

生物统计学： 生物信息中的概率统计模型

2019年秋



Yann LeCun
Geoffrey Hinton
Yoshua Bengio
Andrew Ng

有关信息

- 授课教师：宁康

- Email: ningkang@hust.edu.cn

- Office: 华中科技大学东十一楼504室

- Phone: 87793041, 18627968927

- 课程网页

- <http://www.microbioinformatics.org/teach/#>

- QQ群: 882140516



2019生物统计学
扫一扫二维码，加入该群。

课程安排

- 生物背景和课程简介
- 传统生物统计学及其应用
- 生物统计学和生物大数据挖掘
 - Hidden Markov Model (HMM)及其应用
 - Markov Chain
 - HMM理论
 - HMM和基因识别 (Topic I)
 - HMM和序列比对 (Topic II)
 - 进化树的概率模型 (Topic III)
 - Motif finding中的概率模型 (Topic IV)
 - EM algorithm
 - Markov Chain Monte Carlo (MCMC)
 - 基因表达数据分析 (Topic V)
 - 聚类分析-Mixture model
 - Classification-Lasso Based variable selection
 - 基因网络推断 (Topic VI)
 - Bayesian网络
 - Gaussian Graphical Model
 - 基因网络分析 (Topic VII)
 - Network clustering
 - Network Motif
 - Markov random field (MRF)
 - Dimension reduction及其应用 (Topic VIII)
- 面向生物大数据挖掘的深度学习

研究对象：
生物序列，
进化树，
生物网络，
基因表达

...

