

人工智能技术概述

Introduction to Artificial Intelligence (AI)

王 跃 明

**浙江大学计算机学院
浙江大学求是高等研究院**

2018年10月16日

Some slides are “borrowed” from the Stanford CS221 Course

基本情况

研究领域：人工智能

研究方向：数据挖掘，模式识别，脑机交互

- 2016.12 ~ 至今 ： **浙江大学，求是&计算机学院，教授/博导**
 - 2010.06 ~ 2016.12： **浙江大学，求是&计算机学院，副教授/博导**
 - 2007.12 ~ 2010.06： **香港中文大学，Research Associate**
 - 2007.07 ~ 2007.12： **微软亚洲研究院，视觉计算组，访问学者**
 - 2002.09 ~ 2007.06： **浙江大学，计算机学院，博士**
-

荣誉奖励

- 国际脑机接口**创新研究提名奖**，2016
- 中国高等院校**十大科技进展**，2016
- 中国人工智能学会**吴文俊人工智能科技创新一等奖**，2016
- 林百欣高科技奖，2014
- 中国人工智能学会优秀专委会，2016
- 浙江大学教学成果奖，2012
- 浙江大学求是青年学者，2011



吴文俊人工智能科学技术奖

激励创新 成就荣耀

拟授奖等级	主要完成人	成果名称
一等奖	吴朝晖创新团队	脑机融合的混合智能理论与方法

先来几段视频 ...

Boston Dynamics: Atlas & Big Dog



- 2016年3月，Alphabet已经将旗下机器人公司Boston Dynamics摆上出售的货架了。
 - 2013年，谷歌收购了Boston Dynamics。
-

智能图像增强&理解



Image Inpainting, M. Bertalmío

智能图像增强&理解



智能图像增强&理解




智能图像增强&理解

去雾, *Dehaze, Microsoft Research Asia, K. He.*



智能图像增强&理解

去雾, *Dehaze, Microsoft Research Asia, K. He.*




智能图像增强&理解

去雾, *Dehaze, Microsoft Research Asia, K. He.*



人脸检测以及识别



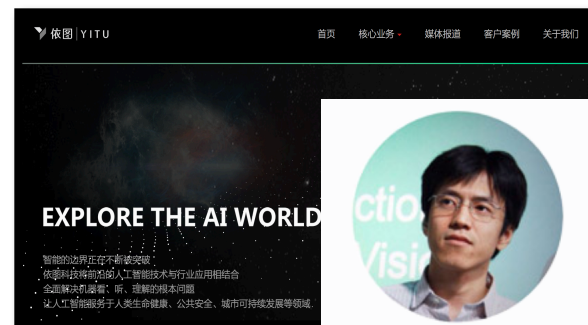
人脸检测以及识别



商汤




旷视




依图

人脸检测以及识别

试试看！



行人检测



语音识别

Siri

1. 我想给浙图打个电话
2. 58除以29等于几？
3. 0除以0等于几？
4. 借我点钱可以吗？
5. 你做我女朋友，行吗？
6. . . .



语音识别

Google Duplex给美发店现场打call，使用了Deepmind的Wavenet技术，使机器的声音与真人基本无异，语气更是令全场惊叹并发笑，顺利与美发店进行了多轮沟通，未现破绽。



下面进入正题 ...

试图回答几个问题


- 什么是AI及AI的发展历史
- 是什么让AI so HOT
- 我国AI的发展规划

周围都发生了什么？



Microsoft creates AI that can read a document and answer questions about it as well as a person

January 15, 2018 | [Allison Linn](#)



Microsoft researchers achieve new conversational speech recognition milestone

August 20, 2017 | By [Xuedong](#)



June 24, 2014

DeepFace: Closing the Gap in Face Recognition Performance in Face

Conference on Computer Vision and Pattern Recognition (CVPR)

By: [Yaniv Taigman](#), [Ming Yang](#), [Marc'Aurelio Ranzato](#), [Lior Wolf](#)

Abstract

In modern face recognition, the conventional pipeline consists of face detection, face alignment, and face classification. We revisit both the alignment step and the representation step, and propose a new pipeline that achieves state-of-the-art performance on the Labeled Faces in the Wild (LFW) benchmark.

If you think AI will never replace radiologists—you may want to think again

May 14, 2018 | [Michael Walter](#) | [Artificial Intelligence](#)



It's one of the most frequently discussed questions in radiology today: What kind of long-term impact will AI have on the profession?

