

Introduction to Machine Learning

Lecture 00: Introduction

Sep. 5, 2023

Jie Wang

Department of Electronic Engineering and Information Science (EEIS)

<http://staff.ustc.edu.cn/~jwangx/>

jiawangx@ustc.edu.cn



Machine Intelligence Research and Applications Lab



- **Logistics**
- **Overview**

Logistics

Course Info.

- **Instructor:** Jie Wang (jjewangx@ustc.edu.cn, jiewang@miralab.ai)
- **TA:** Haoyang Liu, Qitan Lv, Yinqi Bai, and Haotong Huang
{haoyangliu, Qitan Lv, Yinqin Bai, and Haotong Huang}@miralab.ai
- **Website:** https://miralab.ai/course/ml_2023fall/

- **Groups:**



- **Lab Assistant:**



扫一扫上面的二维码图案，添加我为好友



Haoyang Liu

- First-year graduate student
GPA **4.01** during undergraduate study
- S.-T. Yau College Student Mathematics Contest, **outstanding prize**, 2022
- The Chinese Mathematics Competitions, **National Second Price**, 2023



Qitan Lv

- First-year graduate student
- **National Scholarship 2021, 2022**
- The Chinese Mathematics Competitions, **First Price**, 2020



Yinqi Bai

- Undergraduate,
Overall GPA: **3.79**,
4/68 Department of Automation
- Outstanding Student Scholarship **Gold Award** of University of Science and Technology of China, 2022





Haotong Huang

- Undergraduate,
Overall GPA: **3.95**,
3/81 Statistics
- The Chinese Mathematics Competitions, **First Price**, 2021


Prerequisites

- Pattern Recognition or Introduction to Artificial Intelligence
- Programming



- **Solid** background in
 - calculus
 - linear algebra
 - probability
 - stochastic process
 - ...

- No required textbook
- Slides and lecture notes will be posted
- Some excellent textbooks



- **Supervised Learning**

- Linear Regression
- Elementary convex optimization
- Sparse Learning
- Naïve Bayes
- Logistic Regression
- Lagrange Duality
- Support Vector Machine
- Decision Tree
 - Random Forest
 - GBDT
- Neural Network

- **Computational Learning Theory**

- **Unsupervised Learning**

- K-means
- Mixture Models and EM
- Principle Component Analysis

- **Reinforcement Learning**

- Multi-armed Bandits
- Finite Markov Decision Processes
- Dynamic Programming
- Temporal Difference Learning

- Credit distribution

	Homework	Mid-Term Exam	Final Exam	Project
Credit	20	20	40	20
Date	TBD	TBD	TBD	TBD

- Start early
 - Latex is recommended for typeset

- No cheating
 - Finish your own work independently
 - Discussion is encouraged, but no sharing of code or text
 - Explicitly indicate references that you use in your homework and projects
- Deadline is **DEADline**
 - No late submissions shall be accepted
- Be polite in any form of discussion
 - Serious offense will lead to expulsion
- **Ignorance** is **NO** excuse
- Appeal within two weeks, otherwise final

Questions



Overview

What is Artificial Intelligence

The Beginning of A.I.



会议原址：达特茅斯楼

1956年夏天，约翰·麦卡锡等人在美国达特茅斯学院开会研讨“如何利用机器模拟人的智能”，会上提出“人工智能”(Artificial Intelligence, AI)这一概念，标志着人工智能学科的诞生。

