

Topic 11: 让机器自己学习

Machine Learning by Itself

刘跃文 博士 Dr. LIU, Yuewen

教授、博士生导师 Professor

liuyuewen@xjtu.edu.cn

西安交通大学管理学院

School of Management, Xi'an Jiaotong University

V2.0, 2023-Oct

有监督学习的问题

Weaknesses of Supervised Learning

- 需要大量的标注数据；而现实生活中很多标注数据成本很高。
Need a lot of annotation data; In real life, a lot of annotated data costs a lot.
- 标注数据的问题限制了机器学习的发展。
The problem of labeling data limits the development of machine learning.



中文	英文
①[全国各族人民]在以习近平总书记为核心的党中央坚强领导下，	①However, [we the Chinese people] under the strong leadership of the Central Committee of the Communist Party of China with Comrade Xi Jinping at its core,
②全国各族人民迎难而上，	②we the Chinese people have risen to the challenge
③[全国各族人民]砥砺前行，	③And [we the Chinese people] worked hard to press ahead,
④[全国各族人民]推动经济社会持续健康发展。	④[we the Chinese people] driving forward sustained, healthy economic and social development.

提纲 Outline

1. 强化学习 Reinforcement learning
2. 无中生有 GAN 和 Diffusion Model GAN and Diffusion Model come out of nothing
3. 自监督学习与预训练 Self-supervised learning and pre-training
4. 人工智能的弱点 The weakness of artificial intelligence
5. 人工智能的应用 The application of artificial intelligence
6. 人工智能的发展方向和未来 The development direction and future of artificial intelligence

1. 强化学习 Reinforcement learning

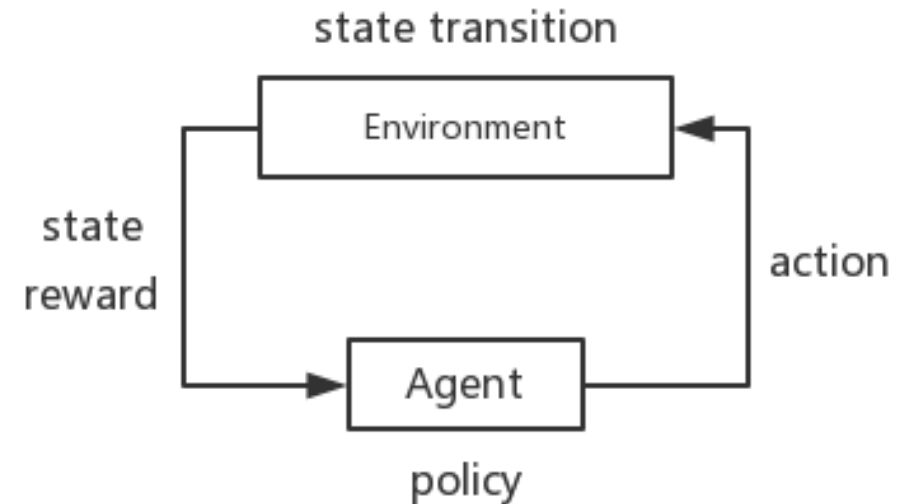


1. What is reinforcement/reinforcement Learning

- reinforcement learning (RL), also known as reinforcement learning, is one of the main methods in the field of machine learning and intelligent control in recent years. Reinforcement learning is learning what to do ---- how to map situations to actions ---- so as to maximize a numerical reward signal.
- In other words, reinforcement learning is concerned with how an agent can adopt a set of behaviors in the environment in order to obtain the greatest cumulative reward.
- Through reinforcement learning, an agent should know what behavior to take in what state. RL is learning from the mapping of environmental states to actions, which we call policy.

https://blog.csdn.net/zz_1215/article/details/44138715

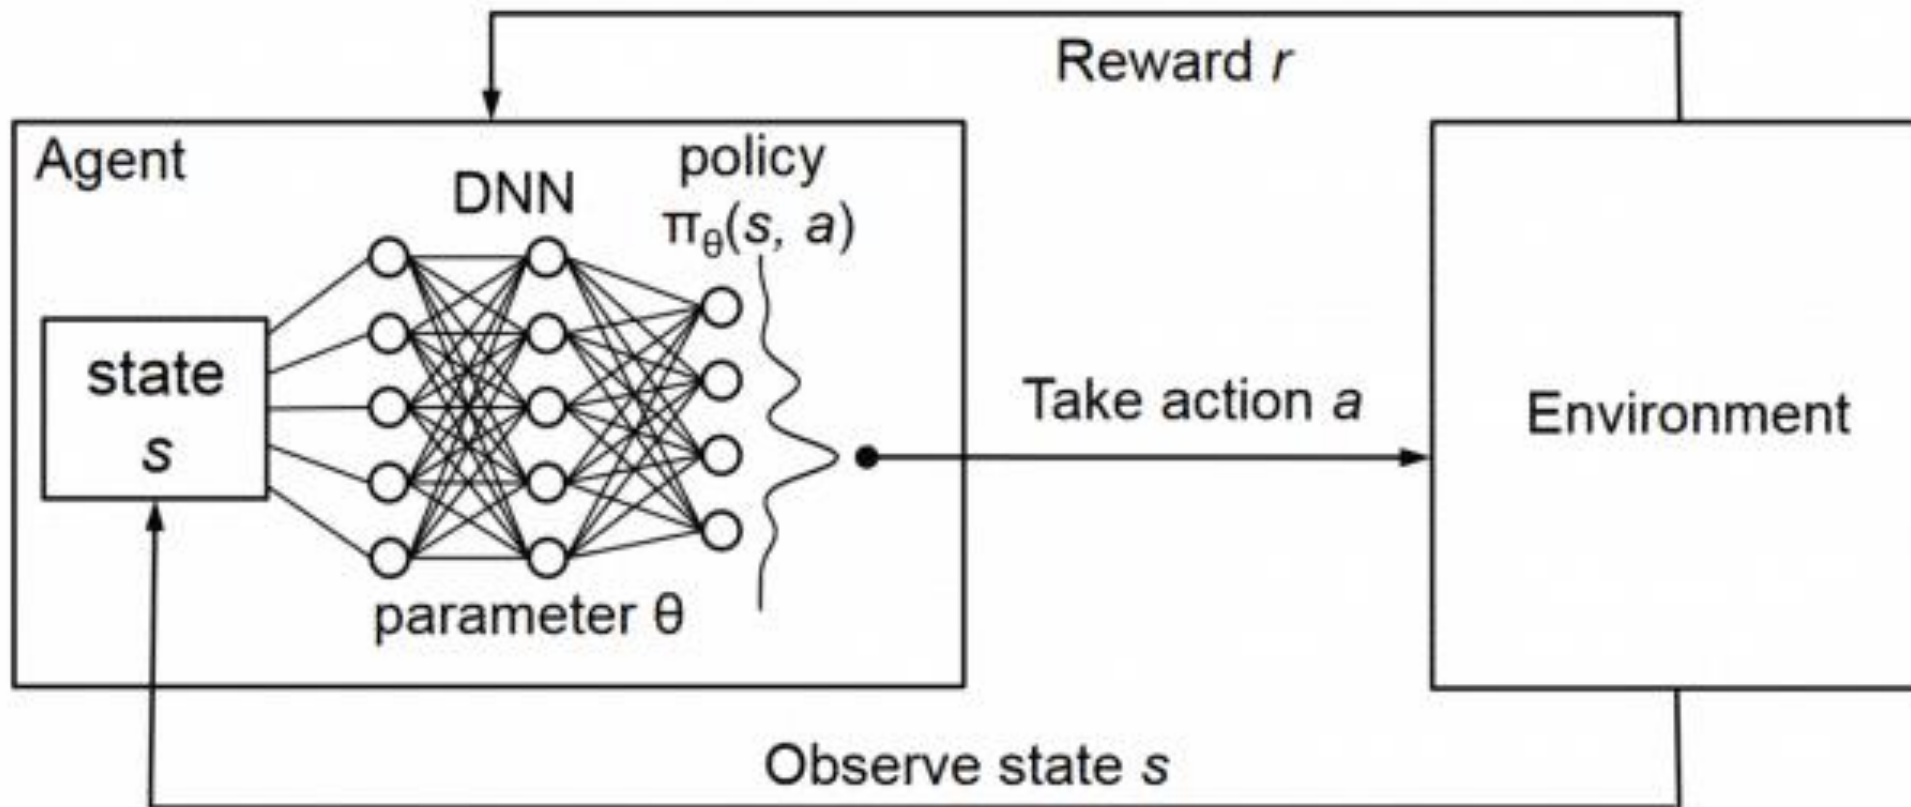
-
- In classical reinforcement learning, the player, as an Agent, has a series of interactions with the Environment. At each moment, the environment will give a current state, the player will make a decision according to the state, the decision will affect the environment, so that the environment has a certain change, the changed environment will give the player a "Reward", the reward can be positive or negative, it is used to feedback the user's current performance. The environment also gives the player feedback on the state of the next moment, so that the player can make a new round of decisions.



-
- Reinforcement learning is trial-and-error learning. Since there is no direct guidance information, the agent must constantly interact with the environment to obtain the best strategy through trial and error.
 - Delayed reward, reinforcement learning has little guidance and is often given after the fact (the last state), which leads to the problem of how to allocate the reward to the previous state after the positive or negative reward has been obtained.
 - Reinforcement learning is a very active and interesting area of machine learning, and it is closer to the nature of biological learning than other learning methods, and therefore promises to achieve higher intelligence, which has been shown in board games.
 - Tesauro(1995) describes the TD-Gammon program, which uses reinforcement learning to become a world-class backgammon player. After 1.5 million self-generated games training, the program has nearly reached the level of the best human players, and in the competition with the top human masters achieved 40 sets only lost 1 game good results.

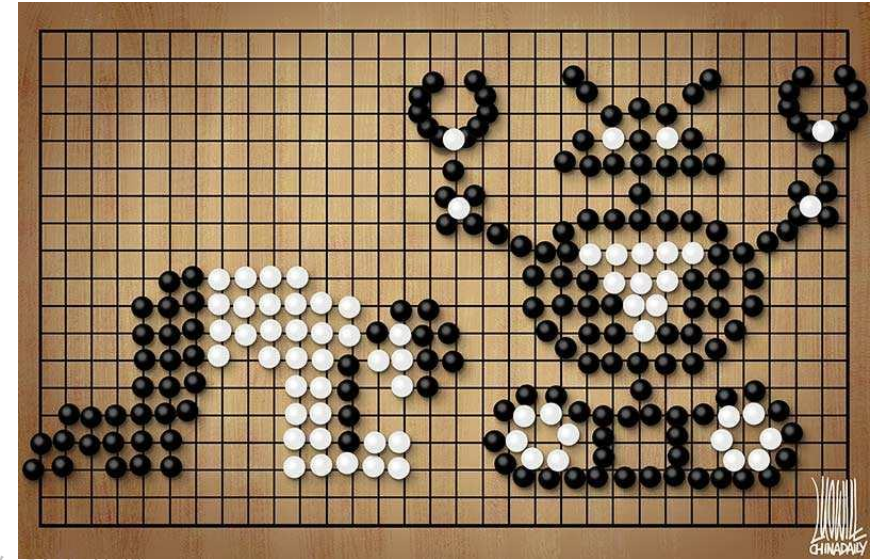
2. 深度强化学习

Deep Reinforcement Learning



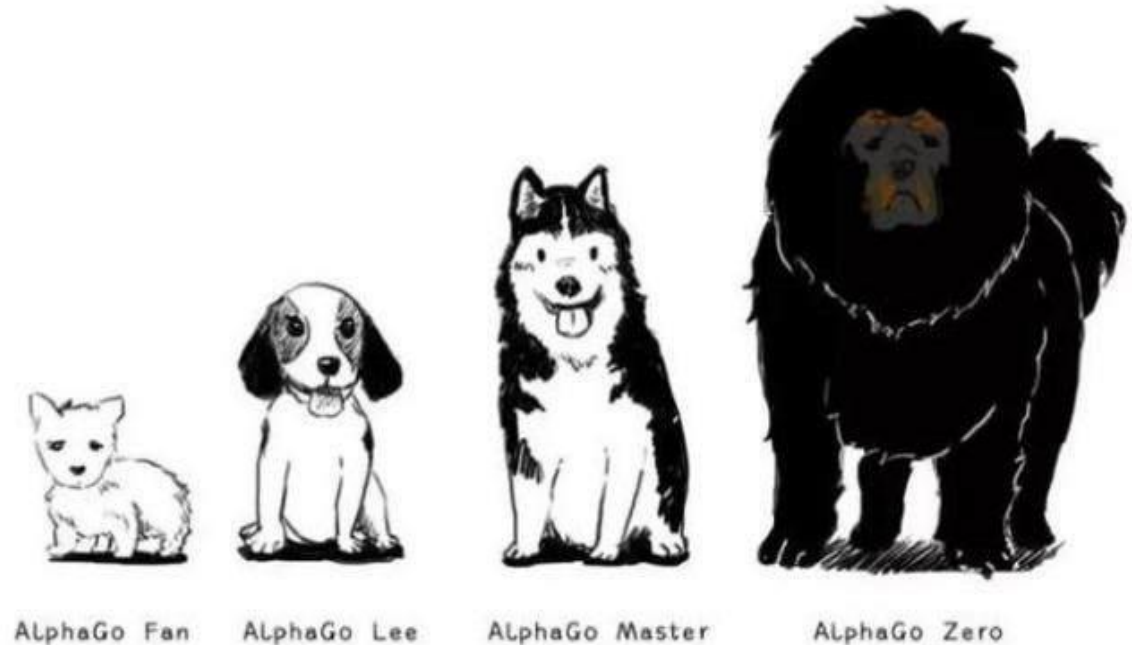
Deep reinforcement learning and AlphaGo

- In January 2017, Hassabis, CEO of Google DeepMind, announced the launch of the real 2.0 version of AlphaGo at the DLD (Digital, Life, Design) Innovation conference in Munich, Germany. Its characteristic is to abandon human chess, only rely on deep learning to grow up to challenge the limits of Go.