It is hard these days to escape hearing about AI — in the news, on social media, in cafe conversations. We see both reports of triumphs of superhuman performance in games such as Jeopardy! (IBM Watson, 2011) and Go (DeepMind's AlphaGo, 2016), as well as on benchmark tasks such as reading comprehension, speech recognition, face recognition, and medical imaging (though it is important to realize that these are about performance on one benchmark, which is a far cry from the general problem).



周围都发生了什么？

Companies



"An important shift from a mobile first world to an AI first world" [CEO Sundar Pichai @ Google I/O 2017]



Created AI and Research group as 4th engineering division, now 8K people [2016]



Created Facebook AI Research, Mark Zuckerberg very optimistic and invested

Others: IBM, Amazon, Apple, Uber, Salesforce, Baidu, Tencent, etc.

周围都发生了什么？

Governments



"AI holds the potential to be a major driver of economic growth and social progress" [White House report, 2016]




Released domestic strategic plan to become world leader in AI by 2030 [2017]



"Whoever becomes the leader in this sphere [AI] will become the ruler of the world" [Putin, 2017]


周围都发生了什么？

AI index: number of published AI papers



周围都发生了什么？

AI index: number of AI startups



OK, 什么是AI?

Artificial intelligence

From Wikipedia, the free encyclopedia

"AI" redirects here. For other uses, see [AI \(disambiguation\)](#) and [Artificial intelligence \(disambiguation\)](#).

Artificial intelligence (AI), sometimes called **machine intelligence**, is [intelligence](#) demonstrated by [machines](#), in contrast to the **natural intelligence** displayed by humans and other animals. In [computer science](#) AI research is defined as the study of "intelligent agents": any device that perceives its environment and takes actions that maximize its chance of successfully achieving its goals.^[1] Colloquially, the term "artificial intelligence" is applied when a machine mimics "cognitive" functions that humans associate with other [human minds](#), such as "learning" and "problem solving".^[2]

Two views of AI



[AI agents](#): how can we re-create intelligence?



[AI tools](#): how can we benefit society?

An intelligent agent

Perception

Robotics

Language



Knowledge


Reasoning

Learning

The Turing Test (1950)

[Turing, 1950. Computing Machinery and Intelligence]

"Can machines think?"



Q: Please write me a sonnet on the subject of the Forth Bridge.

A: Count me out on this one. I never could write poetry.

Q: Add 34957 to 70764.

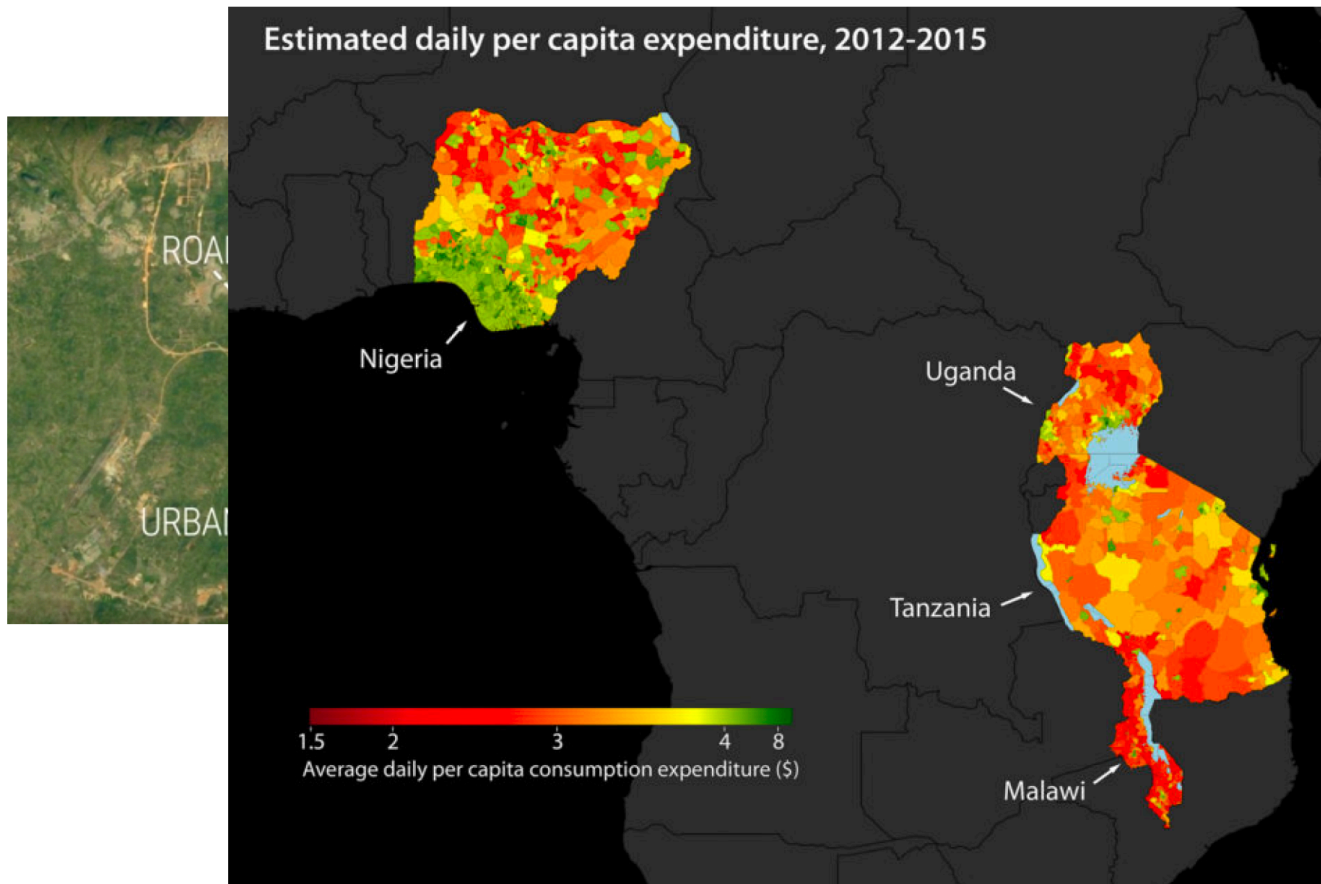
A: (Pause about 30 seconds and then give as answer) 105621.