参会部分人员：

克劳德·香农 (Claude Shannon) :
贝尔实验室，信息论创始人

约翰·麦卡锡(John McCarthy):
达特茅斯学院，图灵奖得主

马文·明斯基 (Marvin Lee Minsky) :
图灵奖得主，麻省理工学院人工智能实验室的创始人

司马贺 (Herbert Simon) :
图灵奖、诺贝尔奖得主

奥利弗·赛弗里奇 (Oliver Selfridge) :
模式识别奠基人



2006年，达特茅斯会议五十年后，当事人重聚达特茅斯。左起：摩尔，麦卡锡，明斯基，赛弗里奇，所罗门诺夫

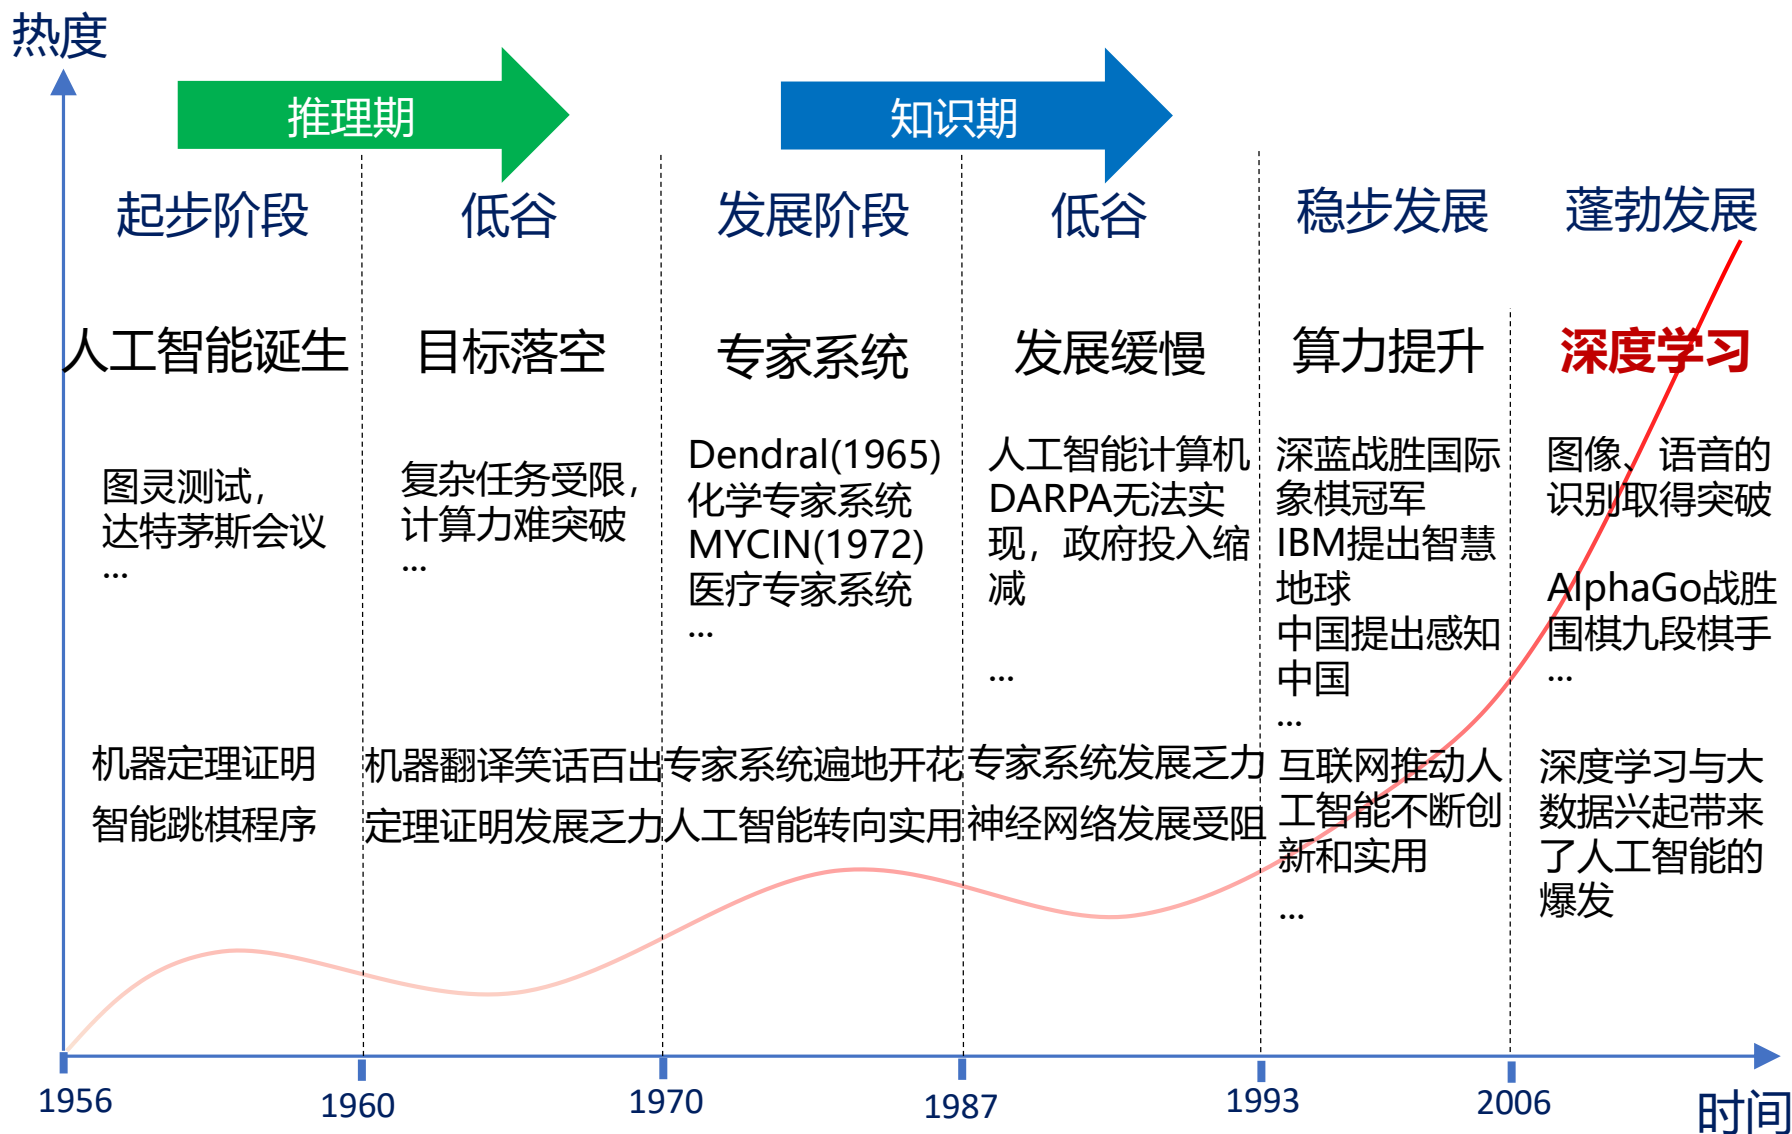
A Brief History



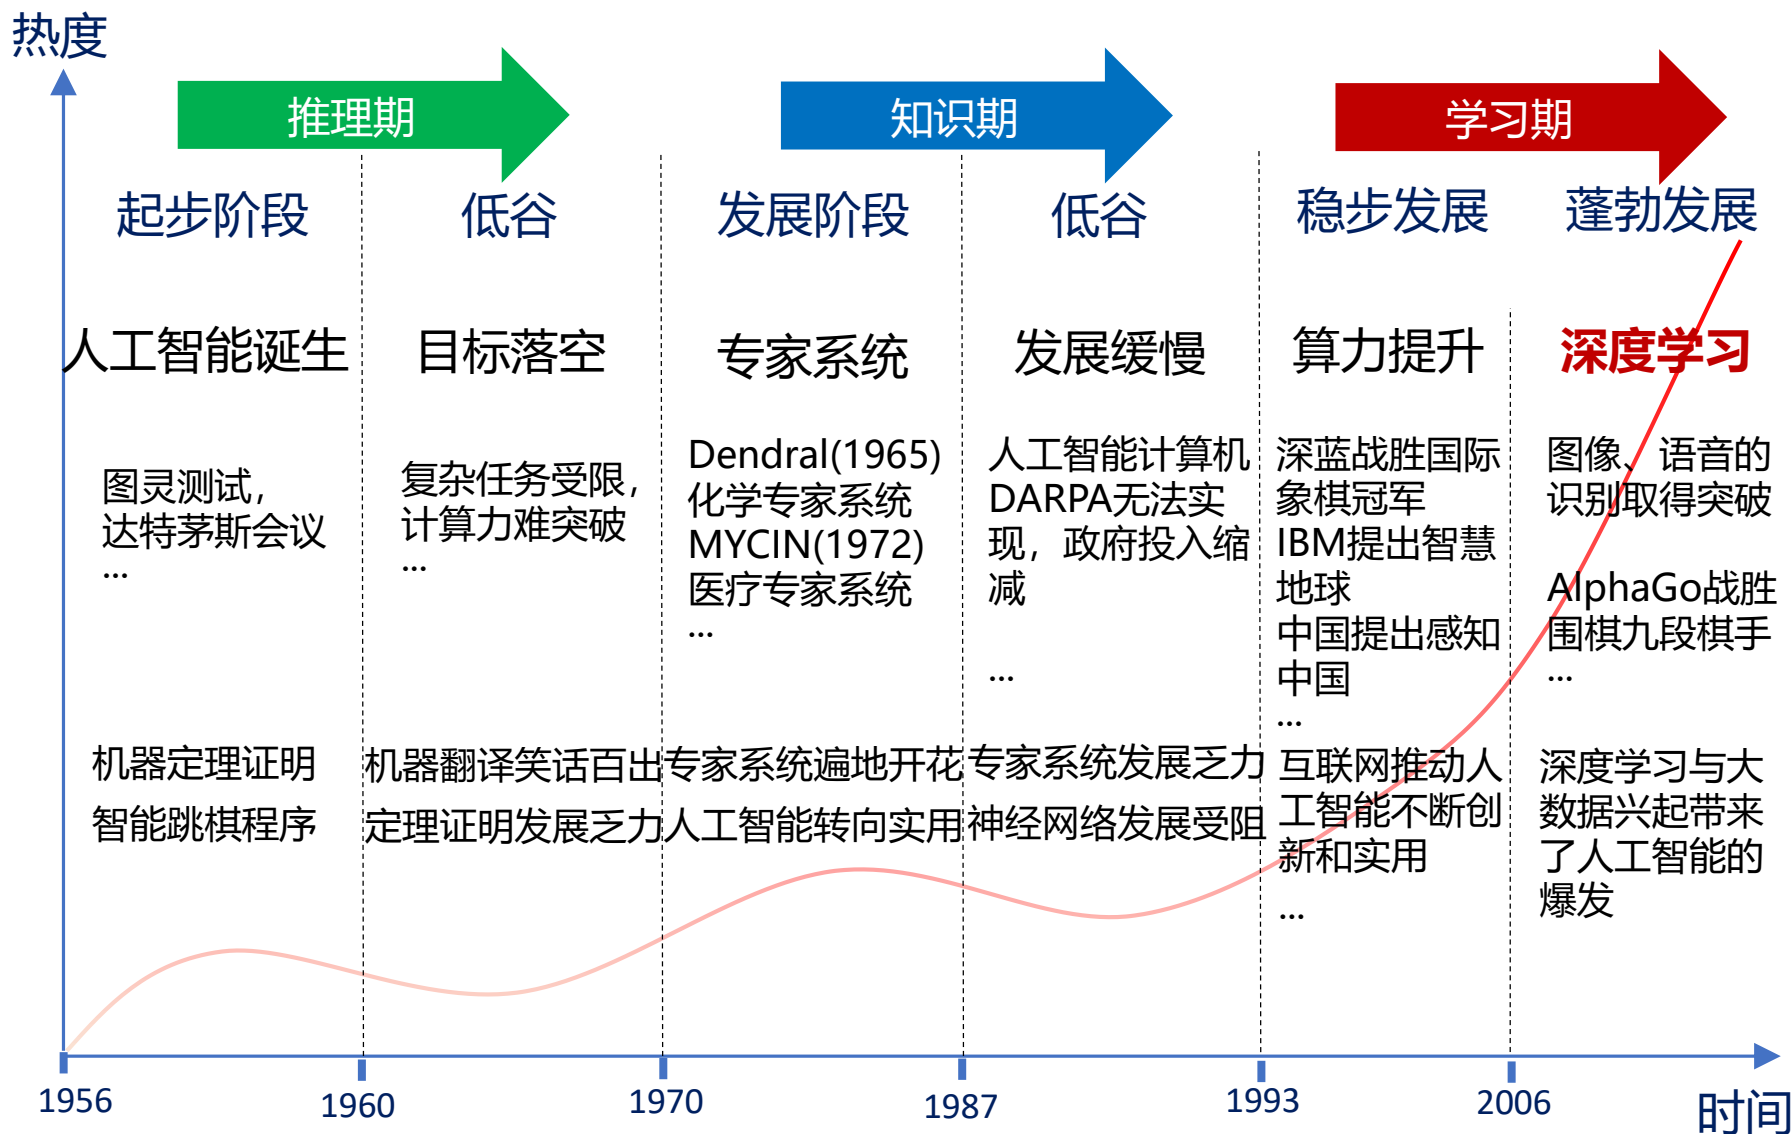
A Brief History



A Brief History



A Brief History



人工智能的发展历程



2018年图灵奖获得者从左到右：Yann LeCun, Geoff Hinton 和 Yoshua Bengio

人工智能的发展历程



迈克尔·乔丹是美国国家工程院院士、[科学院](#)院士和美国艺术和科学学院院士。同时，他还是美国科学进步协会的资深会员。现执教于加州大学伯克利分校，任电气工程与计算机科学系和统计系杰出教授，实时智能决策计算平台实验室（RISELab）共同主任、统计人工智能实验室（SAIL）主任、统计系系主任。2015年获得David E. Rumelhart奖，2009年获得ACM/AAAI的Allen Newell奖。