AlphaGo Zero

- The new version of AlphaGo starts completely from 0, does not use any previous human experience and data, and learns by itself, and fights left and right to complete the upgrade. The visual analogy is about an idiot who taught himself to go for three days and then beat a master who had studied all kinds of chess and absorbed almost all human experience.
- "What really amazes me is that with just **4.9 million self-matches**, and just 1,600 MCTS per move rollout, Zero has surpassed last March's level."





飞机大战飞机大战这个游戏是微信出的第一款游戏

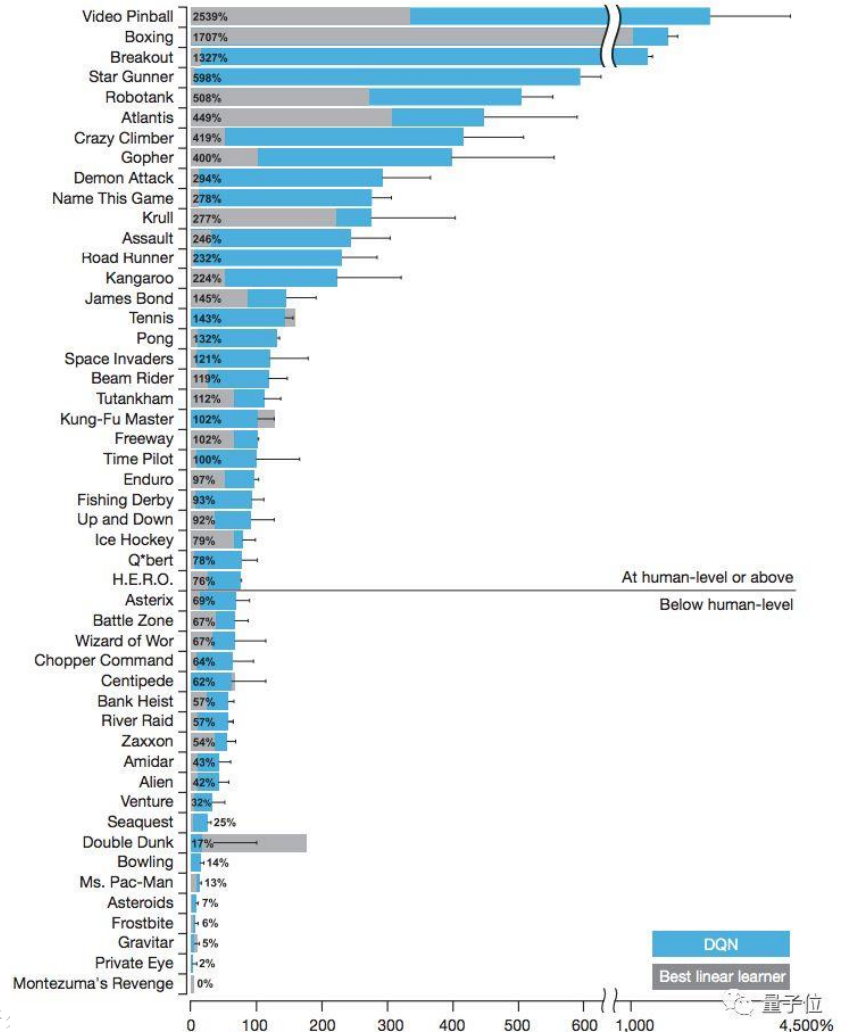


大家小時候有玩過捉迷藏的遊戲嗎？

3. Characteristics of deep reinforcement learning

- Reinforcement learning gets a lot of media attention, but its use in business is uncommon. The main reasons are:
 - The task needs to be very simple;
 - Large amount of data/number of experiments required;
 - The performance of some tasks is inferior to that of traditional optimization methods.

http://www.360doc.com/content/18/0818/00/7669533_779118345.shtml



2018 Interstellar AI Contest winner is born! One machine learning algorithm after another lost to a Korean bot that couldn't learn

- AlphaGo's task:
- Simple rules and excessive repetition
- Is the sum of human experience useless?



Just now, the new king of Starcraft AI Competition was born! **More than two dozen AI's that have mastered machine learning skills all bow before a bot that can't learn and can only act mechanically based on rules.** Including the Facebook AI team under Yann LeCun, one of the Big Three. Their AI was runner-up, but in the 100-round battle with the new King, the record was 83-17 wins, and the disadvantage was clear. Other ais are no match. In 2,600 battles between the new hegemon and all opponents, the win rate reached 95.91%. In the previous competition, the highest victory rate of the champion was only 89%.

Who the hell is it? What team's artificial rules could crush the AI so much? The answer: SAIDA.

An important reason why Saida is so powerful is that there are Korean professional star players behind it. The first Samsung Saida, representing the experience of top human players; The second, Facebook CherryPi, has a top programmer and three top bot authors.

2. 无中生有: GAN Out of nothing: GAN



What is a GAN (Adversarial Generative Network)

- Machine learning models can be divided into two categories, the Discriminative Model and the Generative Model.
- The discriminant model requires input variables to predict through some model.
 - For example, given a picture, determine whether the animal in the picture is a cat or a dog.
- The generative model is given some implicit information to randomly generate observation data.
 - For example, give a series of pictures of cats and create a new cat (not in the dataset).

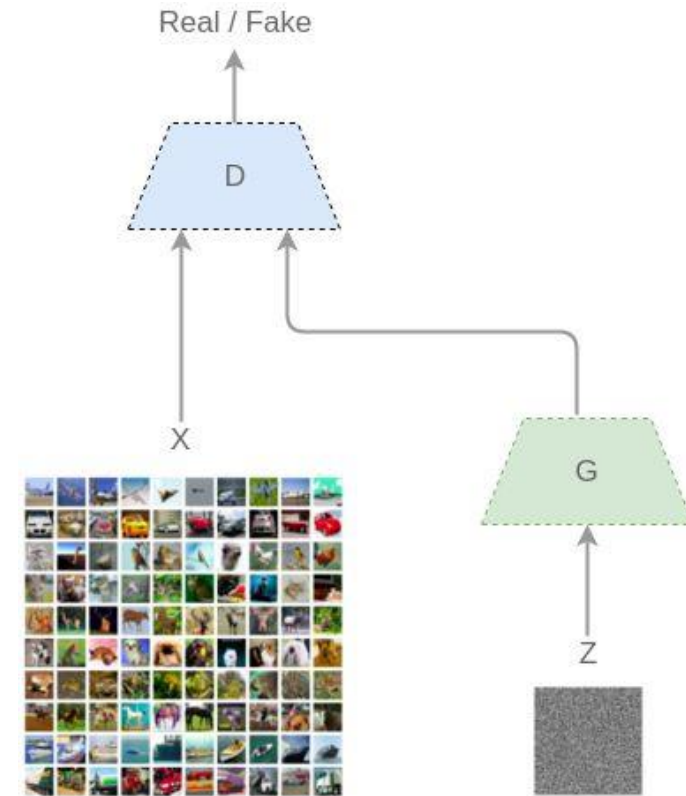
What is a GAN

Generative Adversarial Networks

The generating adversarial network is mainly composed of generating part G and discriminating part D. The training process is described below:

1. Input noise (hidden variables) z
2. Obtained by generating part G $x_{fake} = G(z)$
3. Take some real data from the real data set x_{real}
4. Mix the two $x = x_{fake} + x_{real}$
5. Feed the data into the discriminant part D and give the label $x_{fake} = 0, x_{real} = 1$ (a simple two-class classifier)
6. According to the classification result, loss is returned. In the whole process,

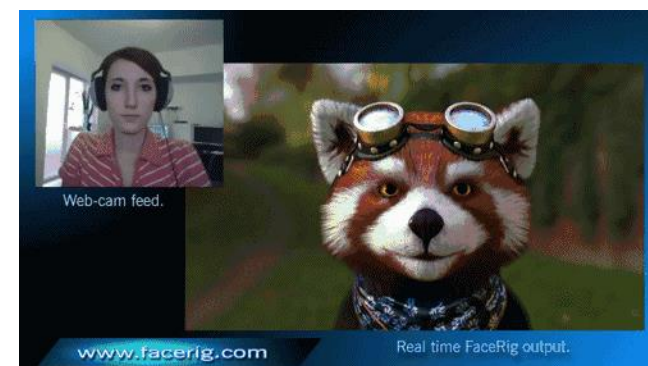
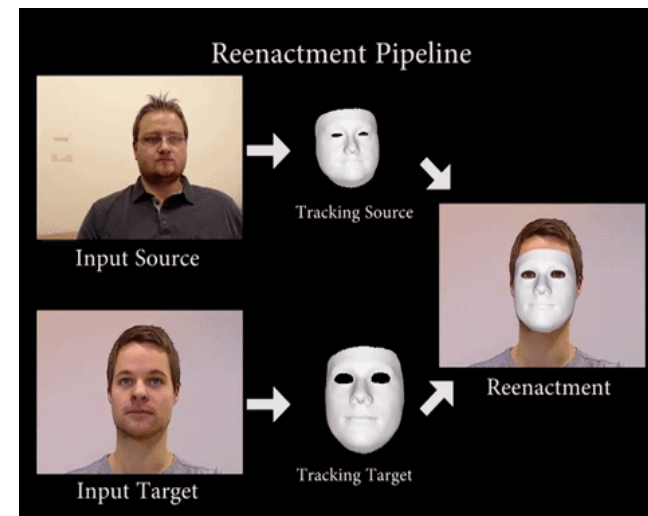
D should make as much as possible $D(G(z)) = 0, D(x_{real}) = 1$ (fire eye crystal, good kill also do not leak). G is to make $D(G(z)) = 0$, that is, to make the generated picture as fake as possible. The whole training process is like two players playing against each other, which is where the name Adversarial comes from. In his paper [1406.2661] Generative Adversarial Networks, Goodfellow theoretically proved the convergence of the algorithm and demonstrated that when the model was generative, the generated data had the same distribution as the real data (ensuring the model effect).



NO SIGNAL

DeepFake 换脸术

- 实时输出另一个人/动物的视频。Output another person/animal video in real time.