Tests behavior — simple and objective






AI tools...

Predicting poverty



Saving energy by cooling datacenters





Fairness in criminal risk assessment

- **Northpointe**: COMPAS predicts criminal risk score (1-10)
- **ProPublica**: given that an individual did not reoffend, blacks 2x likely to be (wrongly) classified 5 or above
- **Northpointe**: given a risk score of 7, 60% of whites reoffended, 60% of blacks reoffended

California just replaced cash bail with algorithms

By [Dave Gershgorn](#) • September 4, 2018





Summary so far

- AI dream of achieving human-level intelligence is ongoing
 - Still lots of open research questions
 - AI is having huge societal impact
 - Need to think carefully about real-world consequences
-

A very brief history

- 1956: Dartmouth workshop, John McCarthy coined "AI"
 - 1960: checkers playing program, Logical Theorist
 - 1966: ALPAC report cuts off funding for translation
 - 1974: Lighthill report cuts off funding in UK
 - 1970-80s: expert systems (XCON, MYCIN) in industry
 - 1980s: Fifth-Generation Computer System (Japan); Strategic Computing Initiative (DARPA)
 - 1987: collapse of Lisp market, government funding cut
 - 1990-: rise of machine learning
 - 2010s: heavy industry investment in deep learning
-

试图回答几个问题

- 什么是AI及AI的发展历史
- 是什么让AI so HOT
- 我国AI的发展规划

AI is Hot partly because of the success of deep learning

Image classification

IMAGENET Large Scale Visual Recognition Challenge 2012 (ILSVRC2012)

Team name	Filename	Error (5 guesses)
SuperVision	test-preds-141-146.2009-131-137-145-146.2011-145f.	0.15315
SuperVision	test-preds-131-137-145-135-145f.txt	0.16422
ISI	pred_FVs_wLACs_weighted.txt	0.26172

<http://www.image-net.org/>

AI is Hot partly because of the success of deep learning

Image

IMAGENET Large Scale Visual Recognition Challenge 2013 (ILSVRC2013)

Team name	Comment	Error
Clarifai	Multiple models trained on the original data plus an additional model trained on 5000 categories.	0.11197
Clarifai	Multiple models trained on the original data plus an additional model trained on other 1000 category data.	0.11537
Clarifai	Average of multiple models on original training data.	0.11743
Clarifai	Another attempt at multiple models on original training data.	0.1215
Clarifai	Single model trained on original data.	0.12535
NUS	adaptive non-parametric rectification of all outputs from CNNs and refined PASCAL VOC12 winning solution, with further retraining on the validation set.	0.12953
NUS	adaptive non-parametric rectification of all outputs from CNNs and refined PASCAL VOC12 winning solution.	0.13303
ZF	5 models (4 different architectures) trained on original data.	0.13511
Andrew Howard	This is an ensemble of convolutional neural networks combining multiple transformations for training and testing and models operating at different resolutions.	0.13555
Andrew Howard	This method explores re weighting the predictions from different data transformation and ensemble members in the previous submission.	0.13564