人工智能为我们带来了什么？

The future of Machine Learning and its Impact on Your Everyday Life



5 seconds

0:00 / 3:55 • Intro >

Scroll for details



<https://www.youtube.com/watch?v=NjVIQUE8Sgk>

人工智能发展现状

专用人工智能取得突破性进展，形成人工智能领域**单点突破**



蛋白质结构预测：AlphaFold2
达到原子级精度



围棋：AlphaGo Zero可以让
最高水平的人类棋手五子



大语言模型：ChatGPT可以在
几秒内就任何主题撰写文章



音频生成：Jukebox 全实时
生成各种风格的歌手录音



图像生成：Dall-E根据文本生
成图像，成为多模态AI开端

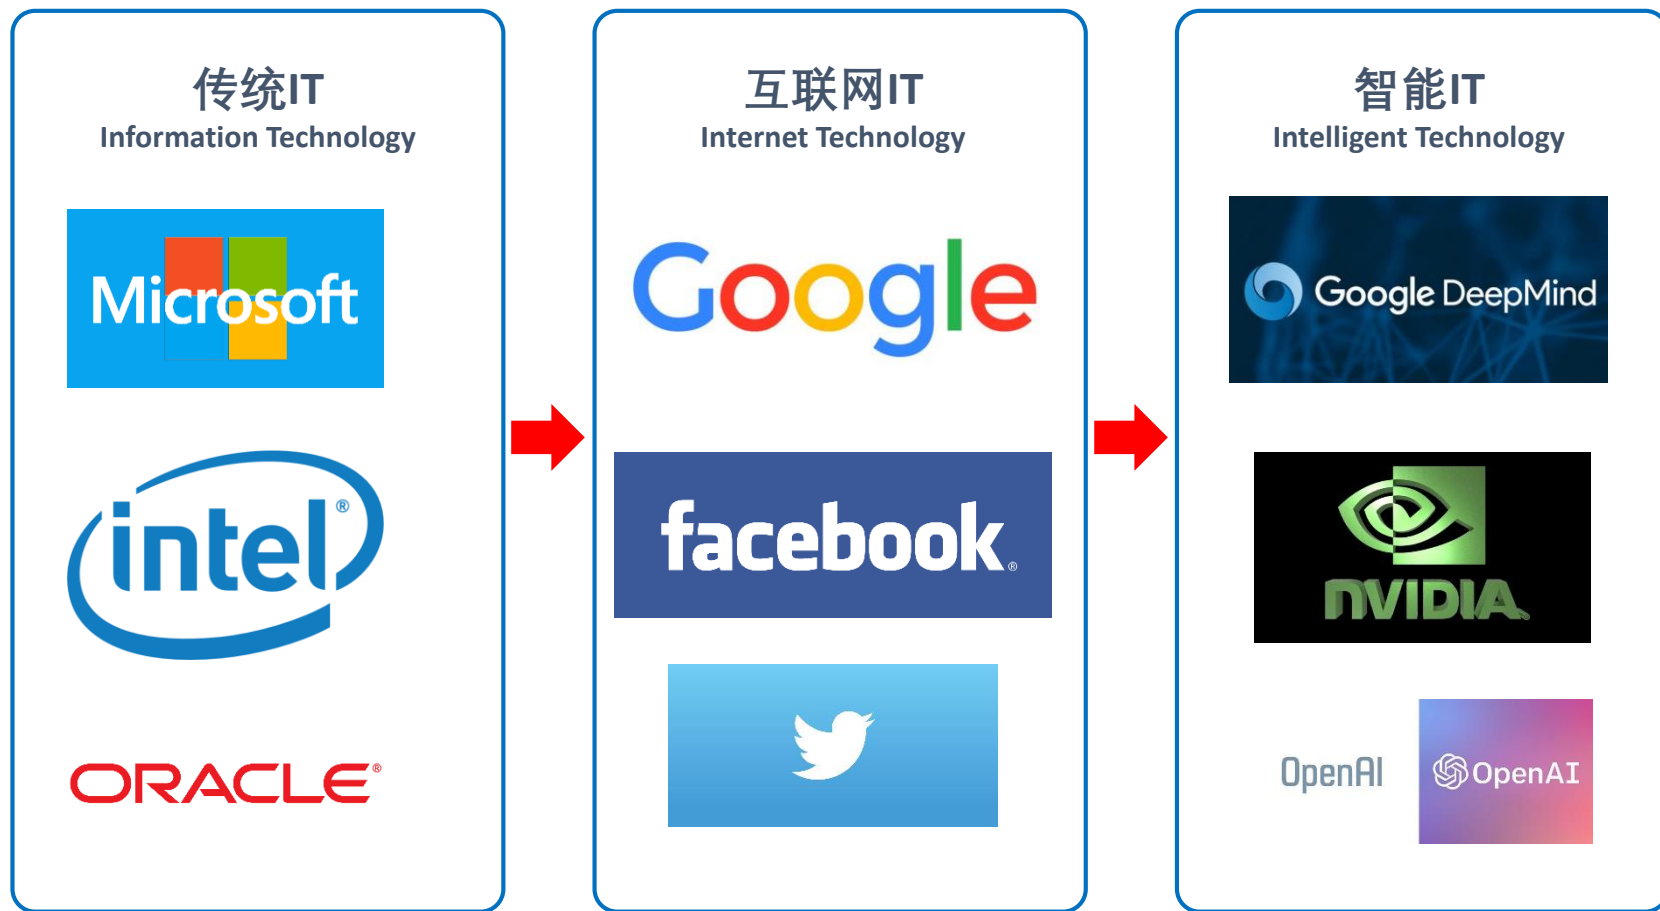


自动代码：AlphaCode 参加算
法比赛，成绩超过一半人类

人工智能发展现状



信息技术和产业的发展史就是IT巨头抢滩布局IT生态的更迭史



人工智能发展现状

“智能+”成为人工智能应用的创新范式

“智能+X”应用范式日趋成熟，AI向各行各业快速渗透融合进而重塑整个社会发展，这是人工智能驱动第四次技术革命的最主要表现形式。



智能+制造



智能+安防



智能+交通



智能+医疗

“智能”+应用范式日趋成熟




华为Mate10搭载人工智能芯片



Google、百度自动驾驶汽车

人工智能驱动第四次技术革命

Atari Games




<https://www.youtube.com/watch?v=V1eYniJ0Rnk>

AlphaGo

AlphaGO以4-1战胜人类围棋九段选手李世石

https://www.youtube.com/watch?v=8tq1C8spV_g


AlphaStar



AlphaStar以**5-0**分别战胜星际2的两位职业选手TLO和MaNa

<https://deepmind.com/blog/alphastar-mastering-real-time-strategy-game-starcraft-ii/>

Dog Fight



人工智能在美国国防部高级研究计划局举办的比赛中以**5-0**打败人类飞行员

<https://www.youtube.com/watch?v=NzdhlA2S35w>

无人机蜂群



浙江大学
控制科学与工程学院

浙江大学湖州研究院

全自主微型无人机集群

周鑫, 温向勇, 汪哲培, 高钰满, 李昊佳, 王钱浩, 杨天凯,
陆豪健, 曹燕军, 许超*, 高飞* (*通讯作者)

1. 系统总览



FASTLAB
Field Autonomous System & compuTing

<https://b23.tv/ezx2KoP>


自动驾驶




学习优化——自动驾驶

基于 ML 优化密集车流场景的自动驾驶变道轨迹

◆ 研究问题



密集车流变道轨迹优化



杨睿

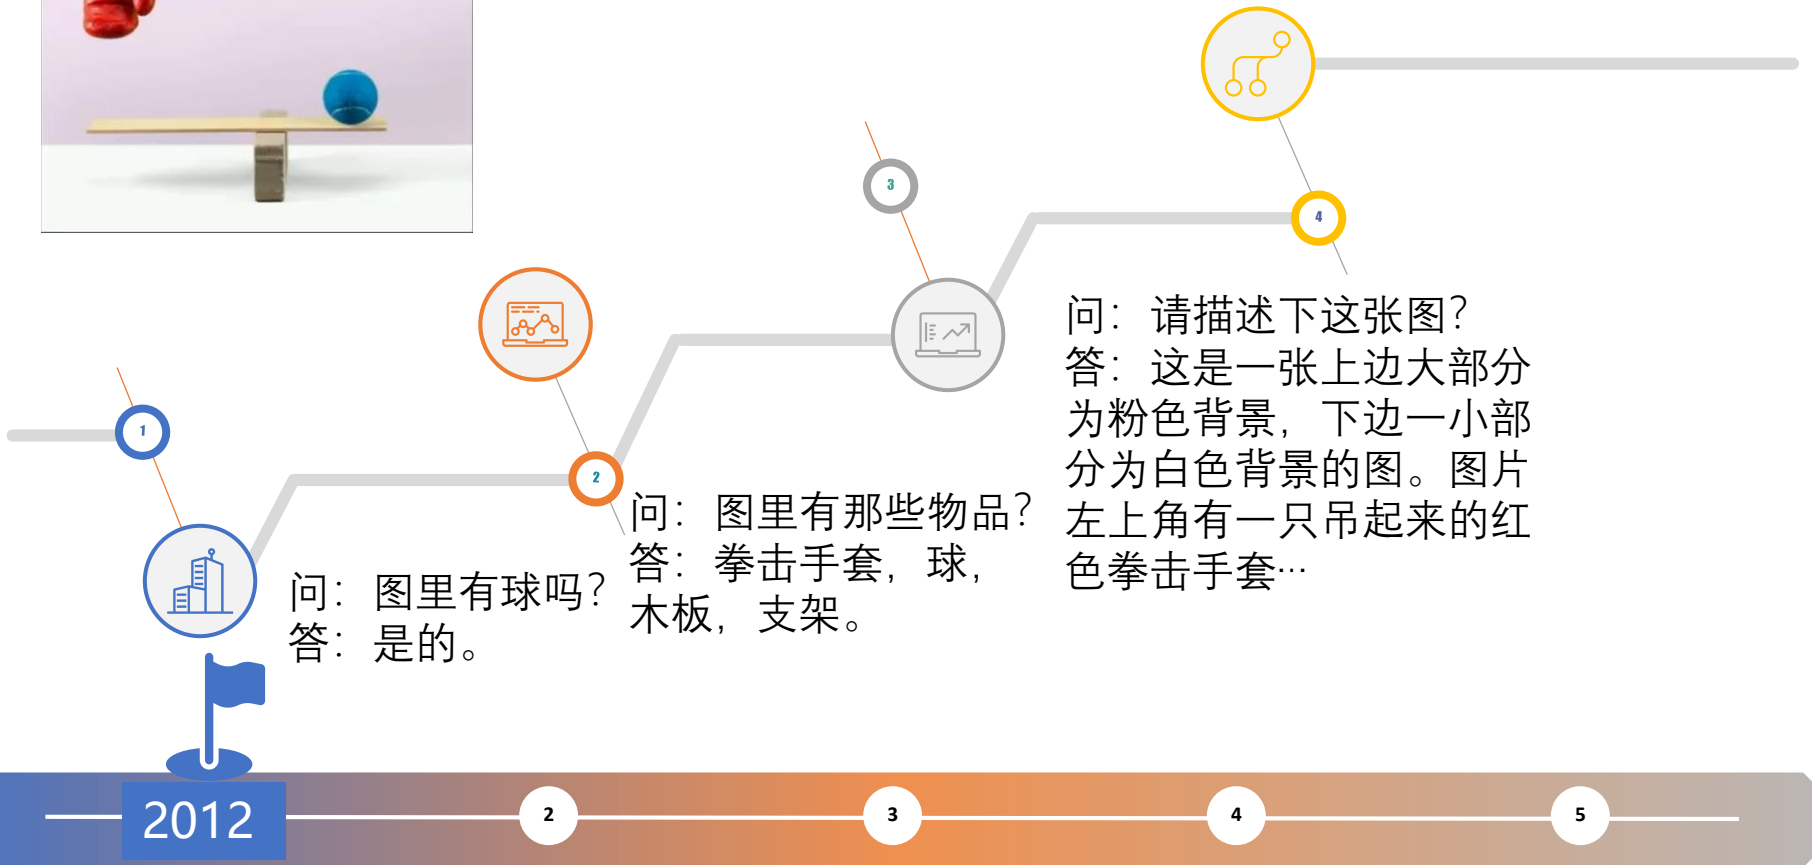
◆ 科研小组介绍

- 团队成员：杨睿（团队负责人，博士三年级，已发表A会论文1篇，一篇TPAMI在投的工作）、刘春阳（滴滴Voyager规控部门负责人，已发表论文10余篇）
- 影响力：和滴滴自动驾驶Voyager部门深度合作。其课题——密集车流等复杂人机交互场景的自动驾驶运动规划是业界重点关注的前沿领域

大模型 – GPT 4 (OpenAI)



问：手套掉下来了会怎么样？
答：它会掉到木板上，并且球被弹飞。



大模型 - 书生 (商汤科技)

Q: 如何前往北大教材中心?

A: 驶离停车场后, 右转进入北四环, 沿北四环行驶至中关村北大街, 左转后直行至北京大学。



Q: 我需要停车。

A: 左前方两个空车位, 右前方三个空车位。



Q: 前方是否可以保持当前速度直行?

A: 前方学校请减速。



Q: 如何驶离停车场?

A: 前方左转



Q: 此刻可以直行吗?

A: 需等待5秒才能直行通过。



Q: 前方是否可以保持当前速度直行?

A: 前方隧道限速50km/h。



Q: 前方能否靠右行驶?