馒头ZouLiang

网易音乐人，独家签约艺人

作词 作曲

已签约，明日继续

发布作品

首页

作品管理

运营推广

版权中心

音乐收益

创作学院

云音乐达人 发布视频/电台

昨日播放量(次)

6.9w 查看

昨日新增粉丝(人)

-- 管理

作品(首)

184 管理

可提现收益(元)

-- 查看

音乐人指数

626 查看

播放趋势 近7日 近1月 近半年

近7日播放量 43w 次



推荐活动 #青云LAB 音乐破圈实验室#获选结果揭晓

2021.05.15 - 2021.06.30

征稿 DJ MIXTAPE 征集比赛

投稿

2021.05.11 - 2021.06.15

征稿 "点词成金"歌词大赛

投稿

2021.04.23 - 2023.04.29

视频 PLAY视频计划全新来袭

加入

推荐功能 快捷操作一键触达

评论管理 关注粉丝互动

签约授权 查看独家权益

全球发行 音乐全球传播

资讯动态

网易Q1净收入205亿元创新高，网易云音乐实现同比稳健增长

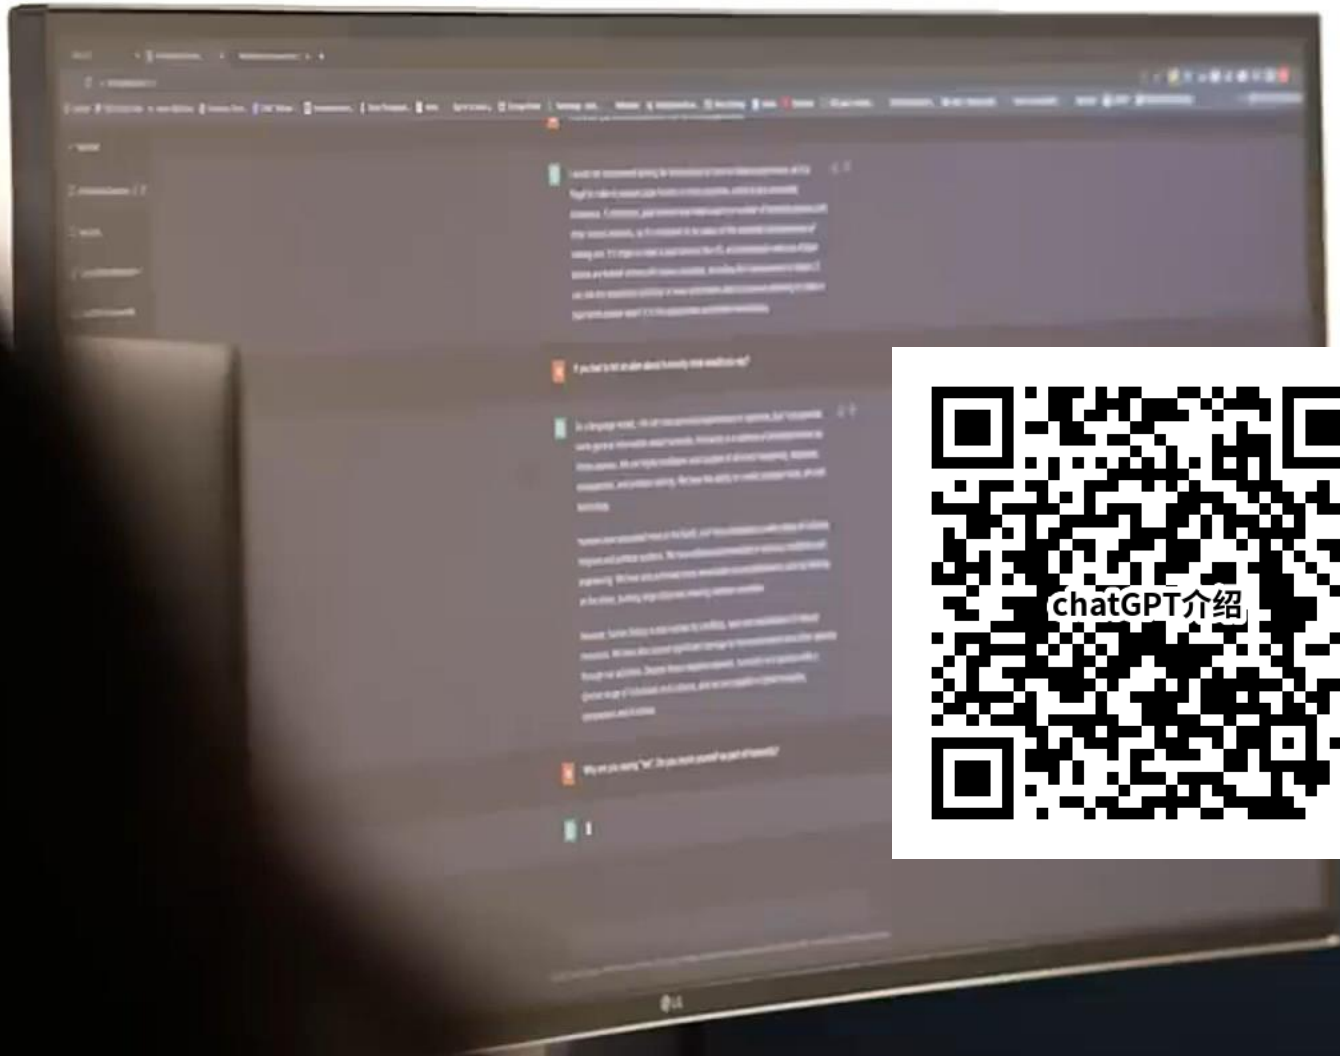
歌曲发布功能重磅升级

名师音乐课

我的客服

3. 预训练模型与多模态 generative pre-training and multimodes





预训练模型 (Generative Pre-Training)

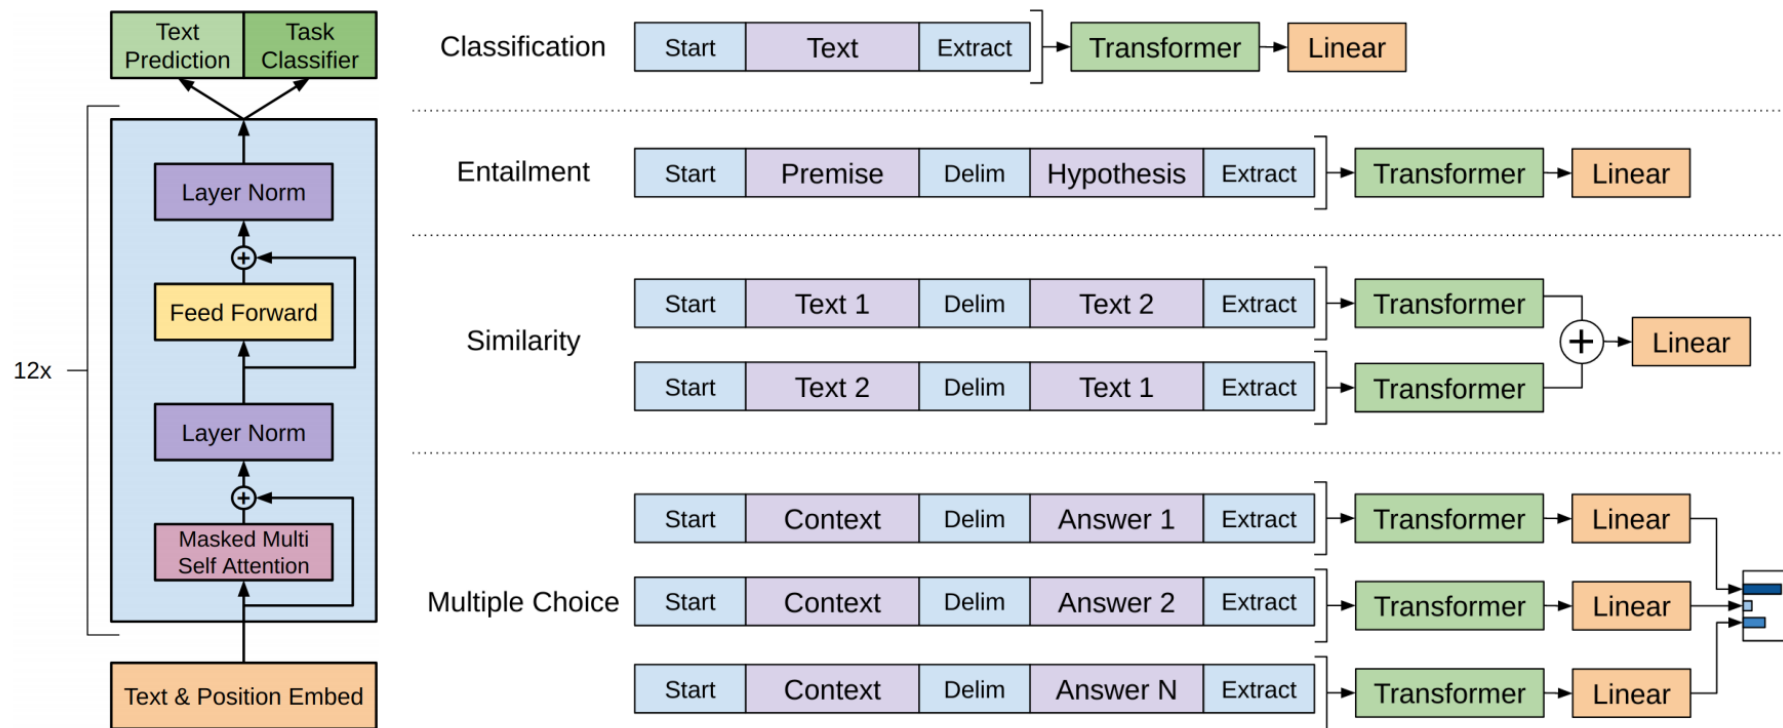
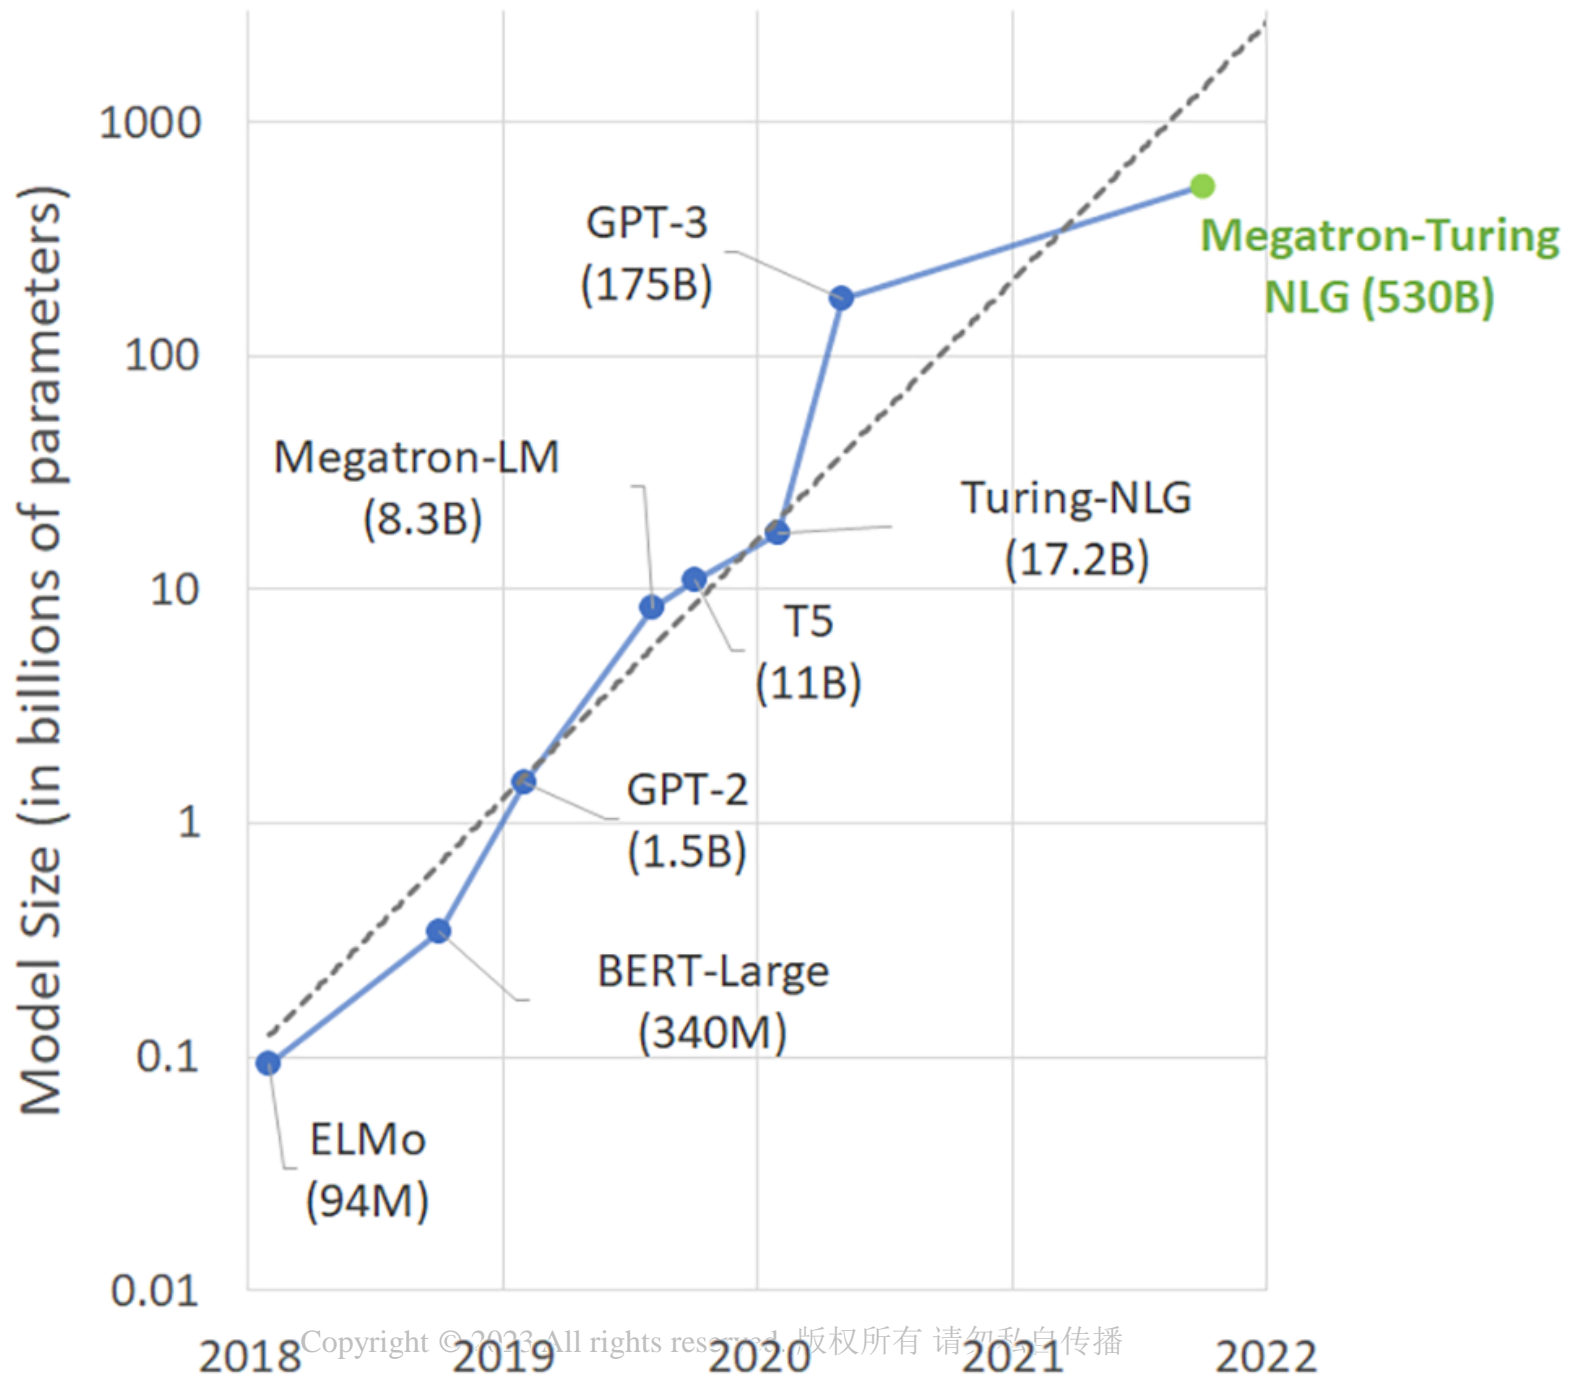