AI is Hot partly because of the success of deep learning

Image classification

IMAGENET Large Scale Visual Recognition Challenge 2014 (ILSVRC2014)

Team name	Entry description	Localization error	Classification error
VGG	a combination of multiple ConvNets (by averaging)	0.253231	0.07405
VGG	a combination of multiple ConvNets (fusion weights learnt on the validation set)	0.253501	0.07407
VGG	a combination of multiple ConvNets, including a net trained on images of different size (fusion done by averaging); detected boxes were not updated	0.255431	0.07337
VGG	a combination of multiple ConvNets, including a net trained on images of different size (fusion weights learnt on the validation set); detected boxes were not updated	0.256167	0.07325
GoogLeNet	Model with localization ~26% top5 val error.	0.264414	0.14828
GoogLeNet	Model with localization ~26% top5 val error, limiting number of classes.	0.264425	0.12724
VGG	a single ConvNet (13 convolutional and 3 fully-connected layers)	0.267184	0.08434
SYSU_Vision	We compared the class-specific localization accuracy of solution 1 and solution 2 by the validation set. Then we chosen better solution on each class based on the accuracy. General speaking, solution 2 outformed solution 1 when there were multiple objects in the image or the objects are relatively small.	0.31899	0.14446
MIL	5 top instances predicted using FV-CNN	0.337414	0.20734
MIL	5 top instances predicted using FV-CNN + class specific window size rejection. Flipped training images are added.	0.33843	0.21023

AI is Hot partly because of the success of deep learning

IMAGENET Large Scale Visual Recognition Challenge 2015 (ILSVRC2015)

MSRA	Kaiming He Xiangyu Zhang Shaoqing Ren Jian Sun	<p>We train neural networks with depth of over 150 layers. We propose a deep residual learning framework [a] that eases the optimization and convergence of extremely deep networks. Our "deep residual nets" enjoy accuracy gains when the networks are substantially deeper than those used previously. Such accuracy gains are not witnessed for many common networks when going deeper.</p> <p>Our localization and detection systems are based on deep residual nets and the "Faster R-CNN" system in our NIPS paper [b]. The extremely deep representations generalize well, and greatly improve the results of the Faster R-CNN system. Furthermore, we show that the region proposal network (RPN) in [b] is a generic framework and performs excellent for localization.</p> <p>We only use the ImageNet main competition data. We do not use the Scene/VID data.</p>
-------------	---	--

Ordered by localization error

Team name	Entry description	Localization error	Classification error
MSRA	Ensemble A for classification and localization.	0.090178	0.03567
MSRA	Ensemble B for classification and localization.	0.090801	0.03567
MSRA	Ensemble C for classification and localization.	0.092108	0.0369
Travis Souther	combined 12 models	0.100007	0.04640

AI is Hot partly because of the success of deep learning

IMAGENET Large Scale Visual Recognition Challenge 2016 (ILSVRC2016)

CUIImage

Hongsheng Li*, Kai Kang* (* indicates equal contribution),
Wanli Ouyang, Junjie Yan,
Tong Xiao, Xingyu Zeng, Kun
Wang, Xihui Liu, Qi Chu,
Junming Fan, Yucong Zhou, Yu
Liu, Ruohui Wang, Shengen
Yan, Dahua Lin, Xiaogang
Wang

Compared with CUIImage submission in ILSVRC 2015, the new components are as follows.