方法：
生物计算与生物统计

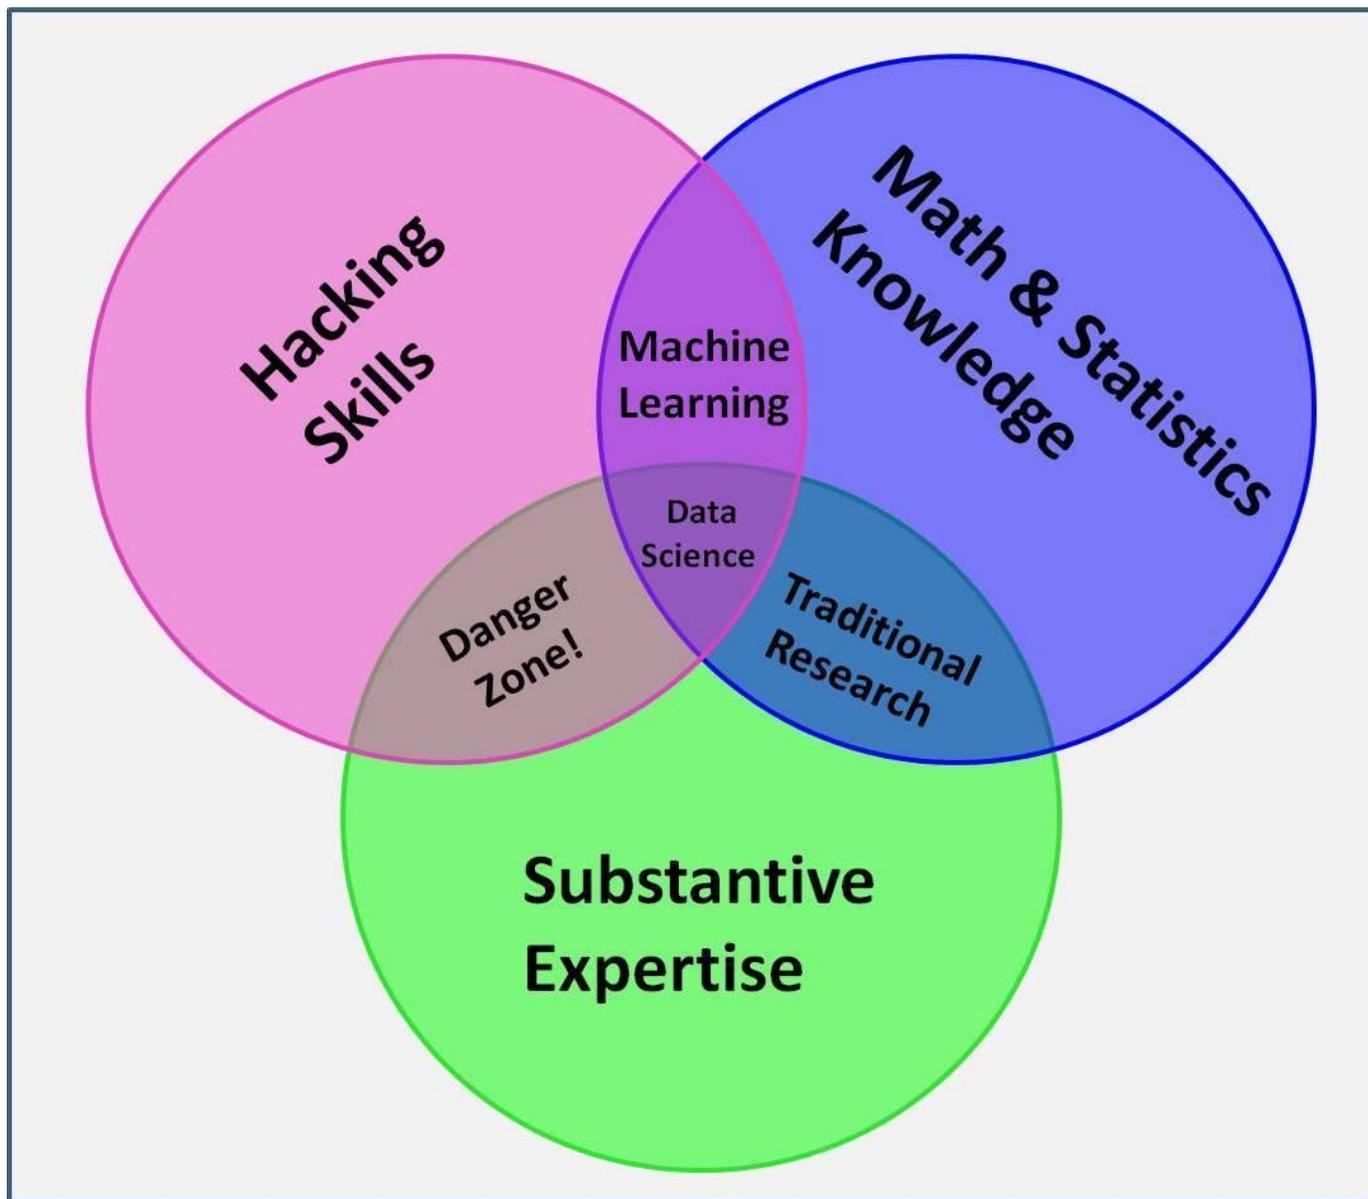
第10章： 面向生物大数据挖掘的深度学习

- 深度学习是什么
- 深度学习的基本方法（以及神经网络算法）
- 生物大数据的深度学习方法
- 生物大数据深度学习的应用案例

Part I

深度学习是什么

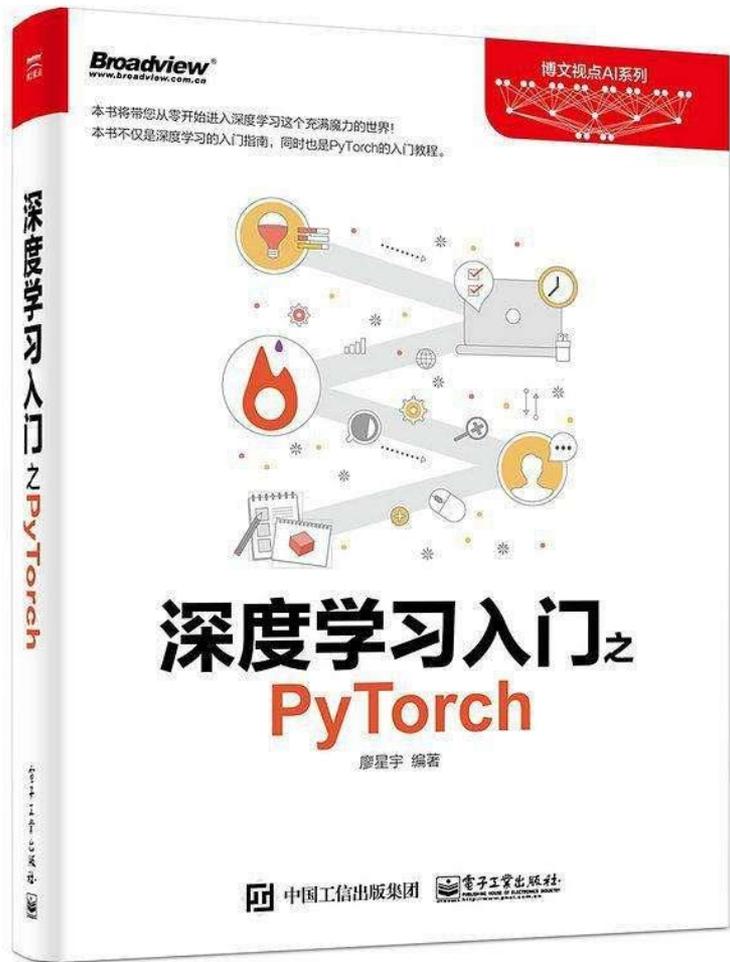
深度学习与数据科学



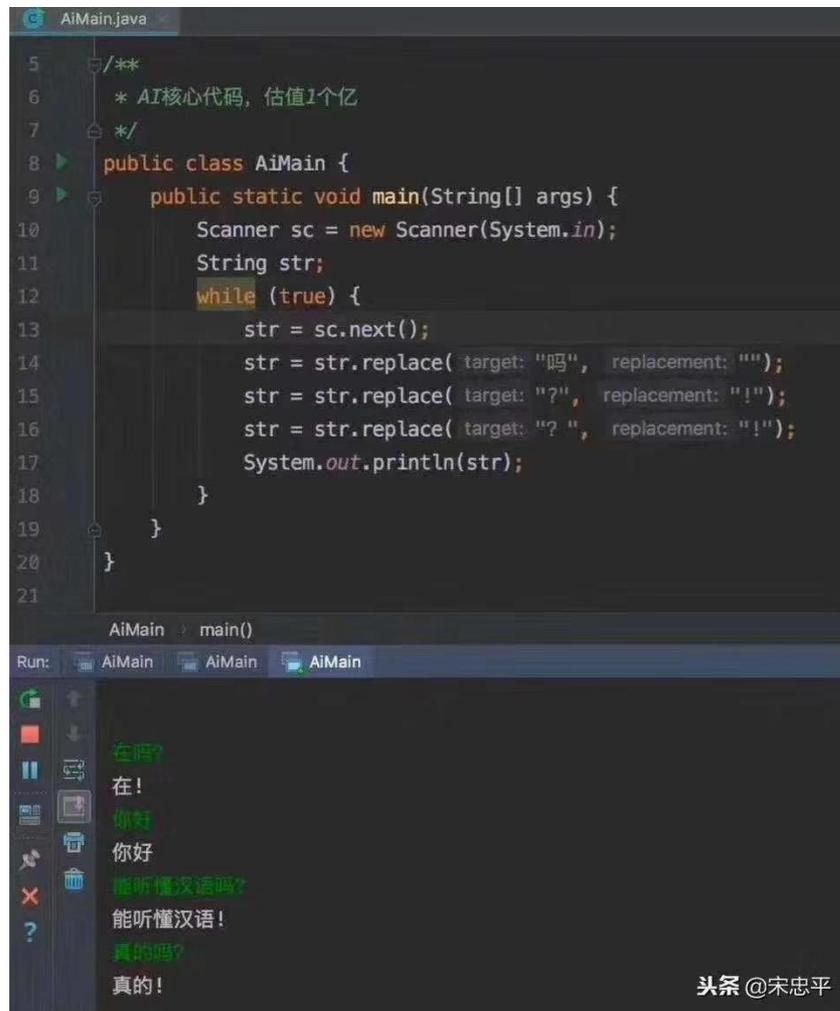