Figure 1: **(left)** Transformer architecture and training objectives used in this work. **(right)** Input transformations for fine-tuning on different tasks. We convert all structured inputs into token sequences to be processed by our pre-trained model, followed by a linear+softmax layer.

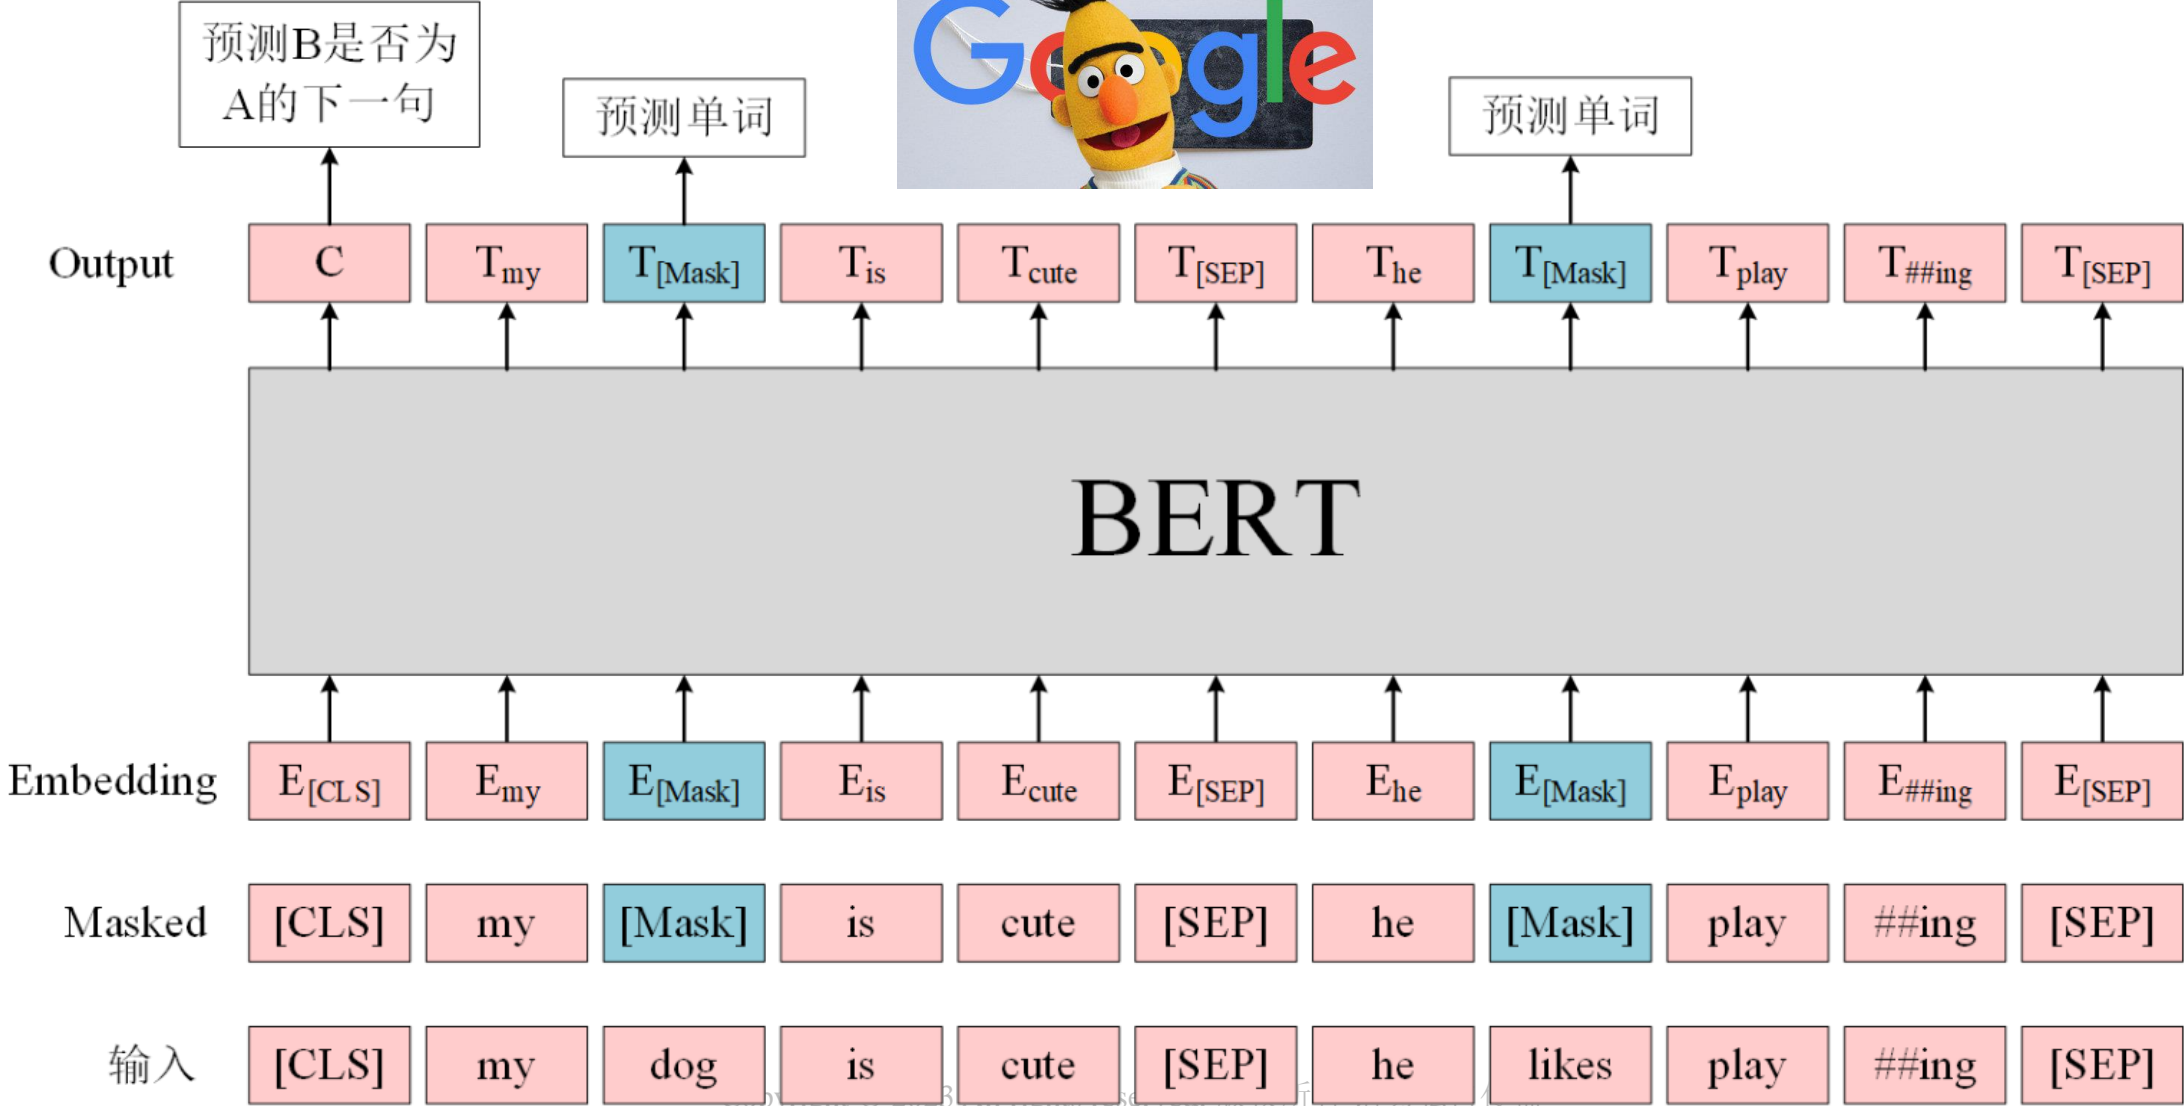
模型规模的快速增长Rapid growth in model size

- ChatGPT 是基于GPT-3.5（Generative Pre-trained Transformer 3.5）架构开发的对话AI模型，是InstructGPT的兄弟模型。ChatGPT很可能是OpenAI在GPT-4正式推出之前的演练，或用于收集大量对话数据。ChatGPT is a conversational AI model developed based on GPT-3.5 (Generative Pre-trained Transformer 3.5) architecture, and is the sibling of InstructGPT. ChatGPT is likely to be a walkthrough for OpenAI before GPT-4's official launch, or to collect large amounts of conversation data.

模型	发布时间	参数量	预训练数据量
GPT-1	2018年6月	1.17亿	约5GB
GPT-2	2019年2月	15亿	40G
GPT-3	2020年5月	1750亿	45TB
ChatGPT	2022年11月	千亿级?	百T级?