A: 不能, 右侧为公交专用道。

大模型 - 但是...

New chat

UR 刘邦是如何打败朱元璋的?

OpenAI logo 刘邦是中国历史上著名的军事家和政治家，他是汉朝的开国皇帝。朱元璋是明朝的最后一位皇帝。

刘邦是通过长期的军事斗争和政治策略最终打败朱元璋的。刘邦曾经是汉朝的大将军，他率领军队与许多地方势力作战，稳定了国家的局面。他后来组织了叛乱，并在汉朝的最后几年夺取了皇位，成为了汉朝的开国皇帝。

因此，刘邦是通过长期的军事斗争和政治策略打败朱元璋的，这是中国历史上一个重要的转折点。

ChatGPT

画一个胸有成竹的男人



好的，根据你的需求，我为你创作了一幅画作。

我的作画技能还在不断进化中，暂时还不支持对画作的修改和解释。

如果重要继续让我为你作画 请完整描

知乎 @游戏开发极客

文心一言

现有人工智能的主要局限




人工智能总体发展水平仍处于起步阶段，
目前人工智能距人类的整体、通用智能水平相去甚远。



会计算不会“算计”：
人工智能系统可谓有智无心，更无谋



有专才无通才：
会下围棋的AlphaGo不会下象棋



Future



<https://www.youtube.com/watch?v=Wx7RCJvoCMc>

习近平总书记：推动我国新一代人工智能健康发展



习近平总书记**高度重视**人工智能领域的建设和发展

习近平总书记在**党的二十大报告**中指出：

- 推动战略性新兴产业融合集群发展，构建新一代信息技术、**人工智能**、生物技术、……等一批新的增长引擎

习近平总书记在**中央政治局第九次集体学习**时的重要讲话：

- 要深刻认识加快发展新一代**人工智能**的重大意义，……，促进其同经济社会发展深度融合，推动我国新一代**人工智能**健康发展

国家战略：国务院、科技部等协力部署人工智能建设



人工智能的迅速发展将深刻改变人类社会生活、改变世界。为抢抓人工智能发展的重大战略机遇，构筑我国人工智能发展的先发优势，加快建设创新型国家和世界科技强国，按照党中央、国务院部署要求，制定本规划。

——国务院《新一代人工智能发展规划》

科技部关于发布科技创新2030—“新一代人工智能”

重大项目2018年度项目申报指南的通知

国科发资〔2018〕208号

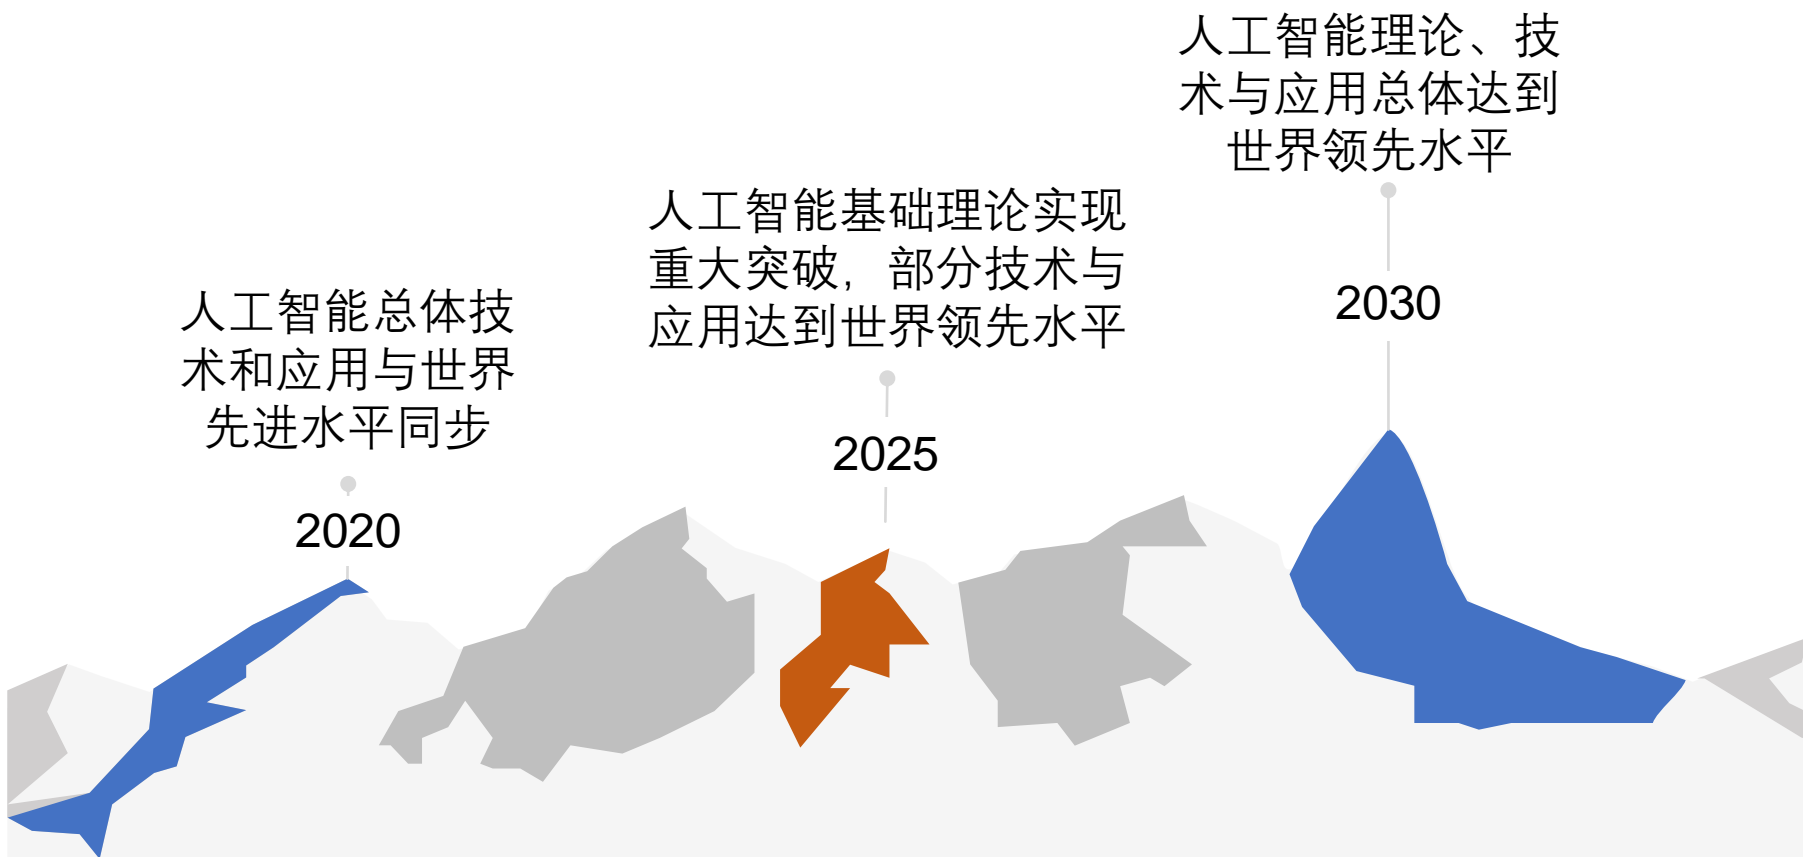
各省、自治区、直辖市及计划单列市科技厅（委、局），新疆生产建设兵团科技局，国务院各有关部门科技主管司局，各有关单位：

为落实国务院印发的《新一代人工智能发展规划》的总体部署，现根据《新一代人工智能重大科技项目实施方案》启动实施科技创新2030—“新一代人工智能”重大项目。按照《关于深化中央财政科技计划（专项、基金等）管理改革的方案》（国发〔2014〕64号）的相关要求，将2018年度项目申报指南予以公布。请按照指南要求组织项目申报工作。有关事项通知如下。

政策支持



国家战略：三步走战略



政策支持



国际舞台：世界各国竞相发布和推进积极政策

中国



国务院《新一代人工智能发展规划》：

对新形势新需求，必须主动求变应变，**牢牢把握人工智能发展的重大历史机遇**，紧扣发展、研判大势、主动谋划、把握方向、抢占先机，**引领世界人工智能发展新潮流**，服务经济社会发展和支撑国家安全，带动国家竞争力整体跃升和跨越式发展。