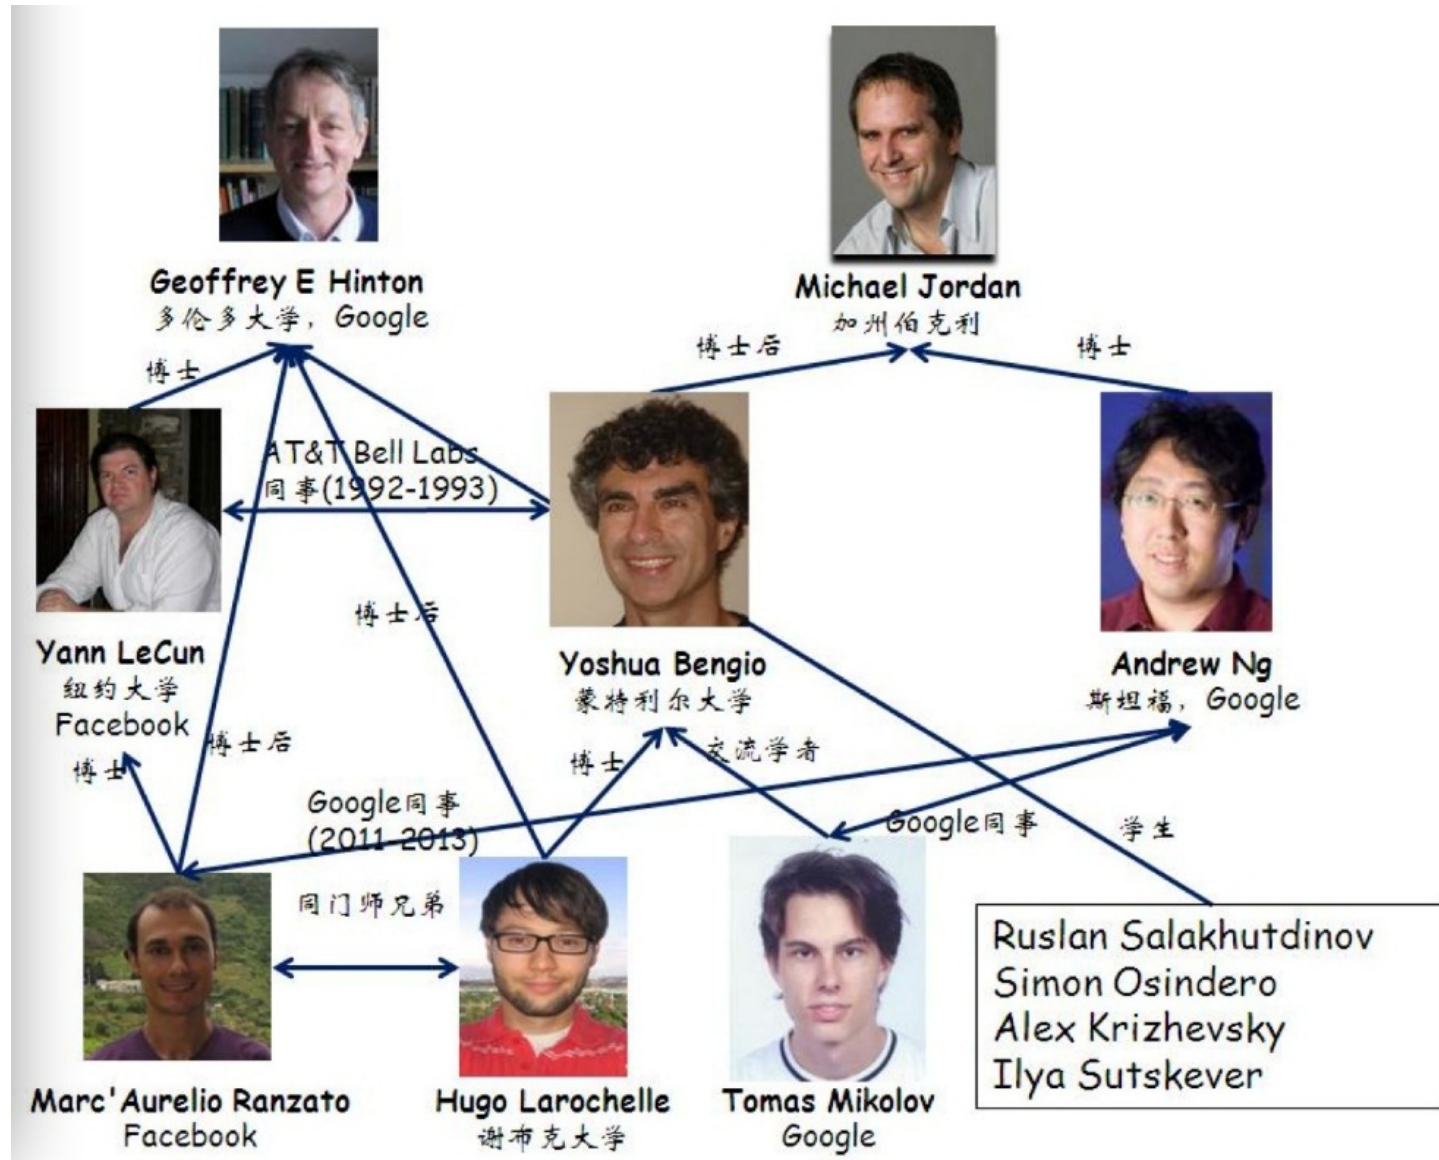
- (1) The models are pretrained for 1000-class object detection task using the approach in [a] but adapted to the fast-RCNN for faster detection speed.
- (2) The region proposal is obtained using the improved version of CRAFT in [b].
- (3) A GBD network [c] with 269 layers is fine-tuned on 200 detection classes with the gated bidirectional network (GBD-Net), which passes messages between features from different support regions during both feature learning and feature extraction. The GBD-Net is found to bring ~3% mAP improvement on the baseline 269 model and ~5% mAP improvement on the Batch normalized GoogleNet.