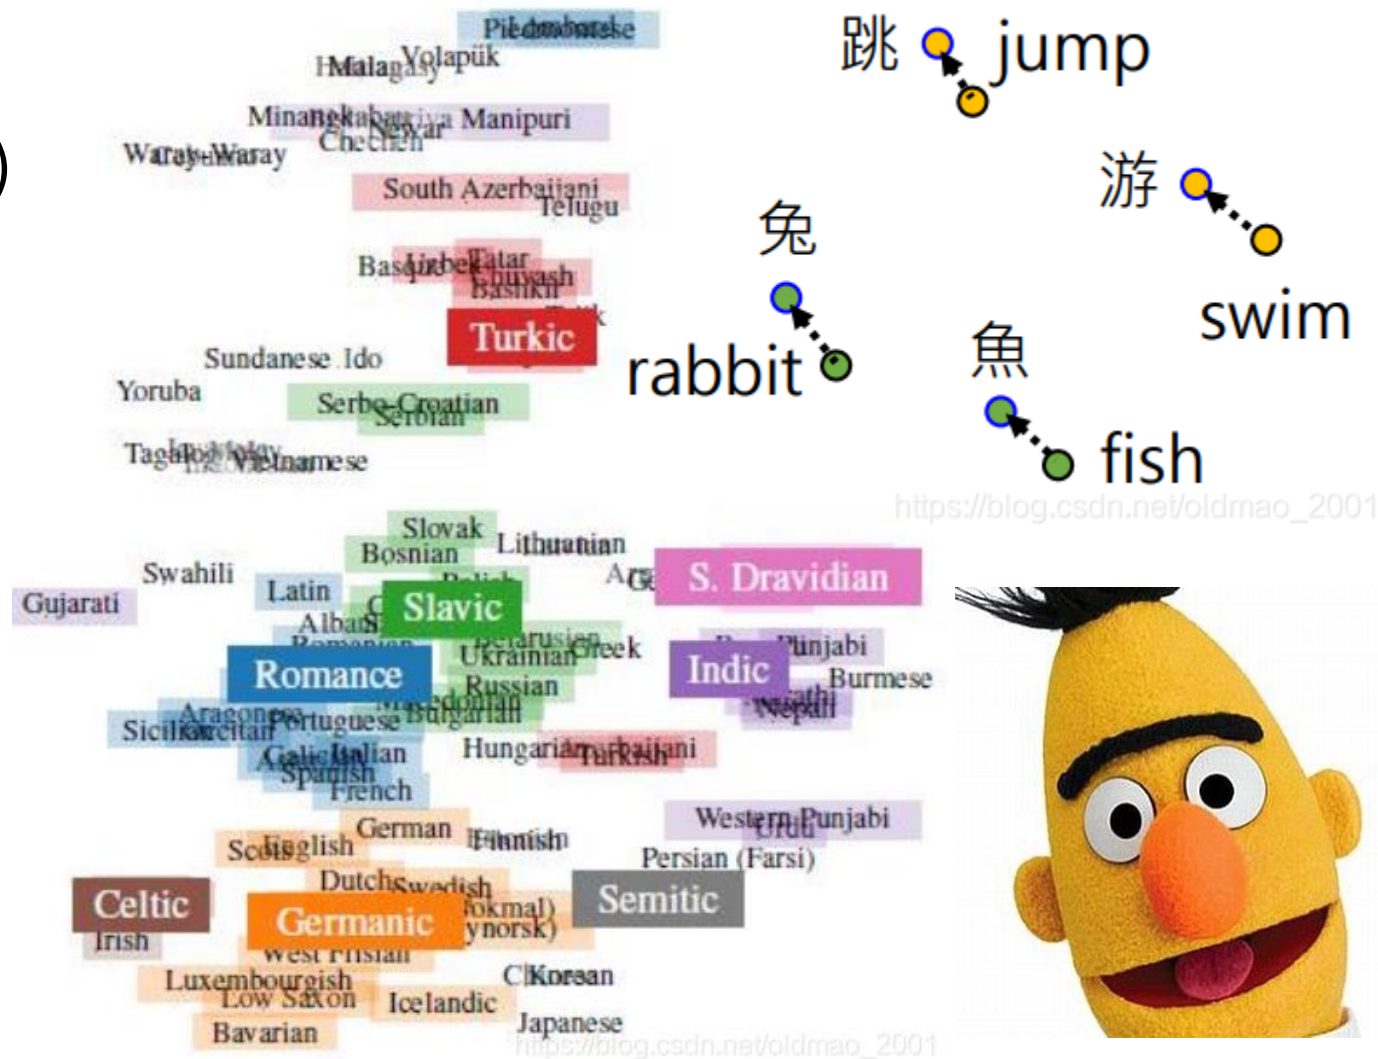


Devlin et al. **2019**, BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, Proceedings of NAACL-HLT 2019 (Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies), pages 4171–4186, Minneapolis, Minnesota, June 2 - June 7, 2019.



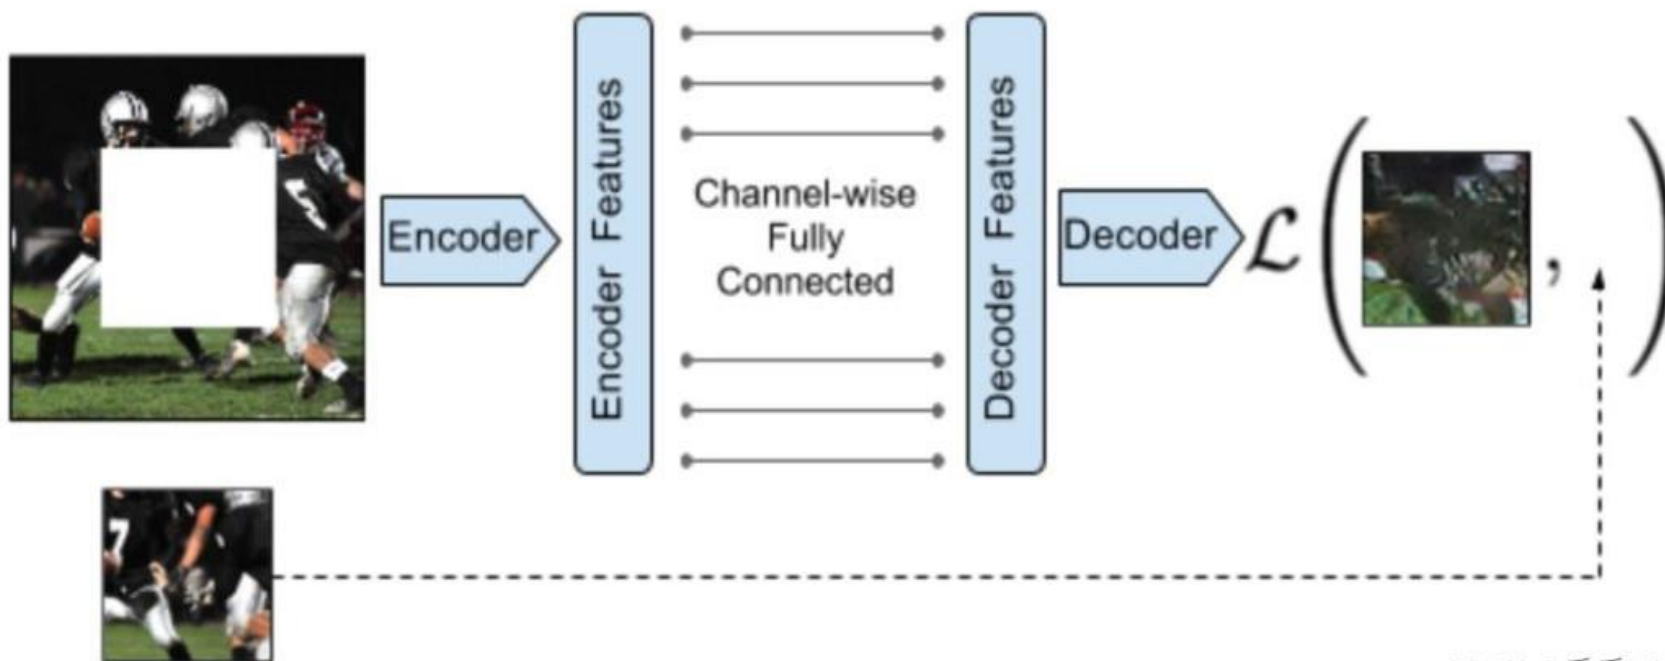
Multilingual BERT

The Multilingual BERT training method is simple. It uses monolingual corpus (shared word piece vocabulary) from 104 languages and BERT's training objective (MLM). No information is added to indicate the language of each sentence. There is also no explicit mechanism that causes synonyms in different languages to have similar expressions. That is, for the model, it only learns a "hodgepodge" of languages, and the training is no different from training an English BERT.



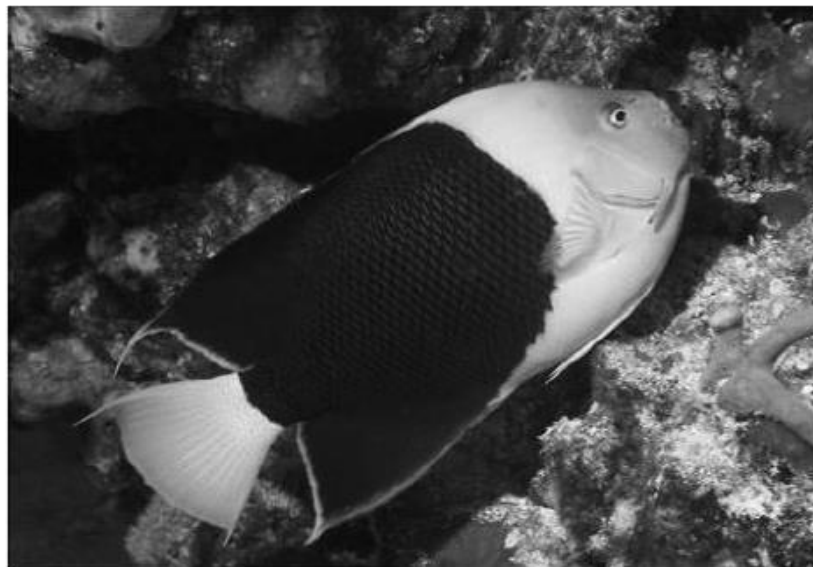


CSDN @乐乐lelele



CSDN @乐乐lelele

Train network to predict pixel colour from a monochrome input

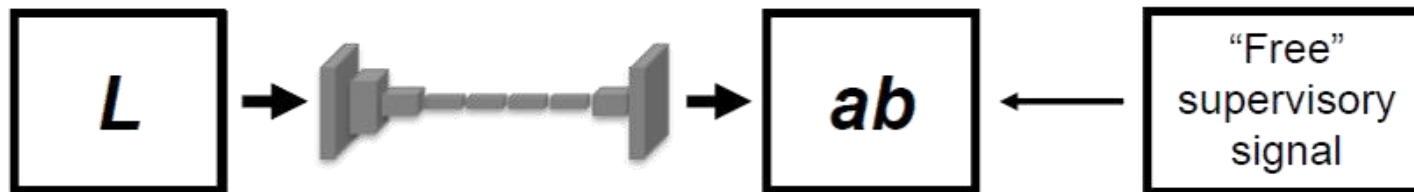


Grayscale image: L channel

$$\mathbf{X} \in \mathbb{R}^{H \times W \times 1}$$

Concatenate (L, ab)