深度学习与数据科学

- 深度学习：一种基于无监督特征学习和特征层次结构的学习方法
- 可能的名称：
 - 深度学习
 - 特征学习
 - 无监督特征学习

真假深度学习



真



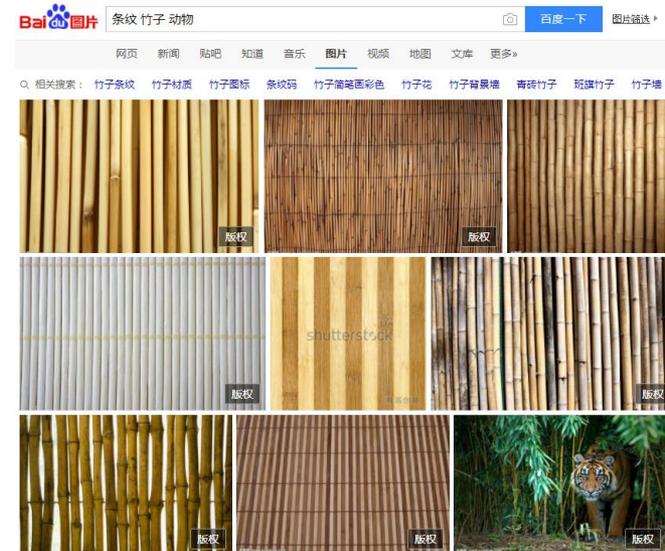
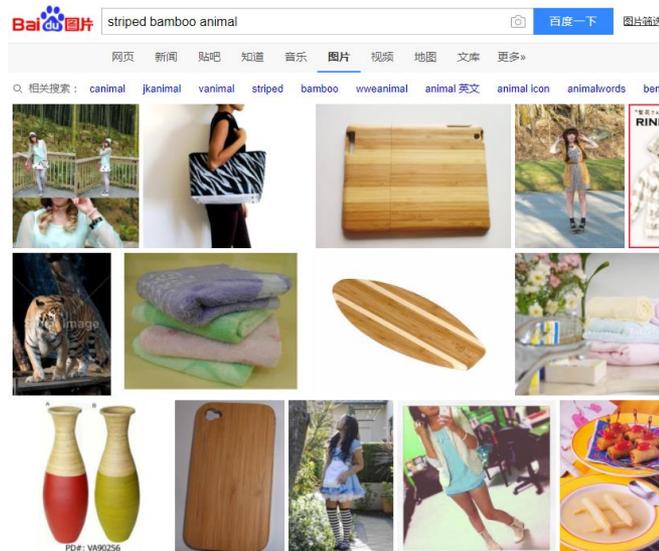
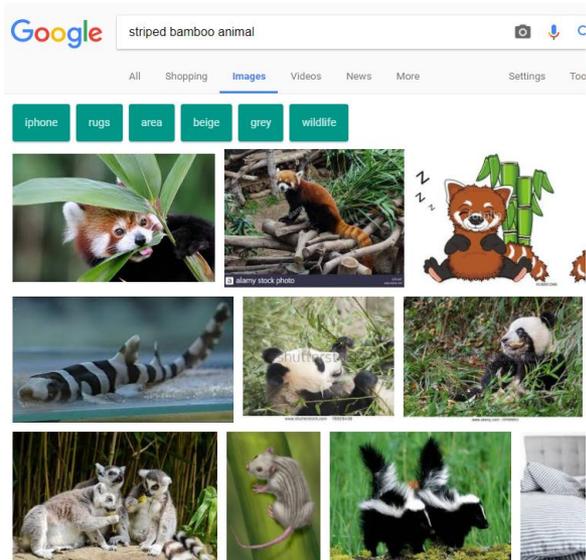
假

搜索引擎的区别