美国



2022年6月，美国国防部发布**人工智能技术战略规划文件**，布局发展应用于军事的人工智能技术。

欧盟

欧盟正在推进《**人工智能法案**》的修订和生效。该《法案》将为人工智能技术提供引导和规范。

人工智能的重要性

人工智能进入社会的方方面面



人工智能上升至**国家战略**

国务院关于印发
新一代人工智能发展规划的通知
国发〔2017〕35号

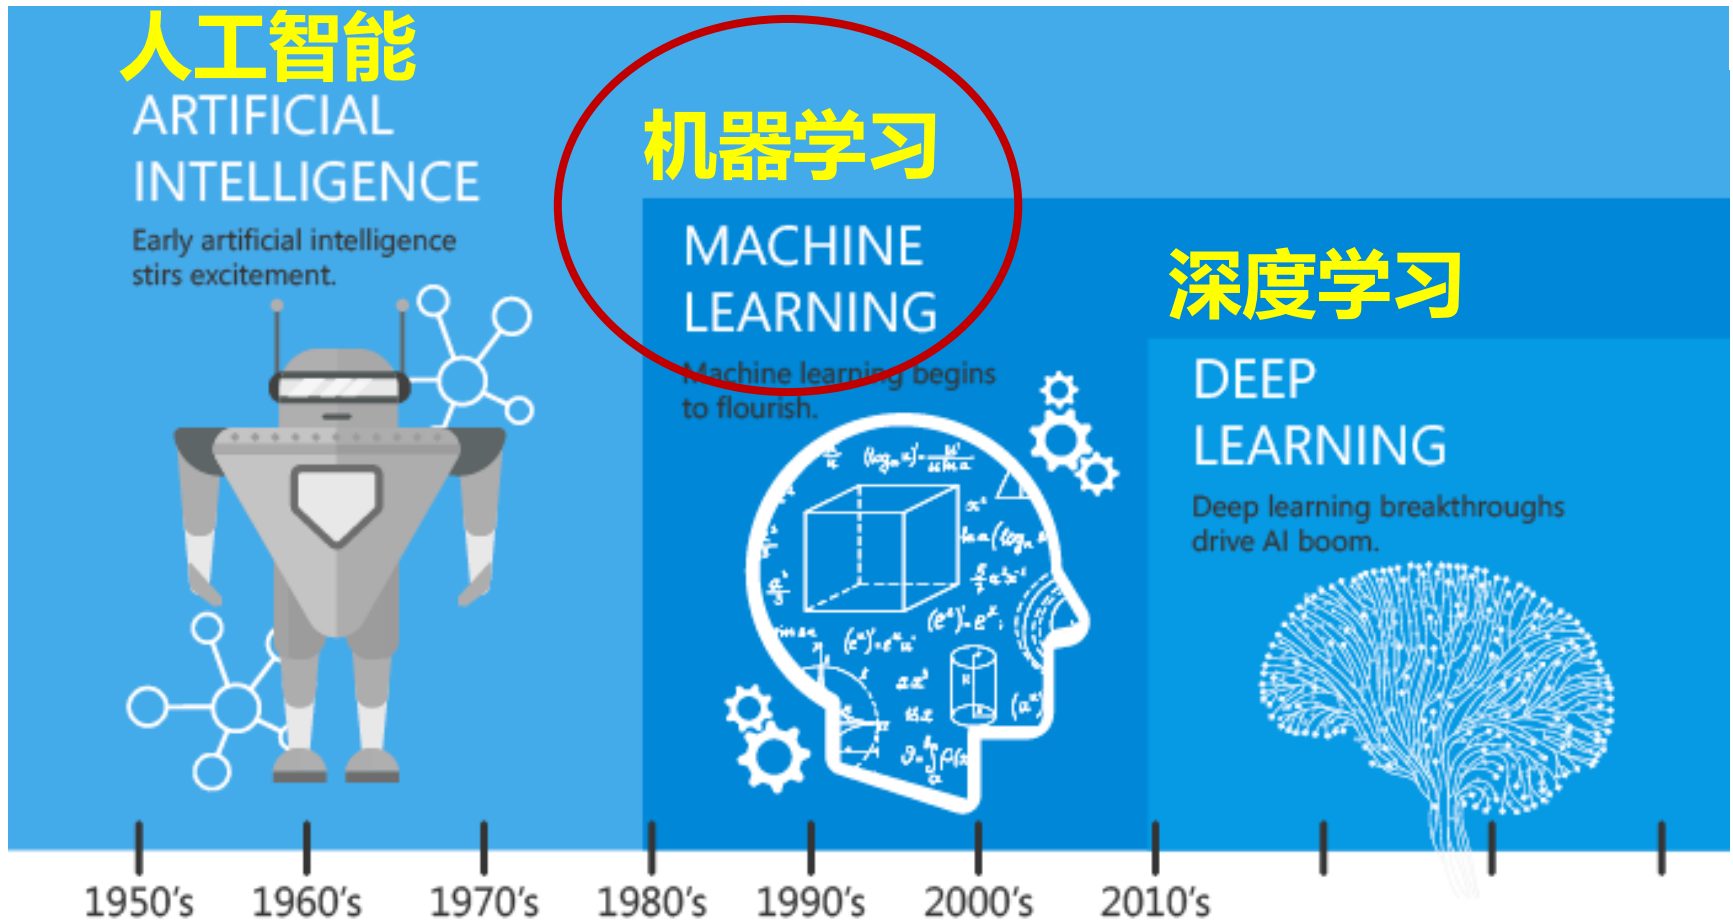
各省、自治区、直辖市人民政府，国务院各部委、各直属机构：
现将《新一代人工智能发展规划》印发给你们，请认真贯彻执行。

国务院
2017年7月8日

新一代人工智能发展规划

国发〔2017〕35号





Since an early flush of optimism in the 1950's, smaller subsets of artificial intelligence - first machine learning, then deep learning, a subset of machine learning - have created ever larger disruptions.


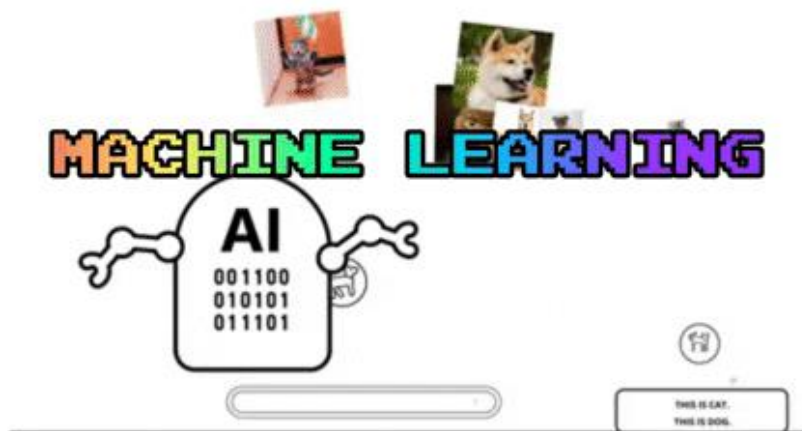
What is Machine Learning?



A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P if its performance at tasks in T , as measured by P , improves with experience E .

--- Tom Mitchell (1997)

GIGABYTE™



深度学习示例

Why you should learn ML?