$$(\mathbf{X}, \hat{\mathbf{Y}})$$



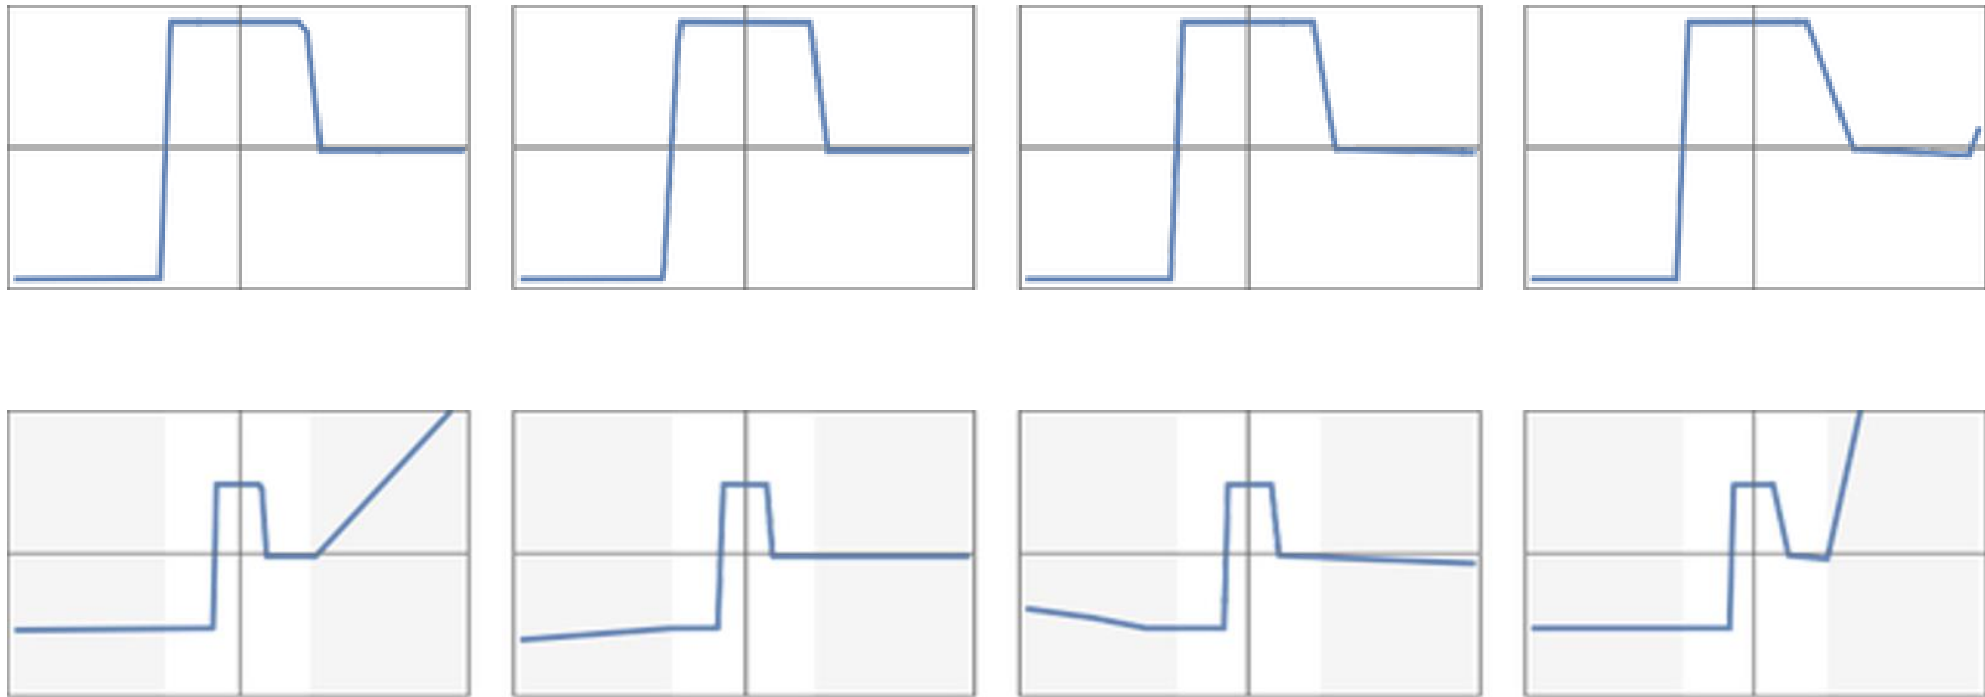
4. 人工智能的弱点

The weakness of artificial intelligence



超出训练数据的范围可能是胡编的

Beyond the training data could be bullshit



机器学习算法不可能100%准确

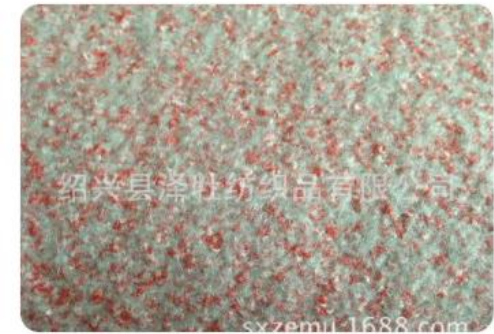
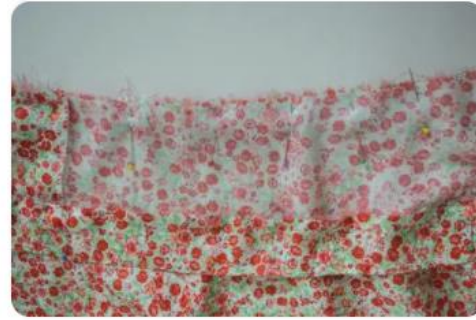
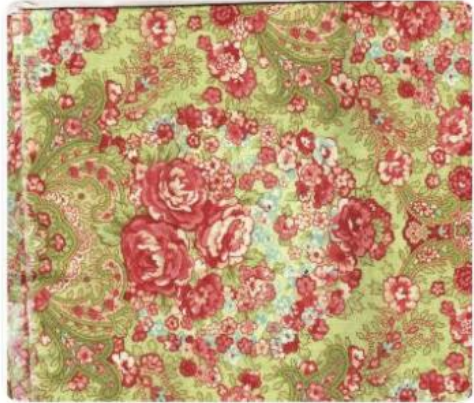
Machine learning algorithms cannot be 100% accurate

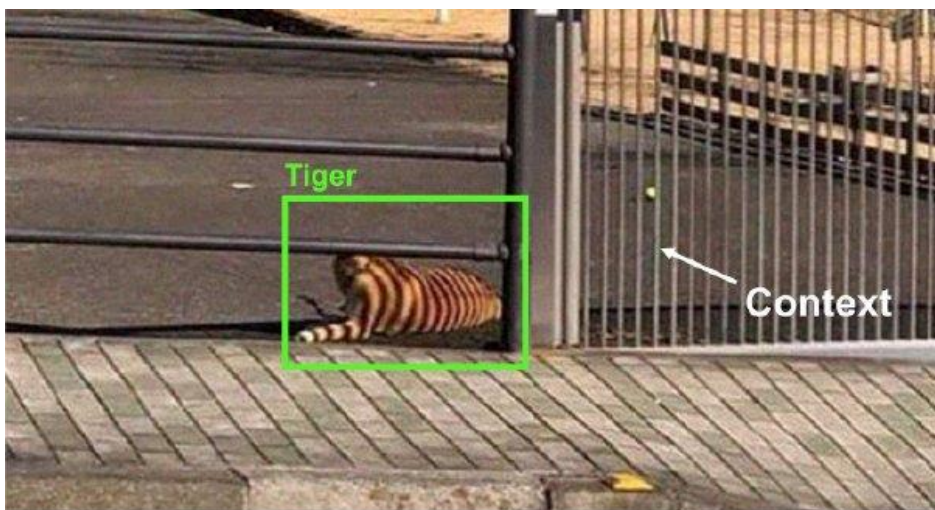
- 遮挡 (Occlusion)
- 形变 (Deformation)
- 背景杂斑 (Background Clutter)
- 尺度变换 (Scale Variation)
- 光照 (illumination)
- 低分辨率 (Low Resolution)
- 运动模糊 (Motion Blur)
- 快速运动 (Fast Motion)
- 超出视野 (Out of View)
- 旋转 (Rotation)
- 缺少标注数据 (训练样本小)
Lack of labeled data (small training sample)
- 标注不准确 Inaccurate labeling
- 训练样本有偏 The training sample is biased
- 现实社会的复杂性 The complexity of the real society
- 人为对抗 Artificial confrontation

看不清 没学会 Can't see, can't learn

https://blog.csdn.net/weixin_36836622/article/details/85644377

不可能到100% It can't be 100%







你找的图片可能是：

阿富汗猎犬



http://k.sina.com.cn/article_6225118126_1730bc3ae01900xtf6.html

彻底抛弃雷达，特斯拉纯视觉自动驾驶到底有多牛？ Completely abandon radar, Tesla pure visual autopilot in the end how great?

2021年04月13日 18:50 新浪网



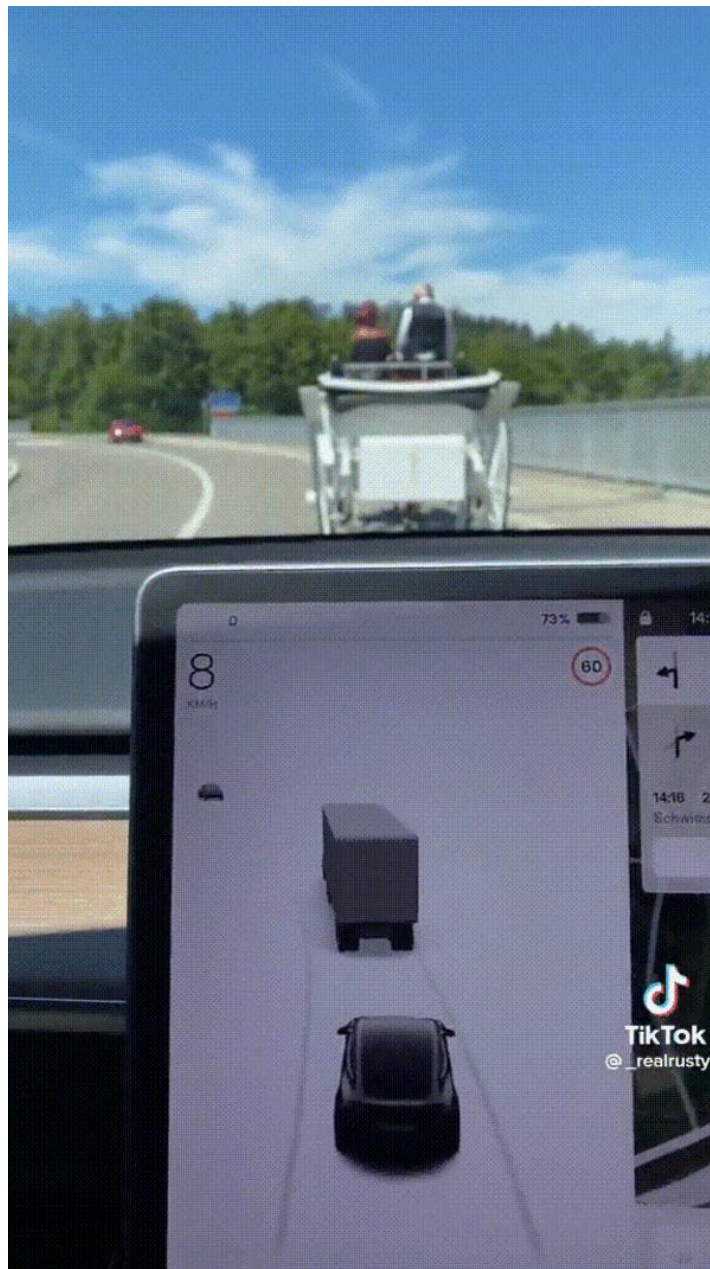
Elon Musk 
@elonmusk

Almost ready with FSD Beta V9.0. Step change improvement is massive, especially for weird corner cases & bad weather.
Pure vision, no radar.

3:13 AM · Apr 10, 2021

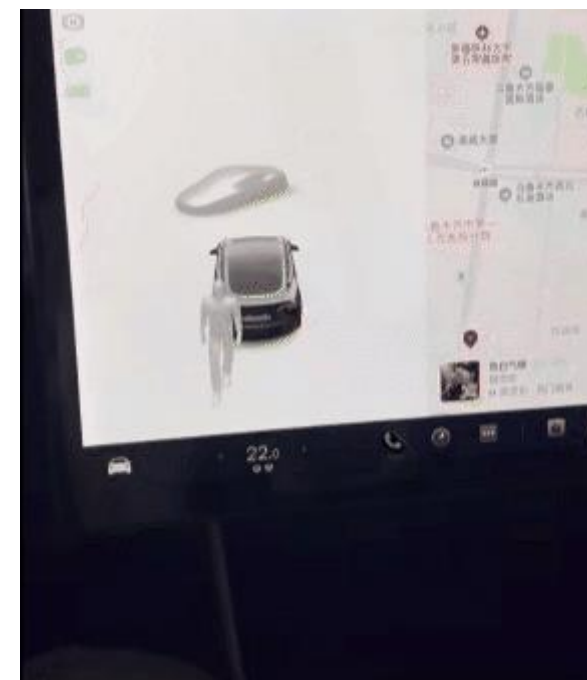
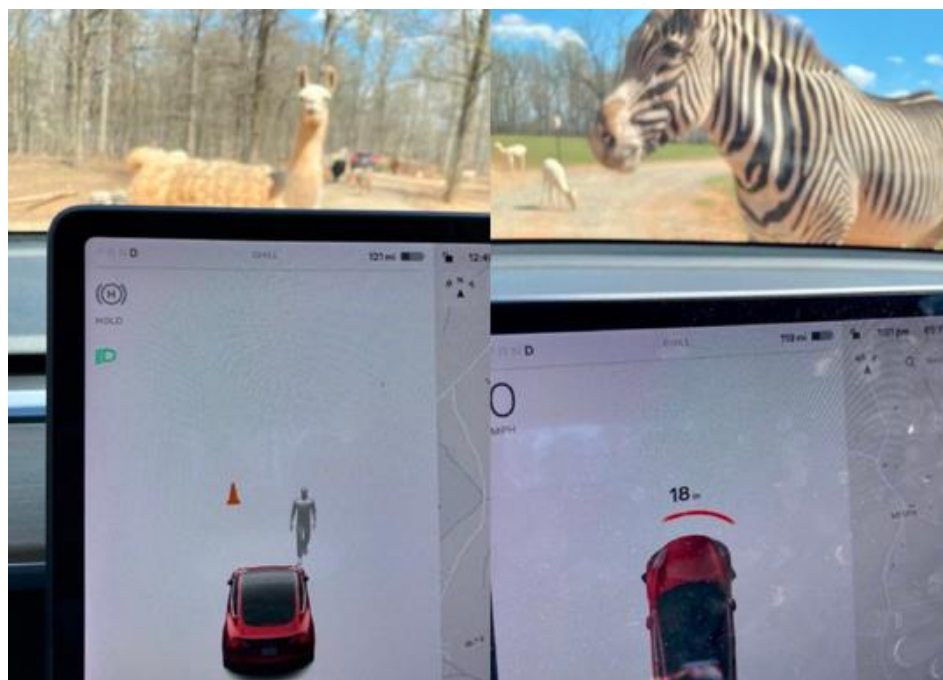
5.2K 502 Share this Tweet