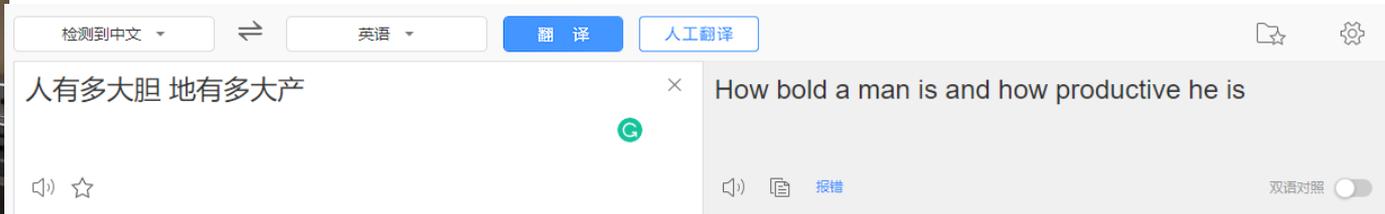
striped bamboo animal

striped bamboo animal

条纹竹子动物



翻译引擎和人类的较量



翻译引擎和人类的较量

< Simple
Florida Georgia Line >

We've been there, it's safe to say it ain't our style
我们也曾有过这种处境，但我们可以肯定地说那不是我们的风格