From then on, EVERYTHING is DEEP LEARNING!



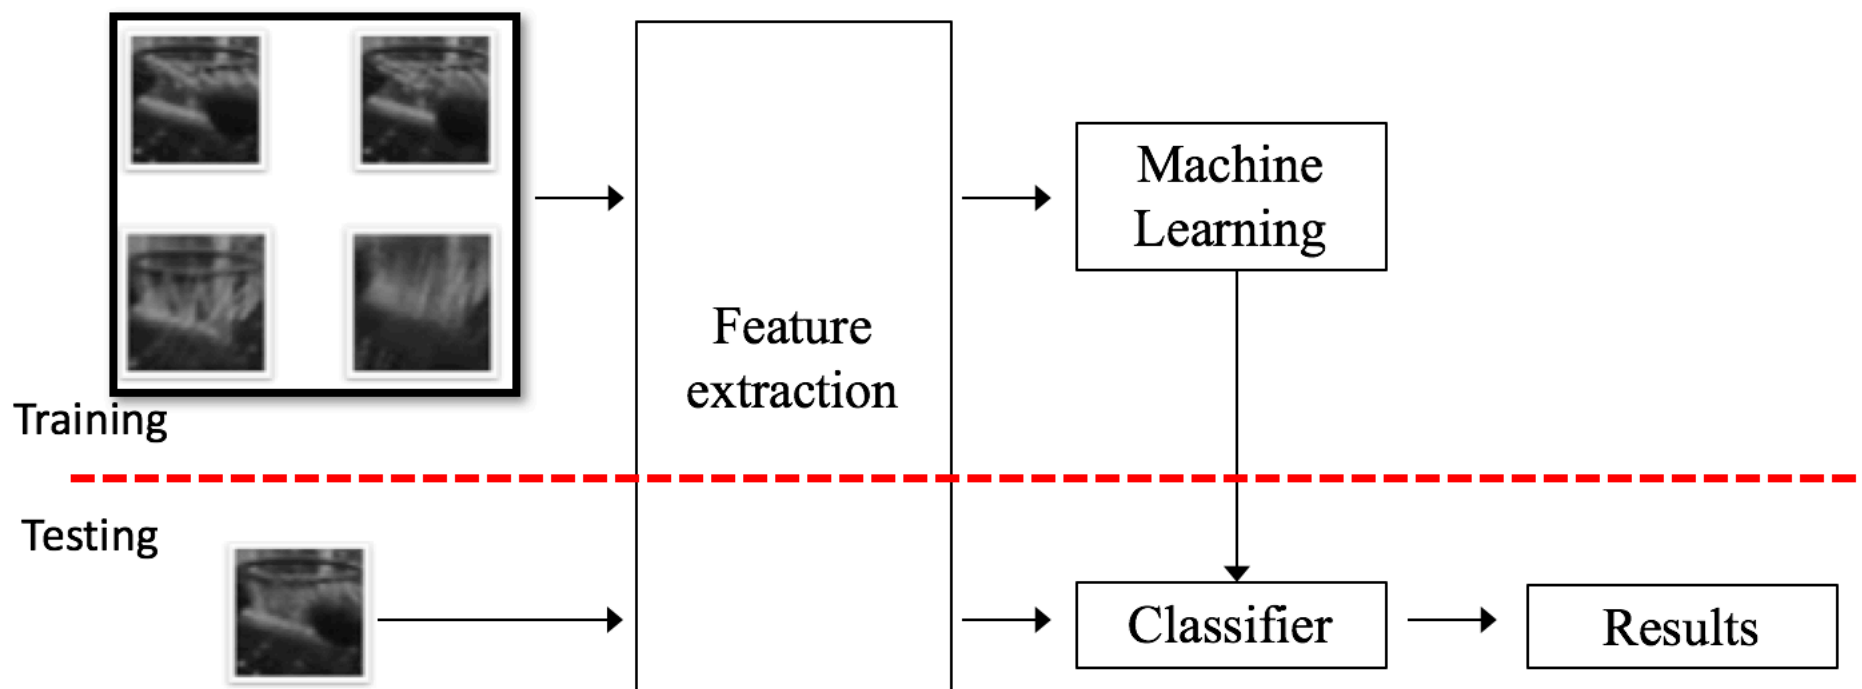
Who is who in deep learning?