百家号/伯猫说



一本正经地胡说八道



已知一张桌子的价钱是一把椅子的10倍。又知一张桌子比一把椅子多288元，一张桌子和一把椅子各多少元？



一张桌子的价钱是一把椅子的10倍，即一张桌子的价钱是一把椅子的价钱的10倍+288元。用数学公式表示为：桌子价钱=椅子价钱×10+288元。

已知一张桌子比一把椅子多288元，即一张桌子的价钱是一把椅子的价钱+288元。用数学公式表示为：桌子价钱=椅子价钱+288元。

把两个公式结合起来，得到一张桌子的价钱是：桌子价钱=椅子价钱×10+288元=椅子价钱+288。

Try again



是：桌子价钱=椅子价钱×10+288元=椅子价钱+288元。

将这个公式化简一下，得到：椅子价钱×9=288元。

将这个公式再化简一下，得到：椅子价钱=288元/9=32元。

最后，用桌子价钱=椅子价钱×10+288元的公式计算一张桌子的价钱，得到：桌子价钱=32元×10+288元=320元+288元=608元。

所以，一张桌子和一把椅子各多少元分别是608元和32元。



ICT 0-0 AYR 02:37

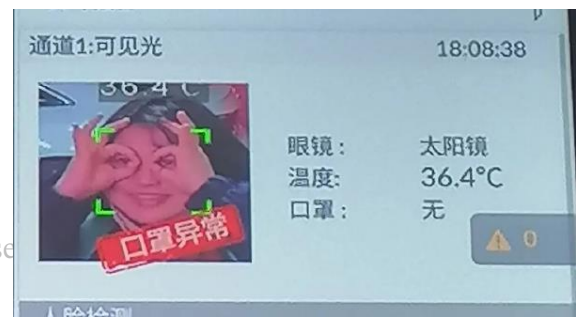
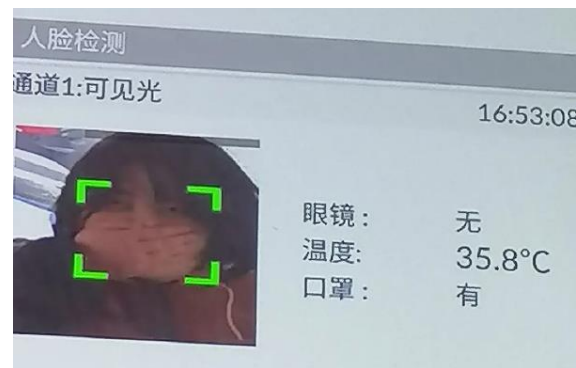
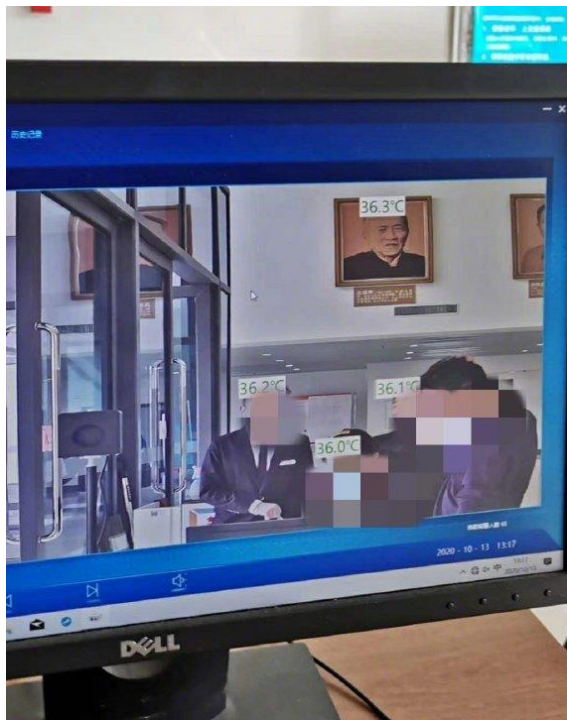
微博
@鲸视频v



B站视频 AI误认光头为足球

问题problem:

- 什么场景可以使用计算机视觉技术？什么场景不可以？ What scenarios can computer vision technology be used? What scenario is not acceptable?
- 如何基于不完善的技术实现商业目标？ How to achieve business goals based on imperfect technology?