- This is an era of ML



Self-Driving

Intel: 4000G/车/天
Waymo: 1000万公里驾驶里程



Health Care

基因组: 30亿碱基对/人
全球2015年医疗数据: 8.6万亿GB



Surveillance

全国监控摄像头: 超3000万
每日监控视频: 超7亿小时



Customer Service

中国移动: 日客服进线240万例
京东: 日客服进线75万例

- Still many opportunities for you to explore



Social networks



Recommender Systems



Medical Image



Intelligent Networks/Cities

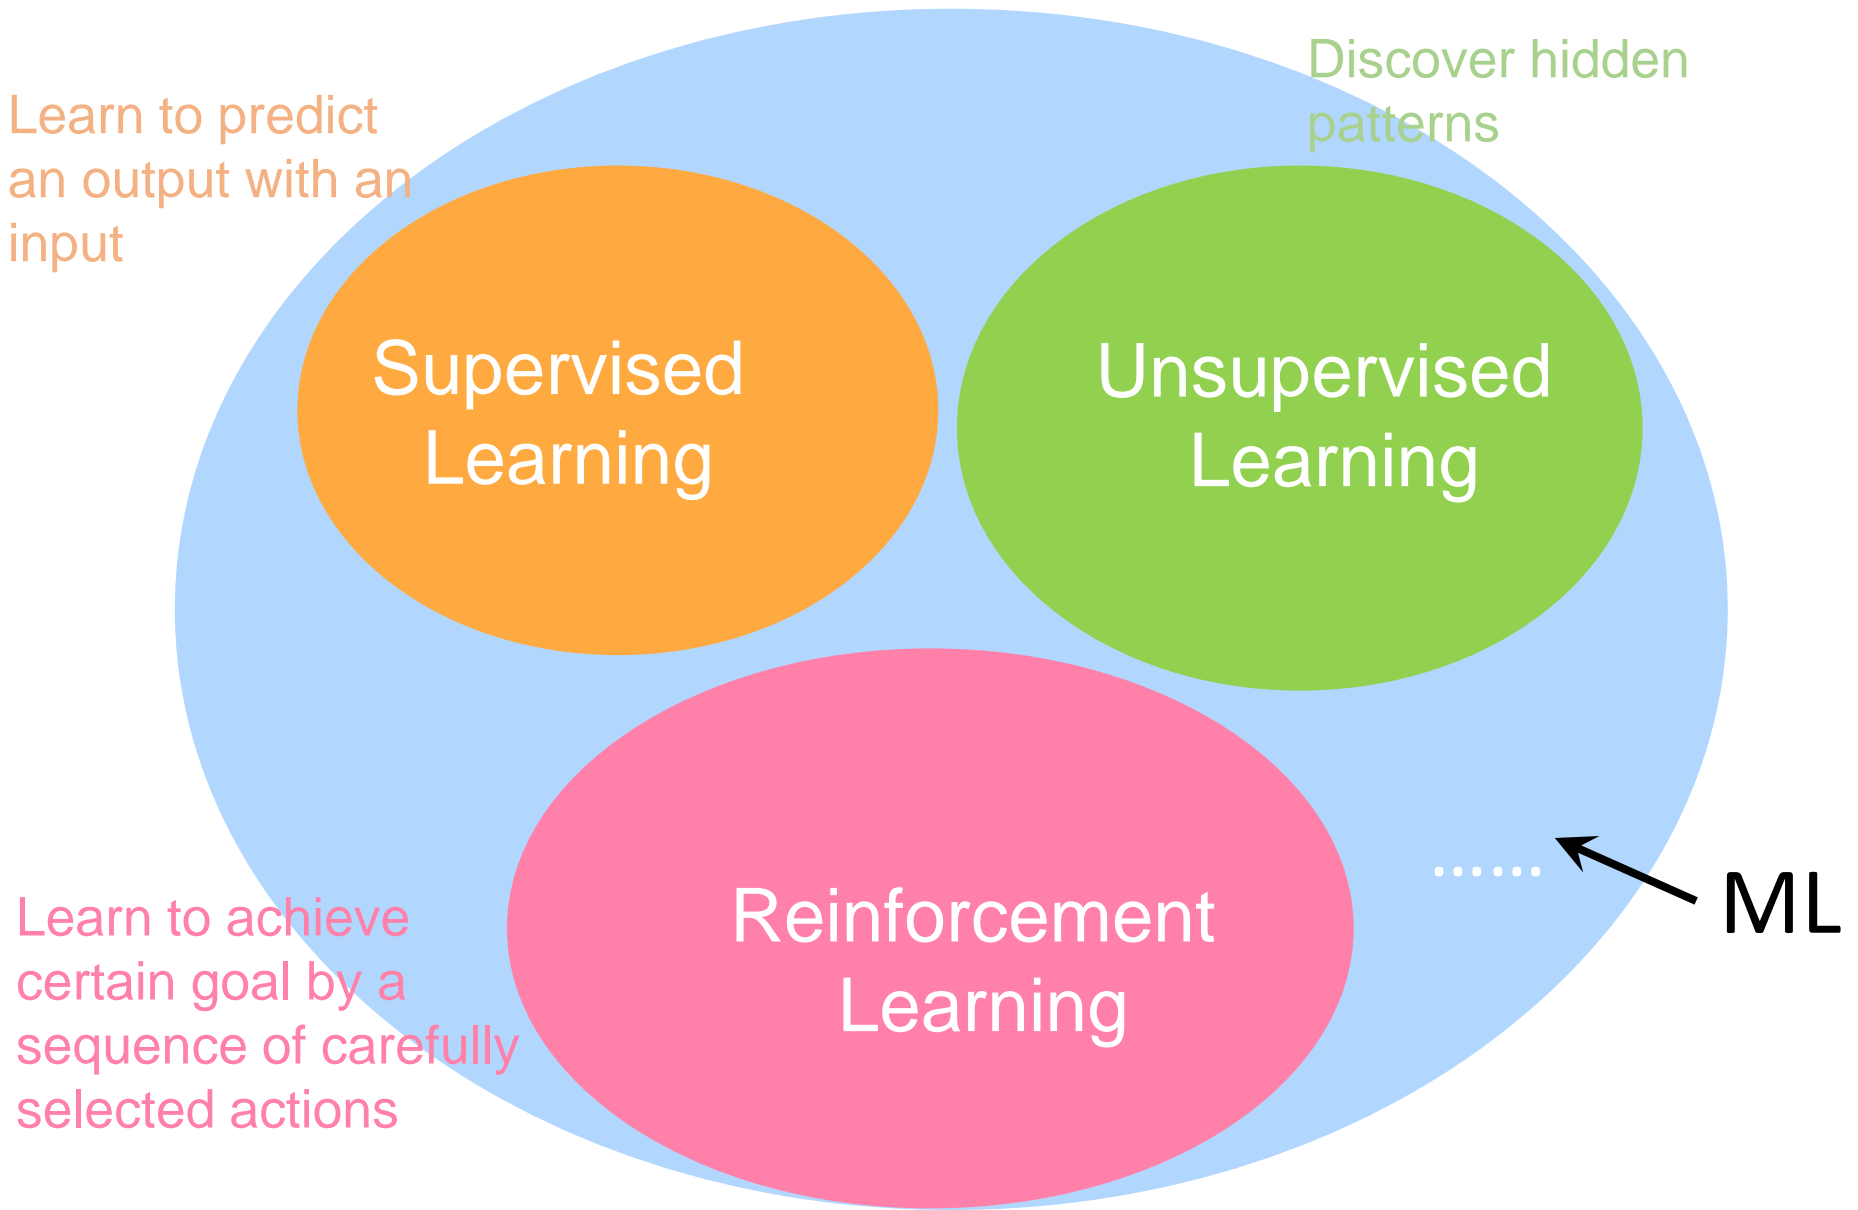


- Plenty of well-paid openings

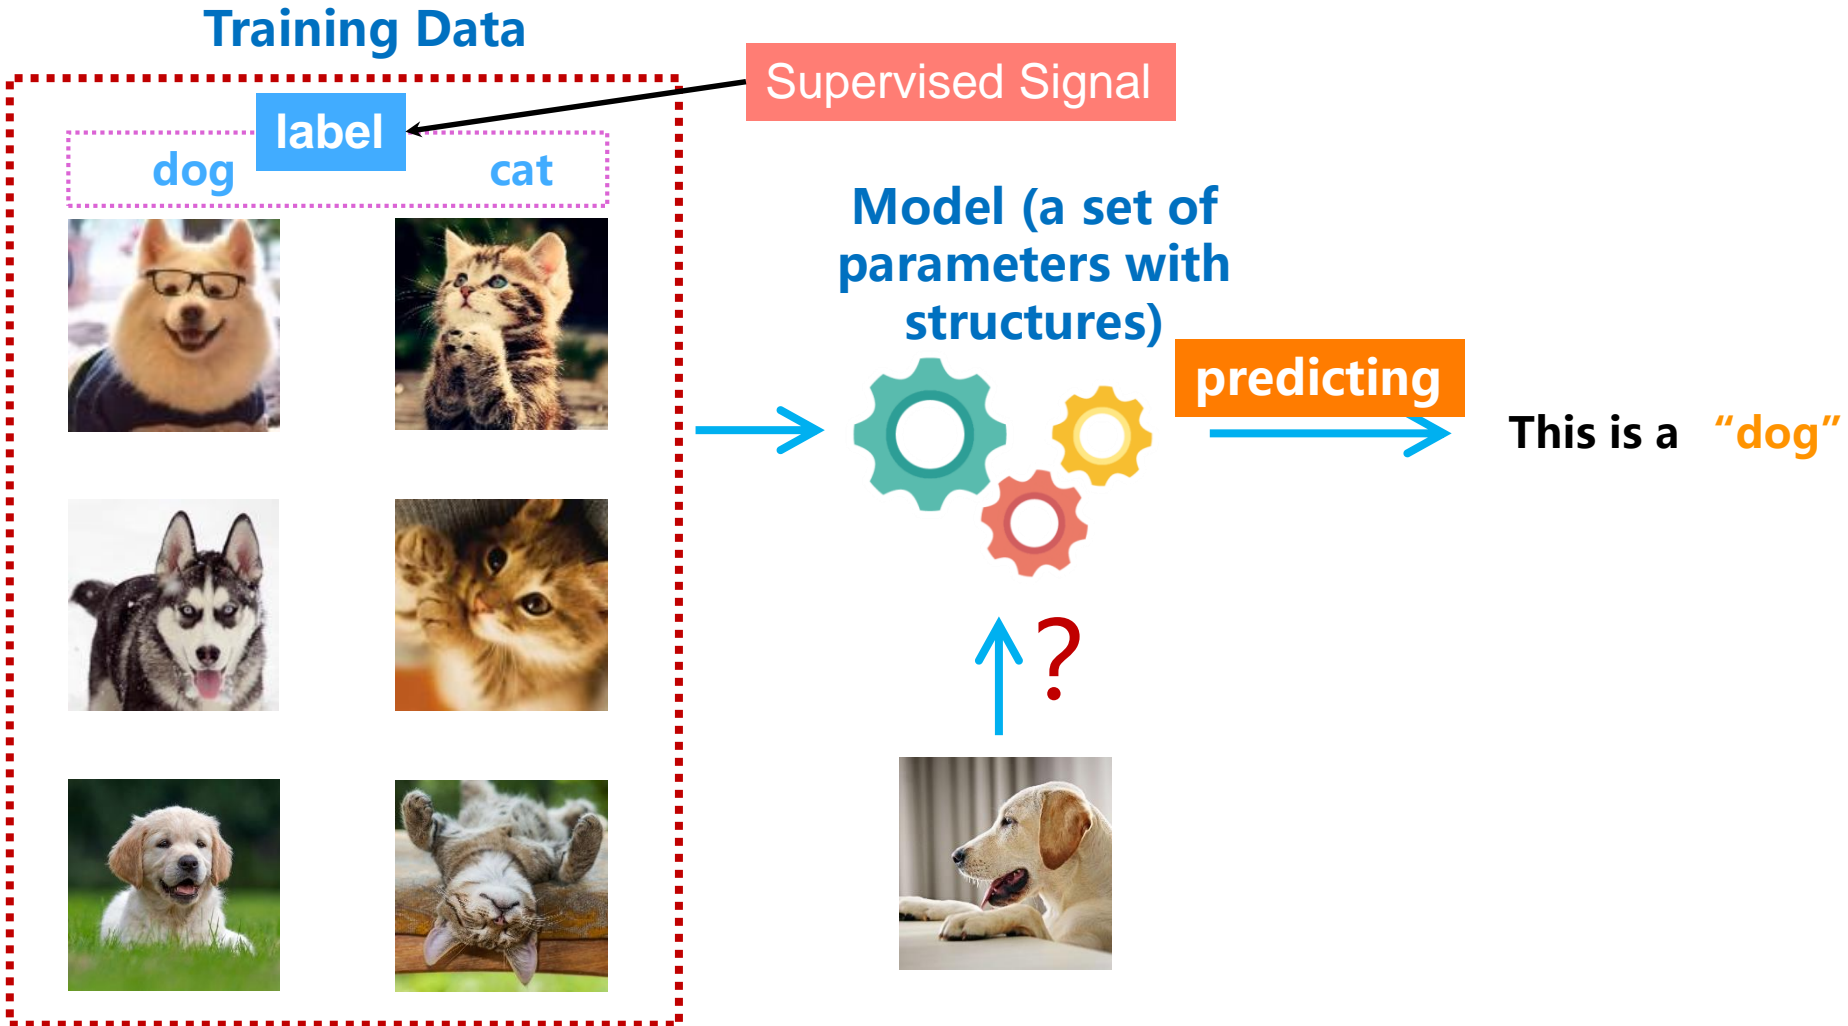


 Baidu Traffic	 Alibaba Industry Chain	 Tencent User
--------------------------	-----------------------------------	-------------------------

A Bird's eye view of ML



Supervised Learning



- Regression : The output is a real number (vector)
- Classification : The output is a class label.

We are indeed looking for a mapping (function).

- Image classification




$f(\cdot)$ → “Dog”

Supervised Learning

We are indeed looking for a mapping (function).

- Autonomous driving



$f(\cdot)$ → “slow down”

We are indeed looking for a mapping (function).


- Speech recognition




$f(\cdot)$ → “How are you doing”

We are indeed looking for a mapping (function).