大白话深度学习Deep Learning - 神经网络




大白话深度学习Deep Learning - 传统模式识别



大白话深度学习Deep Learning - 深度学习



目标函数



An end-to-end framework.

试图回答几个问题

- 什么是AI及AI的发展历史
- 是什么让AI so HOT
- 我国AI的发展规划

我国AI的发展规划 - 新一代人工智能



中华人民共和国科学技术部

Ministry of Science and Technology of the People's Republic of China

[微博微信](#) | [English](#) | [公务邮箱](#) | [加入收藏](#)



站内搜索

[首页](#) | [组织机构](#) | [新闻中心](#) | [信息公开](#) | [科技政策](#) | [科技计划](#) | [办事服务](#) | [公众参与](#) | [专题专栏](#)

信息名称： 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

索引号： 306-33-2018-131

信息类别： 规范性文件2018

发布机构： 科技部

发文日期： 2018年10月12日

文号： 国科发资〔2018〕208号

效力：

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

我国AI的发展规划 - 新一代人工智能

本重大项目的总体目标是：以推动人工智能技术持续创新和与经济社会深度融合为主线，按照并跑、领跑两步走战略，围绕大数据智能、跨媒体智能、群体智能、混合增强智能、自主智能系统等五大方向持续攻关，从基础理论、支撑体系、关键技术、创新应用四个层面构筑知识群、技术群和产品群的生态环境，抢占人工智能技术制高点，妥善应对可能带来的新问题和新的挑战，促进大众创业万众创新，使人工智能成为智能经济社会发展的强大引擎。

我国AI的发展规划 - 中国脑计划



2016年4月7日，中央批准了
“脑科学与类脑研究” 重大项目
立项建议。

科技创新2030——重大项目 实施方案编制工作手册

（脑科学与类脑研究重大科技项目）

附件二

中央审议通过的重大项目立项目标任务

脑科学与类脑研究

到 2020 年，在脑认知原理、类脑计算和脑重大疾病研究以及支撑平台建设方面，从总体上达到同期国际先进水平。到 2030 年，在基础研究方面，特别是脑认知功能相关神经环路工作原理及机制的深入解析方面，产生有重大国际影响的成果；在类脑计算、类脑芯片及脑机智能技术的研制、产业化方面，成为引领国际前沿的主要国家之一，支撑新型高科技产业的可持续发展；在脑重大疾病早期诊断、干预手段研发方面居国际前列，并显著提高此类疾病的疗效（或降低发病率），从而实现我国脑科学研究队伍、研究水平的跨越式、可持续发展。

统筹安排脑科学的基础研究、转化应用和相关产业发展，形成“一体两翼三平台”的布局，即以脑认知原理（认识脑）为主体，以类脑计算和脑机智能（模拟脑）和脑重大疾病诊治（保护脑）为两翼，搭建脑结构与功能研究新技术平台、脑重大疾病临床研究技术平台和类脑研究工程平台等三个关键技术平台；在重要的前沿领域争取与美、