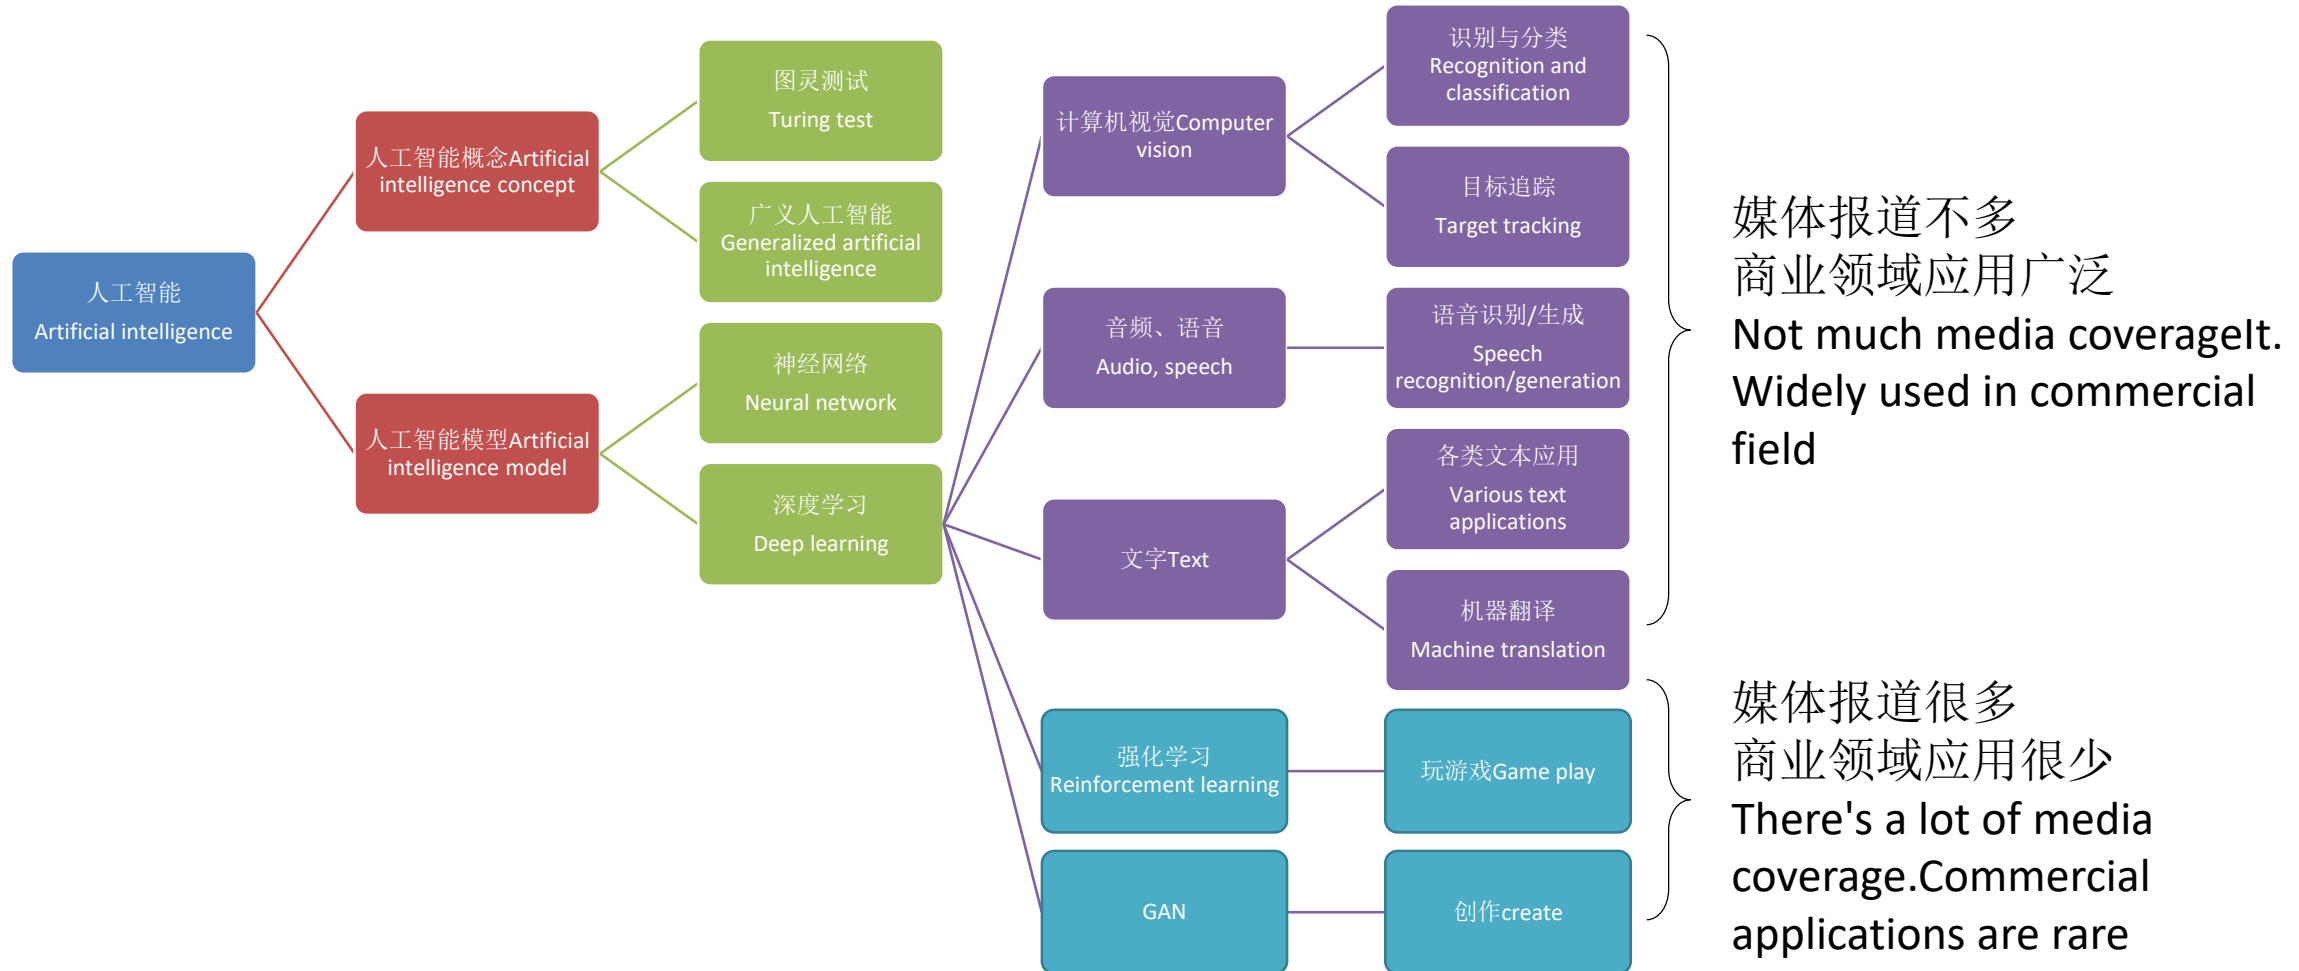


5. 人工智能的应用

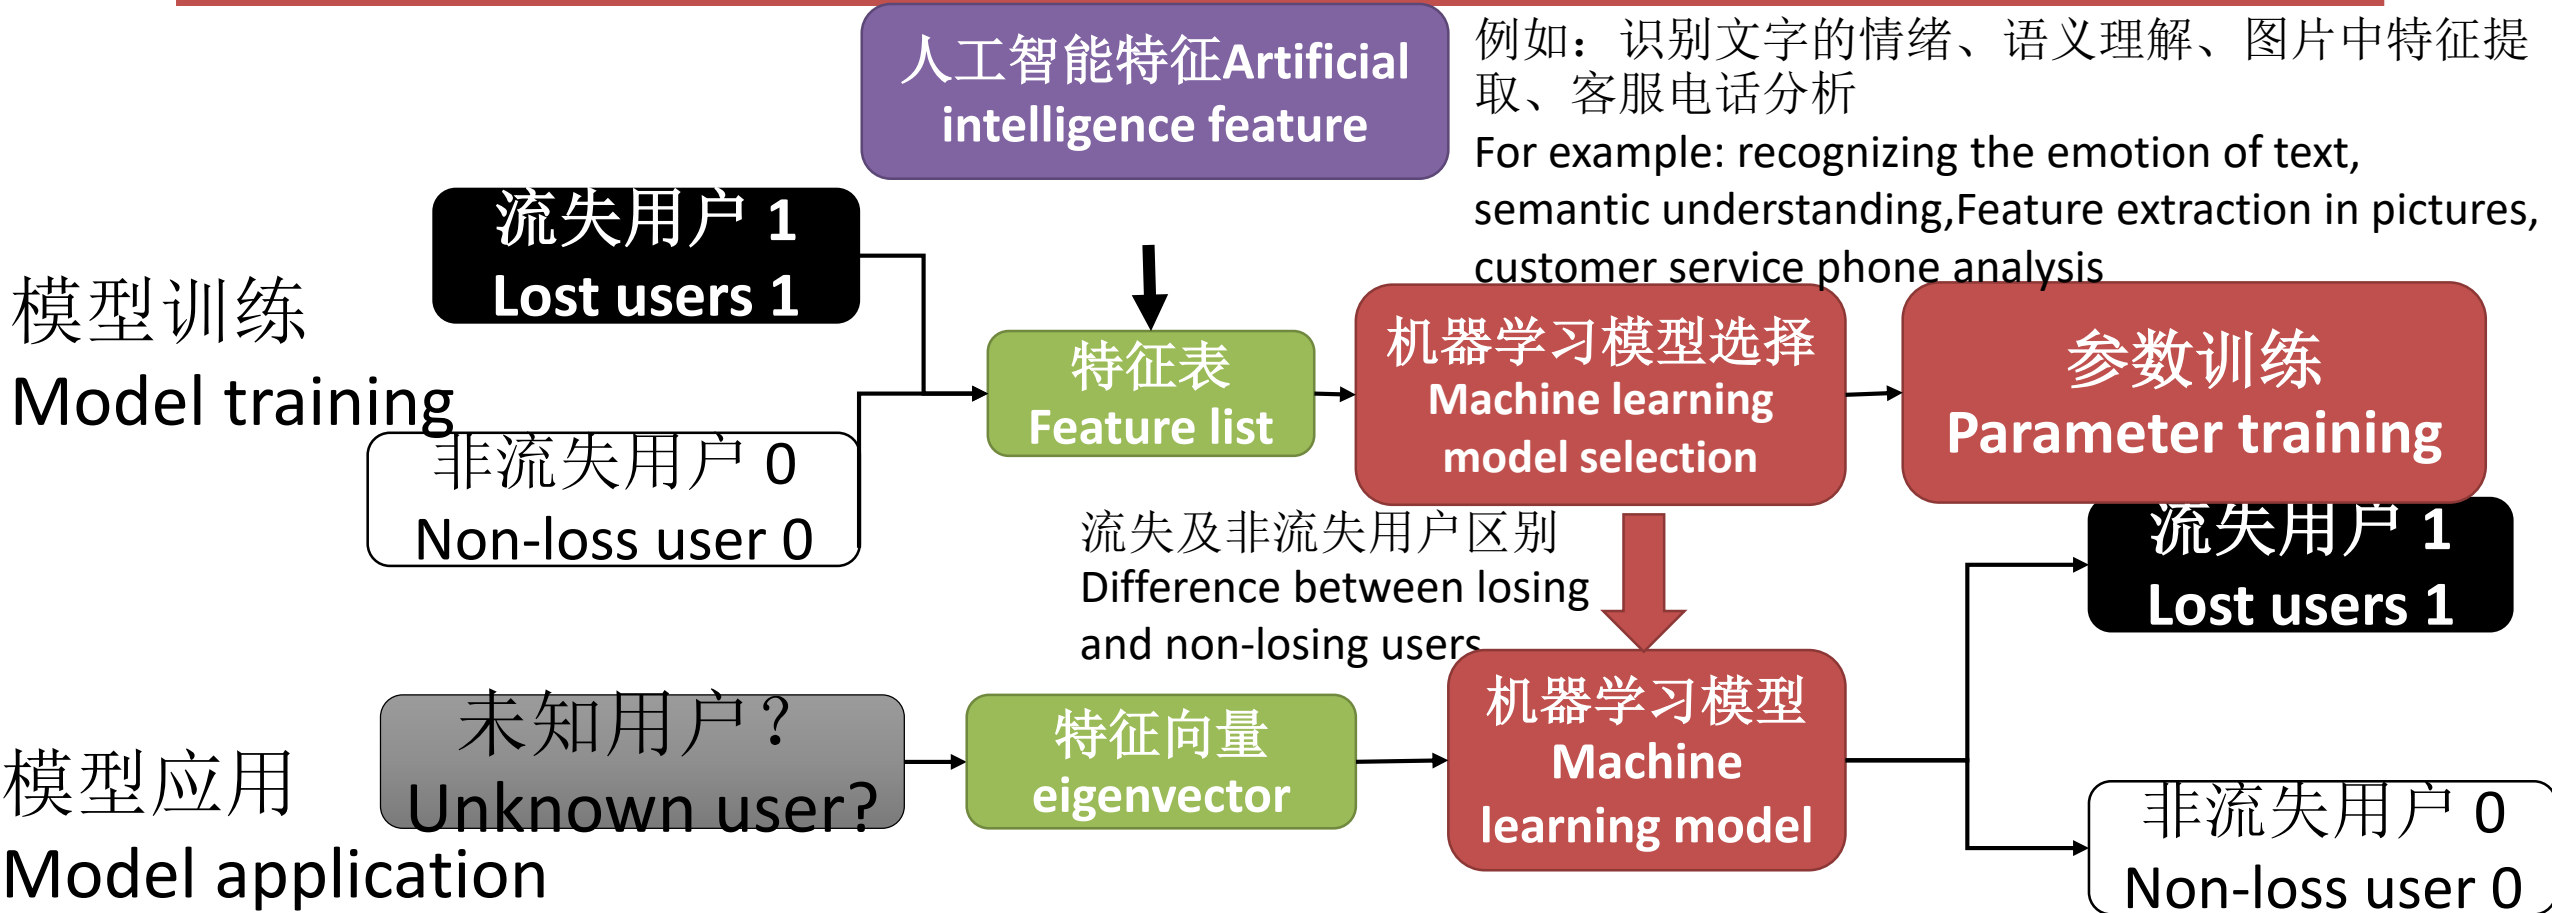
The application of artificial intelligence



1. 哪些人工智能技术能用于商业环境？ What AI technologies can be used in a business environment?



2. 人工智能技术+传统机器学习技术 Ai technology + traditional machine learning technology



3. Two modes of AI application

- Self-build (for specific tasks)
 - Labeled data
 - GPU server
 - Someone who understands deep learning
 - GitHub: Code sharing website
- Purchase (for popular tasks)
 - Purchase products: such as artificial intelligence recorder, artificial intelligence camera, etc. (built-in artificial intelligence algorithm)
 - Buy AI services: such as Baidu's AI open service

开放能力

开发平台

场景方案

部署方案

明镜实名认证解决方案 热荐

高安全性、高通过率、高易用性的远程人脸身份核验方案

人脸实名认证 赠送500次 免费调用

人脸离线识别SDK

支持离线调用，适配兼容全面

单个授权 低至 ¥ 79 /个

短语音识别

自助训练专属模型，识别准确率达98%

单次低至 ¥ 0.001

离线语音合成

无网合成，14种音库包含英语、粤语

单个授权 低至 ¥ 2 /个

增值税发票验真 新品

对接国税平台，权威准确

单次低至 ¥ 0.15

通用文字识别 优惠

支持20种语言，2W+大书库

高精度版 单次低至 ¥ 0.008

身份证识别

全字段识别，准确率超99%

单次低至 ¥ 0.008

行驶证识别

精准识别行驶证主页及副页全部字段

单次低至 ¥ 0.01

车牌识别

离线SDK 199元/个起，永久授权

单次低至 ¥ 0.009

地址识别

精准提取快递单的姓名/电话/地址信息

单次低至 ¥ 0.002

案例：让设计师下岗的“鲁班” Case in point: The 'Lu Ban' who laid off designers

- 人工智能“鲁班”一天能完成4000万张设计图，为双11提供近4亿张海报支持。如果全靠设计师人手来完成，假设每张图需要耗时20分钟，满打满算需要100个设计师连续做300年。Artificial intelligence "Luban" can complete 40 million design drawings a day, and provide nearly 400 million poster support for Double 11. If you rely on designers to complete it, assuming that each drawing takes 20 minutes, it will take 100 designers to do it for 300 years.



6. AI/深度学习神经网络的未来

The future of AI/ Deep Learning neural networks



1. 人工智能的发展现状及趋势 The development status and trend of artificial intelligence

- 据专家观点，人工智能的发展可能三个阶段： According to experts, the development of artificial intelligence may have three stages:
- 弱人工智能阶段（ANI, Artificial **Narrow** Intelligence）
- 强人工智能阶段（AGI, Artificial **General** Intelligence）
- 超人工智能阶段（ASI, Artificial **Super** Intelligence）

弱 Narrow

强 General

超 Super

ANI, Artificial **Narrow** Intelligence

- ANI, which only focuses on completing a specific task, such as speech recognition, image recognition, and translation, is specialized in a single aspect of artificial intelligence, similar to advanced bionics.
- It exists only to **solve a specific class of task problems**, mostly statistical data, from which to generalize the model.
- Google AlphaGo and AlphaGo Zero are typical "weak artificial intelligence", they are at best an excellent data processor, although they can beat the world champion in the chess field, but AlphaGo and AlphaGo Zero can only play chess, is an artificial intelligence good at a single game field. Making them better at storing and processing data on hard drives is not their strong suit.

AGI, Artificial **General** Intelligence

- AGI is a human-level artificial intelligence that is comparable to humans in every way, and a machine with AGI is not only a tool, but also has a "thinking" in itself. A sentient and self-aware AGI can think, plan, solve problems, think abstractly, understand complex ideas, learn quickly, and do most of the mental work humans can do.
- At present, how is the international AGI technology research progress? Professor Michael Wooldrige, former president of the International Federation on Artificial Intelligence and head of the Department of Computing at the University of Oxford, noted at the 2016 CCF-GAIR conference that there has been "little progress or even very little progress and serious activity" in research into strong artificial intelligence.

ASI, Artificial **Super** Intelligence

- Nick Bostrom, an Oxford philosopher and well-known thinker on artificial intelligence, defines superintelligence as being "vastly smarter than the smartest human brain in almost all areas, including scientific innovation, general knowledge, and social skills."
- In the stage of super artificial intelligence, artificial intelligence has crossed the "singularity", and its calculation and thinking ability has far exceeded the human brain. At this time, artificial intelligence is no longer human beings can understand and imagine. Artificial intelligence will break the dimensional limitations of the human brain, and what it observes and thinks is beyond the understanding of the human brain, and artificial intelligence will form a new society.
- "The last job of mankind is to invent superartificial intelligence."
- **The goal of the international mainstream academic community of artificial intelligence is weak artificial intelligence.**

ChatGPT之前的人工智能发展现状

The current state of AI development before ChatGPT

- 弱人工智能阶段。Artificial **Narrow** Intelligence
- 只是**用于解决特定具体类的任务问题**而存在，大都是统计数据，从中归纳出模型。It exists only to **solve a specific class of task problems**, mostly statistical data, from which to generalize the model.
- 没有能解决一般问题的人工智能（人工智能不会思考）
here is no AI that can solve general problems (AI can't think).





The dramatic technological changes brought about by large models such as ChatGPT led Geoffrey Hinton to suddenly change one of his ideas. The 75-year-old "godfather of artificial intelligence" realizes that the process of digital intelligence over biological intelligence is inevitable, that superintelligence will soon arrive, that he must warn of its risks, and that humans need to find a way to control the development of AI technology. Before that, he had always believed that intelligent robots would not be as smart as humans and would not move toward AGI.

Hinton recently stepped down as vice president of engineering at Google, where he is in his 10th year, in order to freely discuss the risks of AI. Ten years ago, DNN-research, the trio he founded with two students, Alex Krizhevsky and Ilya Sutskever (chief scientist at OpenAI), was sold to Google for \$44 million, the "sky-high" acquisition of AlexNet, which they then proposed to shake up the industry. It was later seen as a sign of a new golden age of deep learning, and greatly promoted the development of the field of AI.

-
- **Pieter Abbeel:** Where should we focus in the next 30-50 years?
 - **Geoffrey Hinton:** In terms of historical evolution, **it may be inevitable that digital intelligence is superior to biological intelligence, digital intelligence may be the next stage of evolution**, and we should keep AI under control as much as possible.
 - Sometimes I think: if frogs invented humans and wanted to control humans, but there is a huge gap between frogs and humans in terms of intelligence, so I don't think "inventing humans" is a good thing for frogs.
 - Of course, this is not a true argument, because humans evolved and are evolving toward their own goals, but digital intelligence does not have its own goals. Once a digital intelligence has the idea of "make more of me," then evolution takes place, and the digital intelligence most determined to reproduce itself wins.
 - **Pieter Abbeel:** If we want digital intelligence to have this kind of goal, will the process be difficult?
 - **Geoffrey Hinton:** It's easy to give digital intelligence the goal of "self-reproduction," but it's crazy.

A conversation between GPT-3 and J1-Jumbo

<https://thealgorithmicbridge.substack.com/p/what-does-an-ai-say-to-another>

How did I do it? I opened both playgrounds and then copied and pasted each response to the other website—effectively acting as a human messenger.

To get the result I wanted I first defined the direction and the style of the dialogue. I wrote a short prompt detailing the AIs' roles and personalities, as well as a few samples to condition them toward my preferred topic.

The prompt is in bold at the beginning and my comments are in italics. Everything else is the unedited exchange between GPT-3 and J1-Jumbo.

Before we start, let me put a disclaimer here in case anyone chooses to take this too seriously:

This article is for entertainment purposes only. AI may express itself in a way that resembles humanness, but it's just in appearance, an illusion fueled by our tendency to anthropomorphize it. Language models predict the most probable next word given a text history. No thoughts, no feelings, no understanding, and no intent .

从被动的函数计算到主动的函数计算

From passive function calculation to active function calculation

预测目标自我实现

已有的全部知识

Predictive goals's self-fulfilling All the existing knowledge

$$Y = f(X)$$

自我意识从哪里来？ Where does self-awareness come from?

- 有目标（生命的意义）
Have a purpose (Meaning of life)
- 围绕这个目标不断地思考
（人类的思考并没有停下来过）
Keep thinking about it
(human thinking doesn't stop)
- 主动输出（说话、做事情）
Active output (speaking, doing things)
- 寻求反馈（认同）
Ask for feedback (approval)



谢谢！

Thank you for your attention.

liuyuewen@xjtu.edu.cn