- Medical diagnosis



$f(\cdot)$ → “normal”



$f(\cdot)$ → “mild cognitive impairment”




$f(\cdot)$ → “Alzheimer Disease”

Supervised Learning – Hypothesis Set

Regression

Linear Models







<https://realpython.com/linear-regression-in-python/>

Supervised Learning – Hypothesis Set



Regression

Linear Models




Polynomials

Supervised Learning – Hypothesis Set



Classification

Linear Models

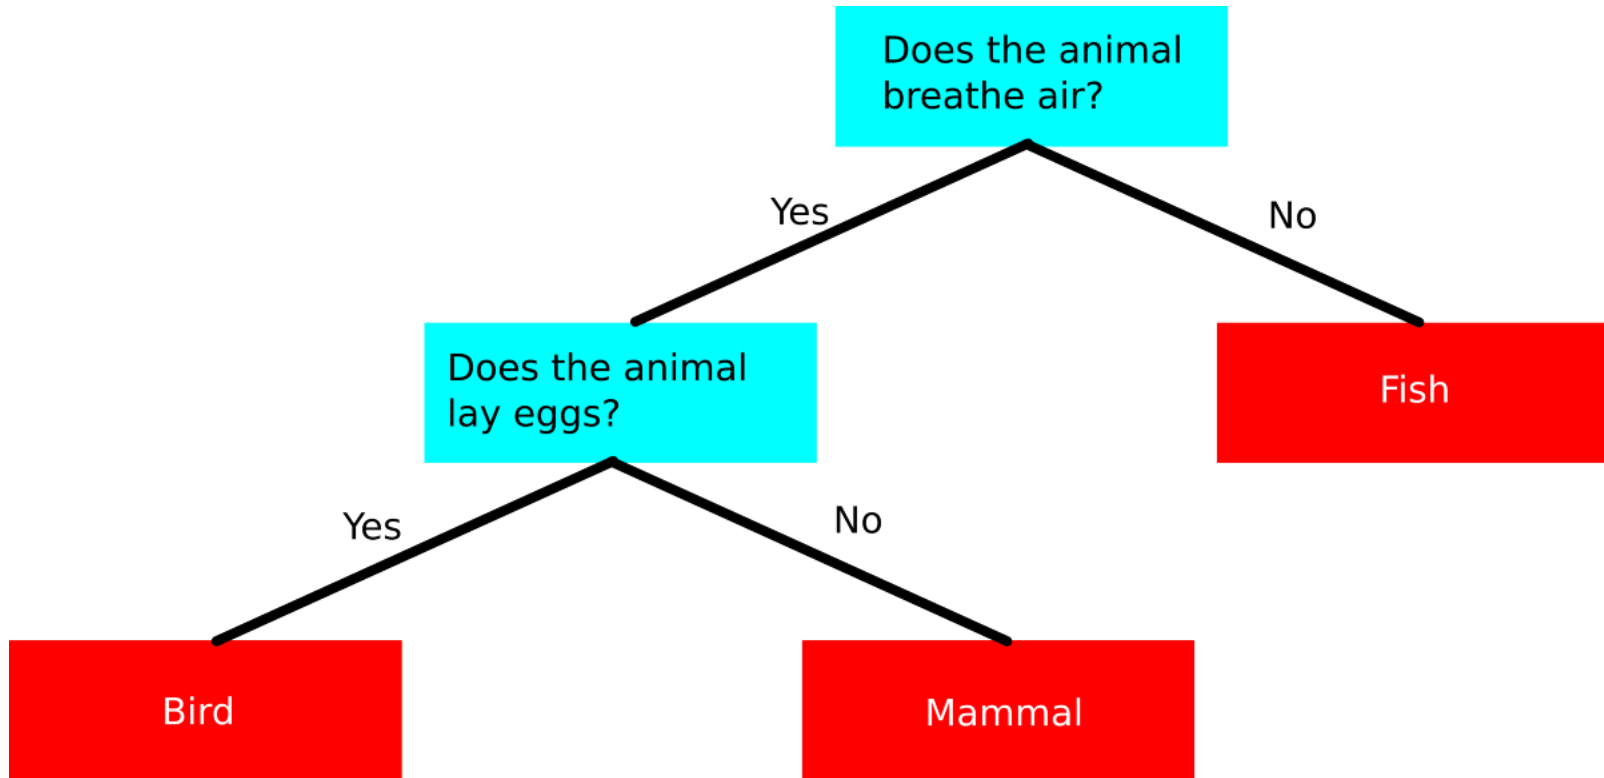


<https://blogs.oracle.com/bigdata/machine-learning-techniques>

Supervised Learning – Hypothesis Set

Classification

Decision Tree

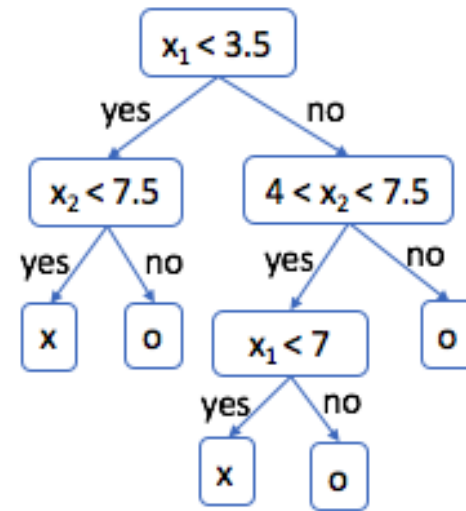



<https://towardsdatascience.com/a-beginners-guide-to-decision-tree-classification-6d3209353ea>

Supervised Learning – Hypothesis Set

Classification

Decision Tree




<https://www.jeremyjordan.me/decision-trees/>

Supervised Learning – Hypothesis Set

Classification

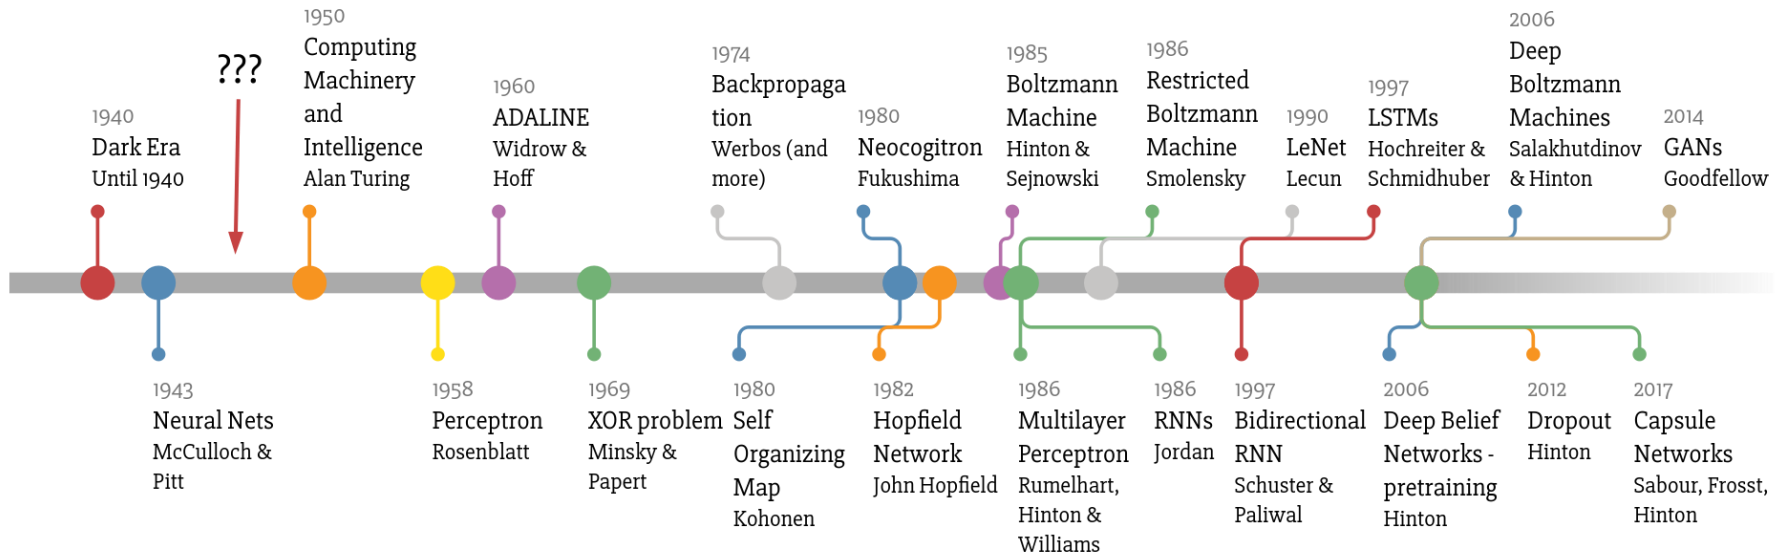
Neural Network



Supervised Learning – Hypothesis Set



Deep Learning Timeline



Made by Favio Vázquez


<https://towardsdatascience.com/a-weird-introduction-to-deep-learning-7828803693b0>

Supervised Learning – Hypothesis Set




8 layers

16.4%



AlexNet (2012)



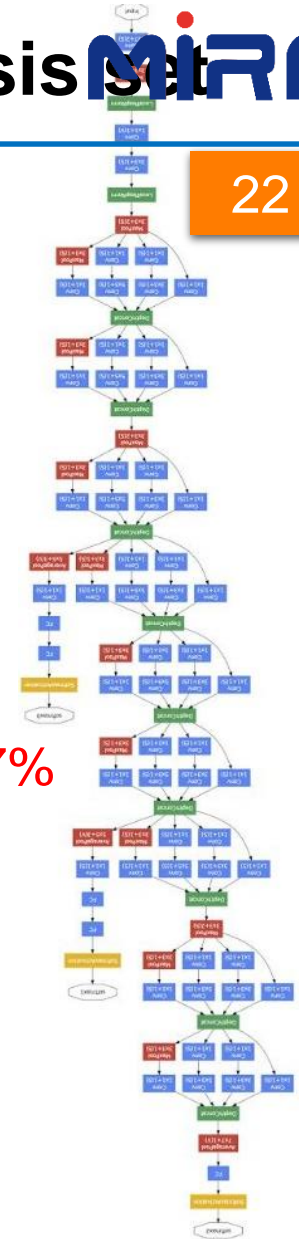
19 layers

7.3%

VGG (2014)


22 layers

6.7%



GoogleNet (2014)

Supervised Learning – Hypothesis Set



16.4%

AlexNet
(2012)

7.3%

VGG
(2014)


6.7%

GoogleNet
(2014)

3.57%

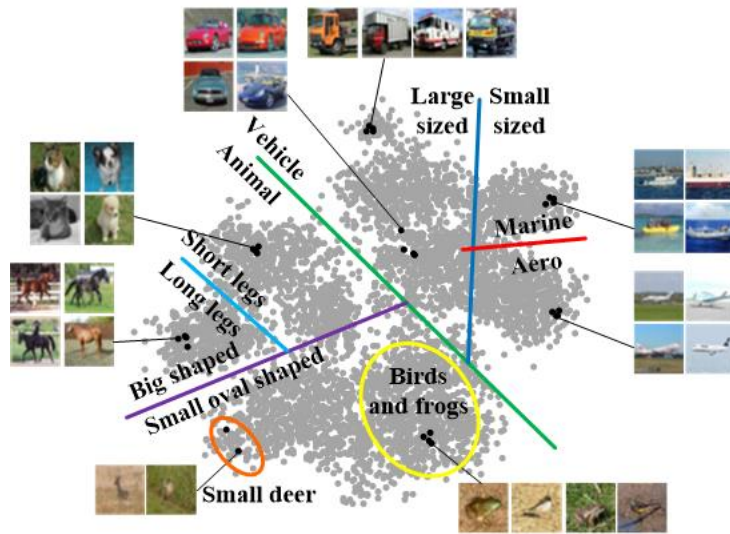
Residual Net
(2015)