It's just that simple, S-I-M-P-L-E
就是如此简单，鸡一安简，的安单

Simple as can be
尽可能的简单

It's just that simple, S-I-M-P-L-E
就是如此简单，鸡一安简，的安单

Simple as can be
尽可能的简单

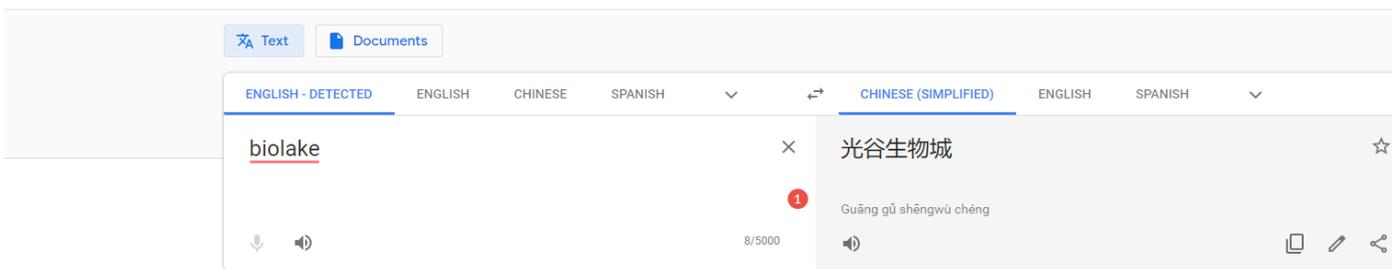


翻译引擎的区别

biolake



Google Translate



biolake



光谷生物城



翻译的区别（互有优劣）

它很难吃。

谷歌：It is hard to eat.

百度：It tastes terrible.

我要下班。

谷歌：I want to get off work from work.

百度：I'm going to get off work.

The image shows a side-by-side comparison of two translation services: Baidu Translate (top) and Google Translate (bottom). Both are translating the same English text from a news article about a council meeting in Charlottesville. The English text on the left includes a title, a short text paragraph, and two longer paragraphs. The Baidu Translate results on the right show a title and a long, somewhat awkward Chinese paragraph. The Google Translate results on the right show a title and a more natural, readable Chinese paragraph.

Baidu Translate (Top):

检测到英语 → 中文 → 翻译 人工翻译

title: Anger Boils Over at Charlottesville Council Meeting

text: Anger boiled over at the first Charlottesville City Council meeting since a white nationalist rally in the city descended into violent chaos, with some residents screaming and cursing at councilors Monday night and calling for their resignations.

Scores of people packed the council's chambers, and The Daily Progress reported Mayor Mike Signer was interrupted by shouting several times in the first few minutes of the meeting. As tensions escalated, the meeting was halted. Live video showed protesters standing on a dais with a sign that said, "Blood on your hands."

标题：愤怒沸腾的夏洛茨维尔会议
正文：愤怒爆发了第一次夏洛茨维尔市议会会议以来，在市一个白人民族主义者集会演变成暴力骚乱，一些居民的尖叫和咒骂议员星期一晚间呼吁辞职。许多人挤满了会议室，每日进展报告说，市长Mike Signer在会议的头几分钟被人多次叫喊。随着紧张局势升级，会议被叫停。现场视频显示，示威者站在一个标志说，“你手上的血。”在与群众交谈后，议员Wes Bellamy说，安理会将放弃议程，并集中在群众的关注，该报报道。
喇叭，大喊大叫的，还骂人，然后轮流解决委员会，一些挫折，领导人授予的8月12日的集会变成了暴力的许可证。其他人批评警方对这一事件的反应，该事件引来数百名白人民族主义者和其他反示威者。
阅读更多

Google Translate (Bottom):

翻译 中文(简体) 英语 日语

title: Anger Boils Over at Charlottesville Council Meeting

text: Anger boiled over at the first Charlottesville City Council meeting since a white nationalist rally in the city descended into violent chaos, with some residents screaming and cursing at councilors Monday night and calling for their resignations.