我国AI的发展规划 - 中国脑计划

国家战略需求

认知障碍相关
重大脑疾病诊治

儿童青少年
脑智开发

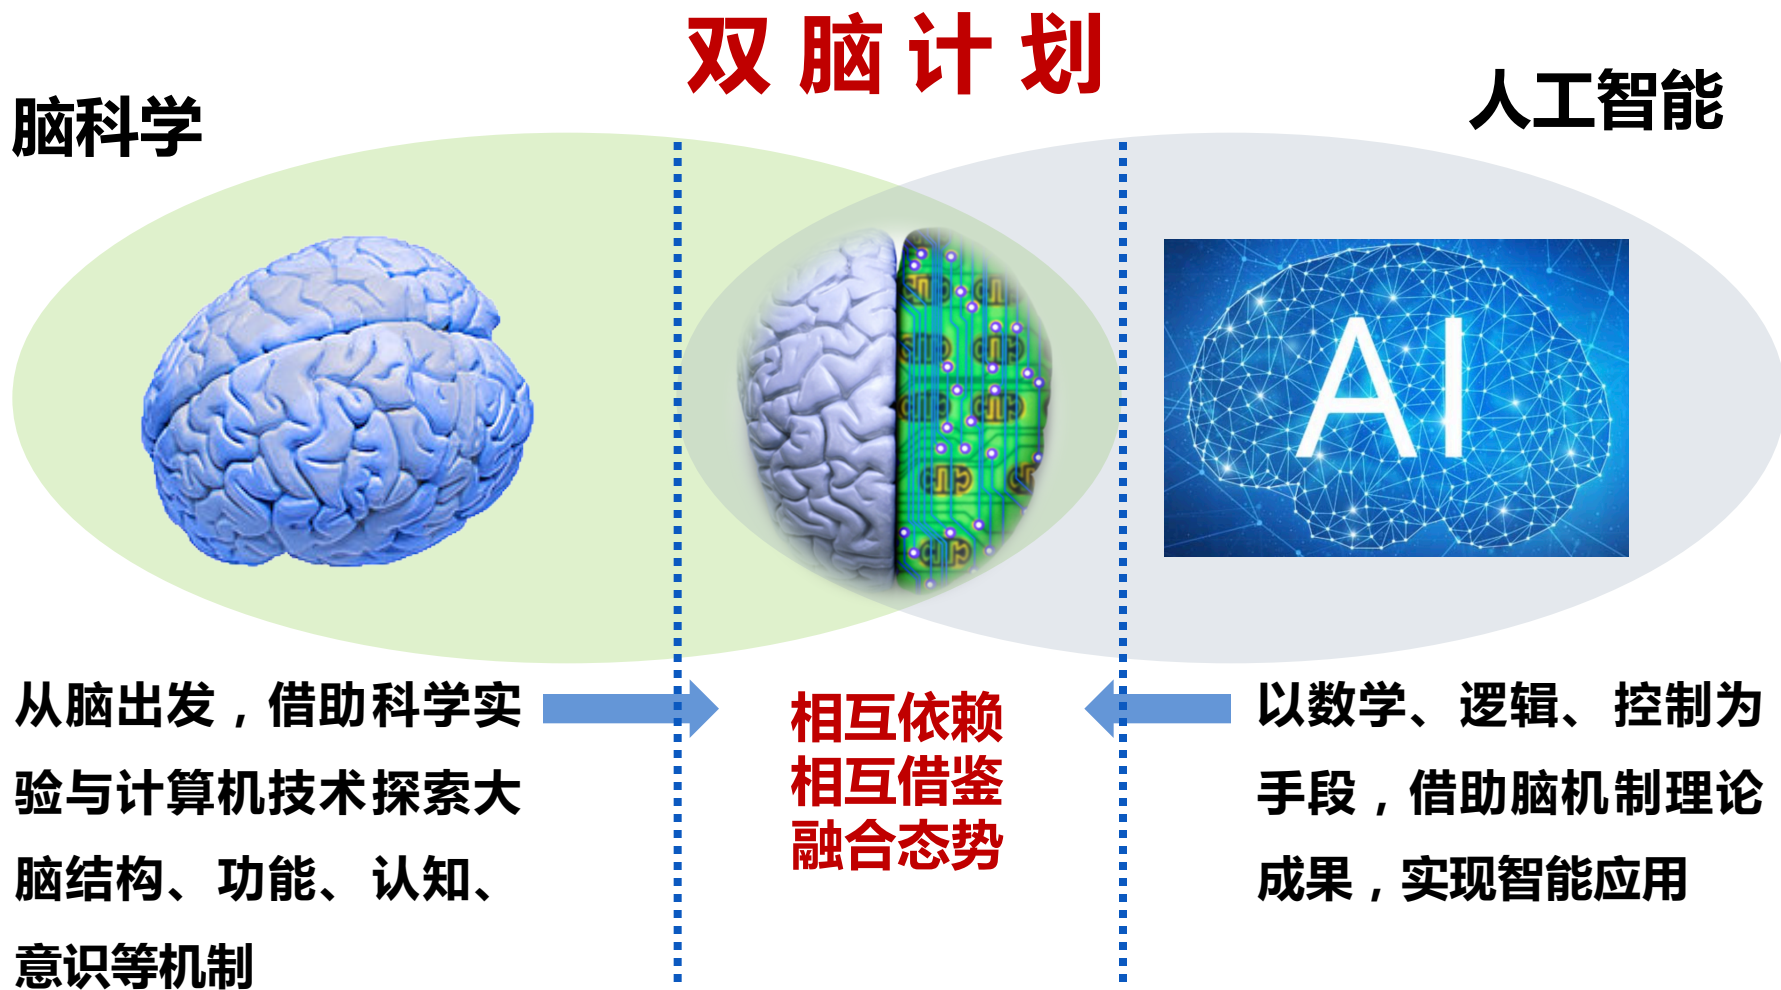
类脑计算与脑机
智能技术

脑认知功能解析
国际大科学计划

技术平台

世界科学前沿

我国AI的发展规划 - 浙江大学双脑计划



聚焦影响双方互进重大科学问题，交互探索，融合创新！

THANKS, Q&A
