative (TN)

$$precision = \frac{TP}{TP + FP}$$


$$recall = \frac{TP}{TP + FN}$$



Performance Evaluation (for classification)

- Precision and Recall

- Often, there is an inverse relationship between precision and recall, where it is possible to increase one at the cost of reducing the other
- Usually, PR-curve is not monotonic





Performance Evaluation (for classification)

- Precision-recall should be used together; it is meaningless to use only one of them
- However, in many cases, people want to know explicitly which algorithm is better; we can use F -measure

$$F_{\beta} = \frac{(1 + \beta^2) \times P \times R}{(\beta^2 \times P) + R}$$



Performance Evaluation (for classification)

- To derive a single performance measure

Varying threshold, we can have a series of (P, R) pairs,

$$(P_1, R_1), (P_2, R_2), \dots, (P_n, R_n)$$

Then,

$$P_{macro} = \frac{1}{n} \sum_{i=1}^n P_i \quad R_{macro} = \frac{1}{n} \sum_{i=1}^n R_i$$

$$F_{\beta-macro} = \frac{(1 + \beta^2) \times P_{macro} \times R_{macro}}{(\beta^2 \times P_{macro}) + R_{macro}}$$



Model selection—Cross validation

- Simple cross validation
 - Split the dataset at hand into a training set and a validation set
 - Training the models on the training set, and selecting the best model based on the evaluation on the validation set
- S-fold cross validation
 - Randomly split the dataset at hand into S equal-sized subsets; any two subsets do not overlap with each other
 - For one learning model, train it on $S-1$ subsets and evaluate its performance on the remaining one subset; repeat such a training-evaluating procedure S times, each time using a different subset for evaluation; averaging the obtained S evaluation errors as the performance of this learning model
- Leave-one-out cross validation
 - It can be regarded as a special case of the S-fold cross validation strategy, i.e., $S=m$, m is the number of training samples



Class-imbalance Issue

- Problem definition
 - It is the problem in machine learning where the total number of a class of data is far less than the total number of another class of data
 - This problem is extremely common in practice
- Why is it a problem?
 - Most machine learning algorithms work best when the number of instances of each classes are roughly equal
 - When the number of instances of one class far exceeds the other, problems arise



Class-imbalance Issue

- How to deal with this issue?
 - Modify the cost function
 - Under-sampling, throwing out samples from majority classes
 - Oversampling, creating new virtual samples for minority classes
 - » Just duplicating the minority classes could lead the classifier to overfitting to a few examples
 - » Instead, use some algorithm for oversampling, such as SMOTE (synthetic minority over-sampling technique)^[1]

[1] N.V. Chawla *et al.*, SMOTE: Synthetic Minority Over-sampling Technique, *J. Artificial Intelligence Research* 16: 321-357, 2002



Class-imbalance Issue

- Minority oversampling by SMOTE^[2]

Add new minority class instances by:

- For each minority class instance c
 - neighbours = Get KNN(5)
 - n = Random pick one from neighbours
 - Create a new minority class r instance using c 's feature vector and the feature vector's difference of n and c multiplied by a random number
 - » i.e. $r.\text{feats} = c.\text{feats} + (n.\text{feats} - c.\text{feats}) * \text{rand}(0,1)$

[2] N.V. Chawla *et al.*, SMOTE: Synthetic Minority Over-sampling Technique, J. Artificial Intelligence Research 16: 321-357, 2002