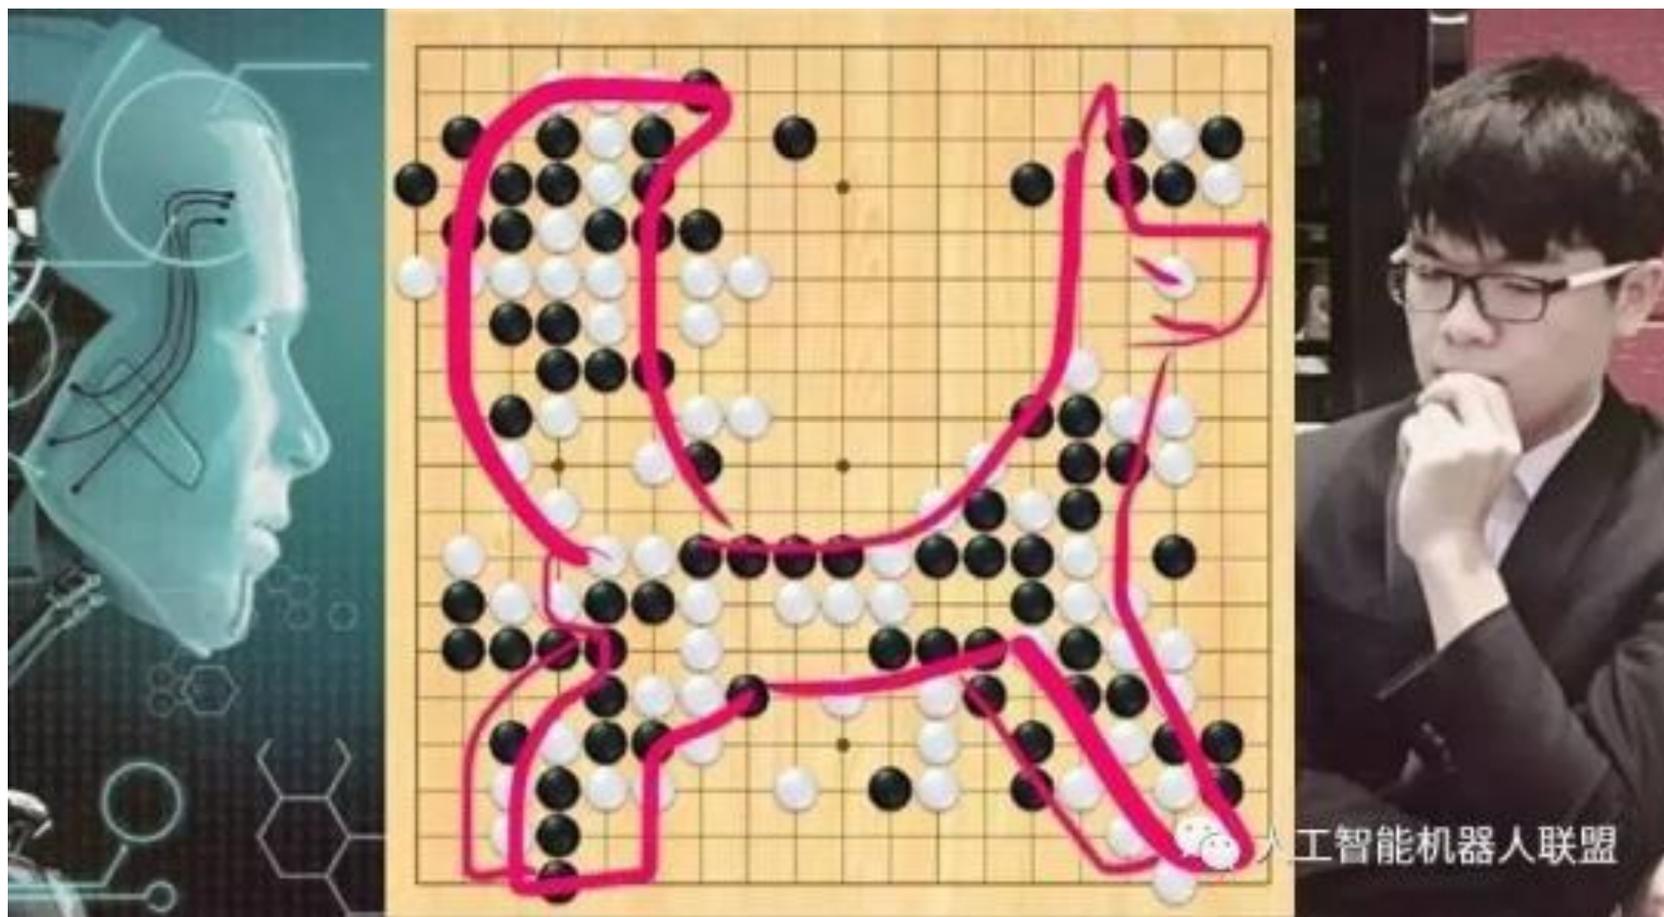
Scores of people packed the council's chambers, and The Daily Progress reported Mayor Mike Signer was interrupted by shouting several times in the first few minutes of the meeting. As tensions escalated, the meeting was halted. Live video showed protesters standing on a dais with a sign that said, "Blood on your hands."