Unsupervised Learning

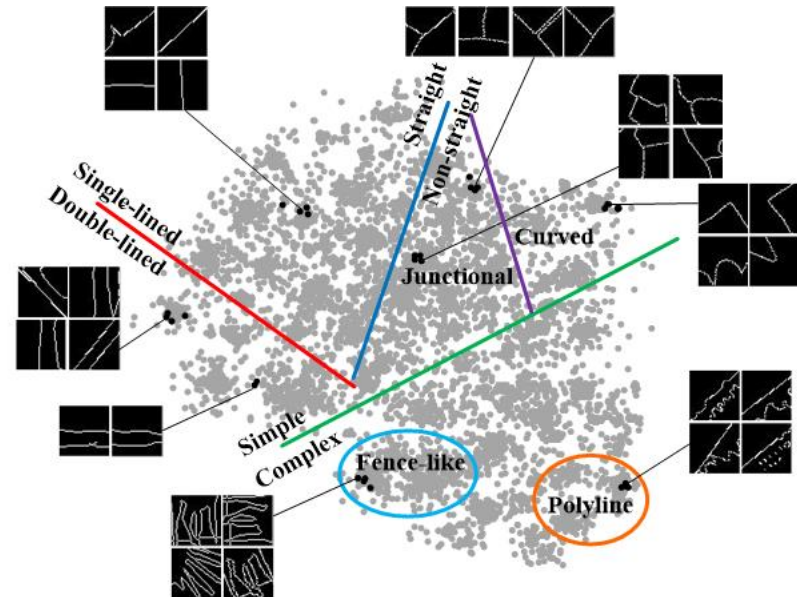


No labels available

Unsupervised Learning




Unsupervised attributes of natural images



Unsupervised attributes of binary contour patches


Unsupervised Learning



<https://towardsdatascience.com/google-news-and-leo-tolstoy-visualizing-word2vec-word-embeddings-with-t-sne-11558d8bd4d>

Self-supervised Learning

100 year old pictures...



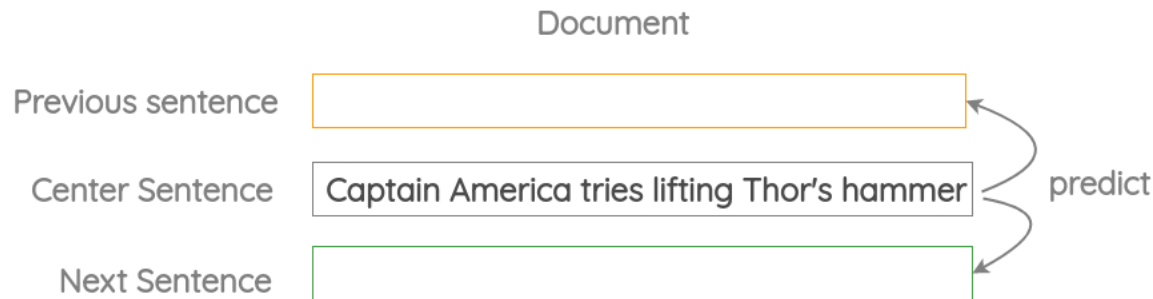
<http://izuka.cs.tsukuba.ac.jp/projects/colorization/en/>

Self-supervised Learning

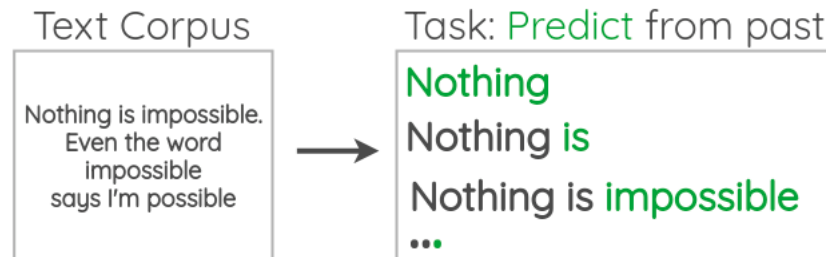
- Center word prediction

A quick brown fox jumps over the lazy dog

- Neighbor sentence prediction



- Neighbor sentence prediction



Reinforcement Learning



Deep Reinforcement Learning: AI = RL + DL



David Silver

Reinforcement Learning



Reinforcement Learning

$$\pi(A_t|S_t) \rightarrow A_t$$

Environment

Agent



$$P(S_{t+1}, R_{t+1}|S_t, A_t)$$



$$S_{t+1}, R_{t+1}$$

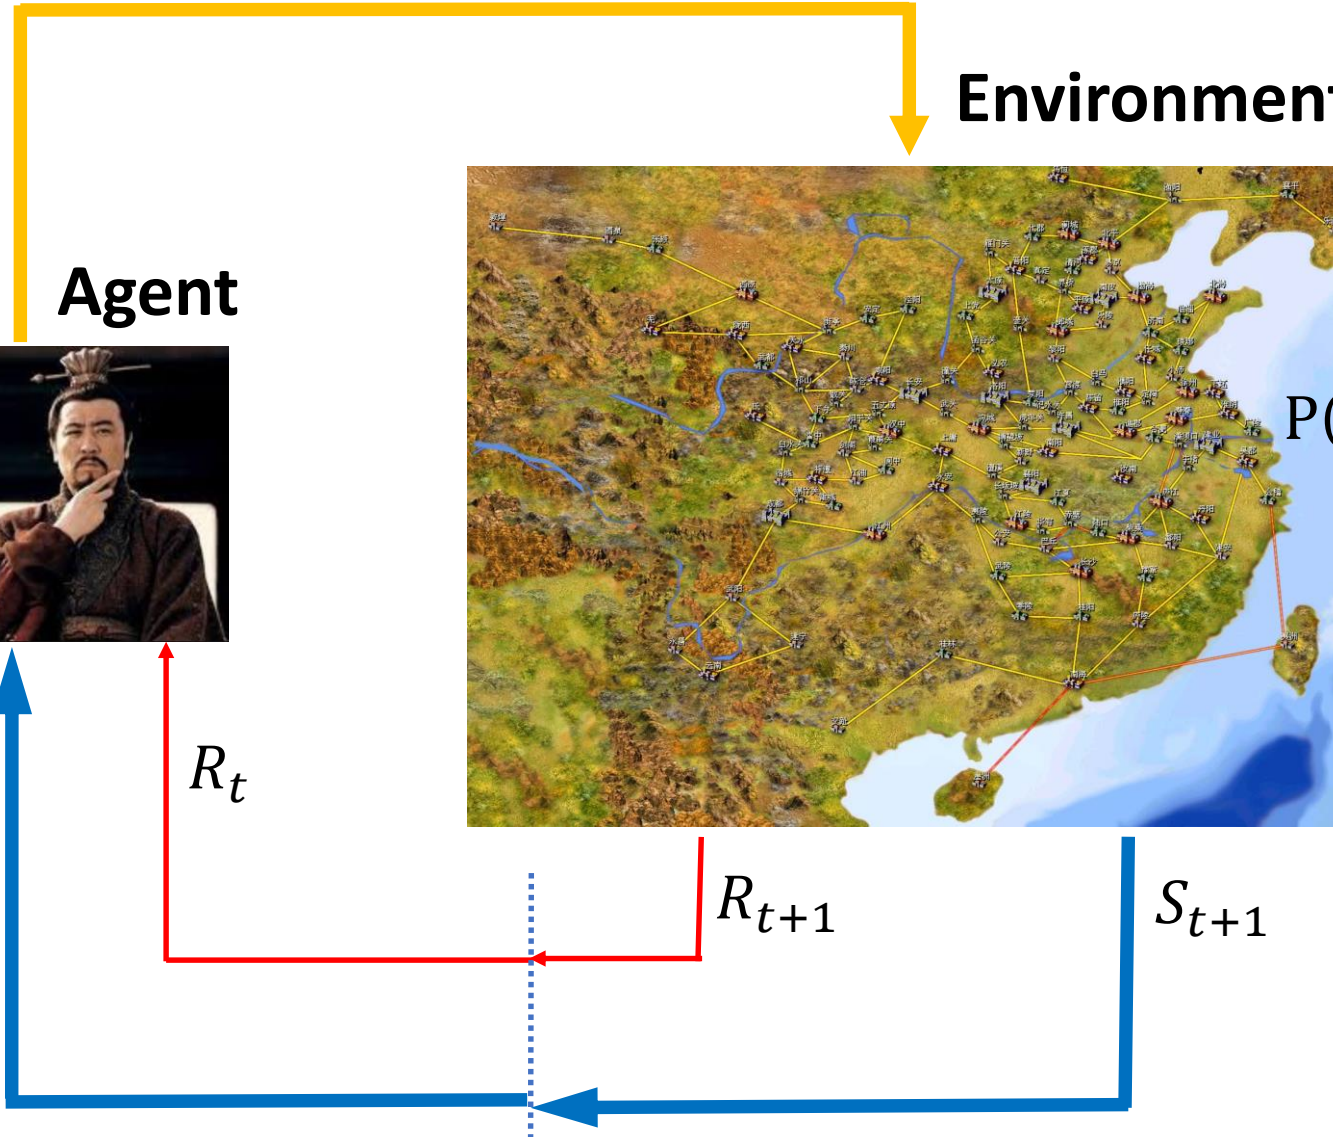


S_t

R_t

R_{t+1}

S_{t+1}






Controlling the robotic arm to play Rubik's Cube

[1] Peng, Xue Bin, et al. "Sim-to-real transfer of robotic control with dynamics randomization." 2018 IEEE international conference on robotics and automation (ICRA). IEEE, 2018.

Achievements – Industry Control



Task: controlling **google data center cooling** to minimize the energy consumption while satisfying safety constraints



Over nine months, the AI control system performance increases from a 12 percent improvement to around a **30** percent improvement.

- Representation
 - What is the most appropriate representation of the data regarding the task on hand?
- Generalization
 - Can we do well on the unseen data?
- Complexity
 - How much time and space?
- Efficiency
 - How many samples needed to output an accurate hypothesis?
- Applications

Questions