After talking with members of the crowd, Councilor Wes Bellamy said the council would drop its agenda and focus on the crowd's concerns, the newspaper reported.

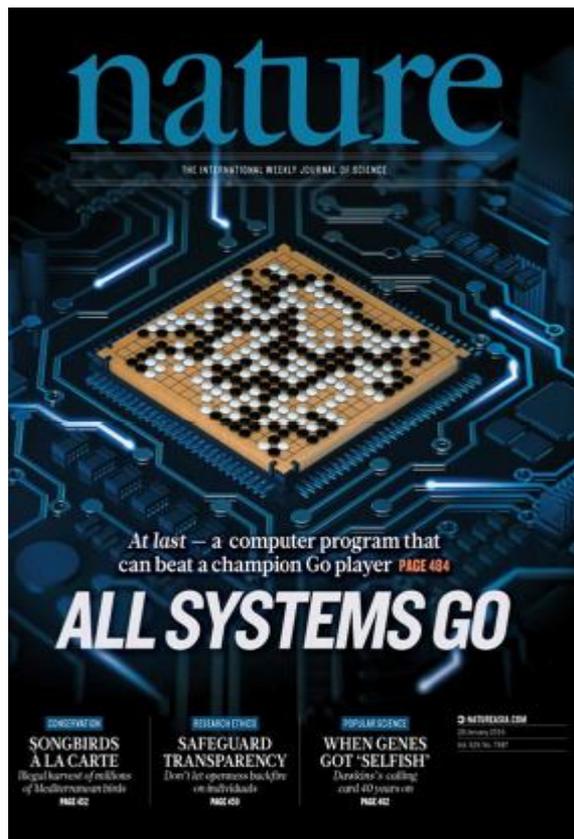
Speakers, some yelling and hurling profanities, then took turns addressing the council, some expressing frustration that leaders had granted a permit for the Aug. 12 rally that had turned violent. Others criticized the police response to the event, which drew hundreds of white nationalists and other counter-protesters.

标题：愤怒在夏洛茨维尔理事会会议上沸腾
文字：在第一届夏洛茨维尔市议会会议上，愤怒者沸腾起来，因为白城民族主义集会在这个城市发生暴力混乱，一些居民星期一晚上在议员尖叫和骂声，并呼吁辞职。
几分钟的人们挤满了议会的众议院；“日报”报道说，市长迈克·派克（Mike Signer）在会议的前几分钟多次喊叫中断。随着紧张局势升级，会议停止了。现场视频显示抗议者站在一个戴着标志的表示说，“血在你手上。”
报纸报道，在与人群中成员交谈之后，议员韦斯·贝拉米（Wes Bellamy）表示，该会议将放弃其议程，重点关注人群的关切。
发言人，一些大喊大叫和肆意的表演言论，然后轮到议会，有些人表示沮丧，领导人已经给予8月12日暴力许可。其他人批评警方对这一事件的反应，并提出了数百名白人民族主义者和其它反抗抗议者。

AlphaGo



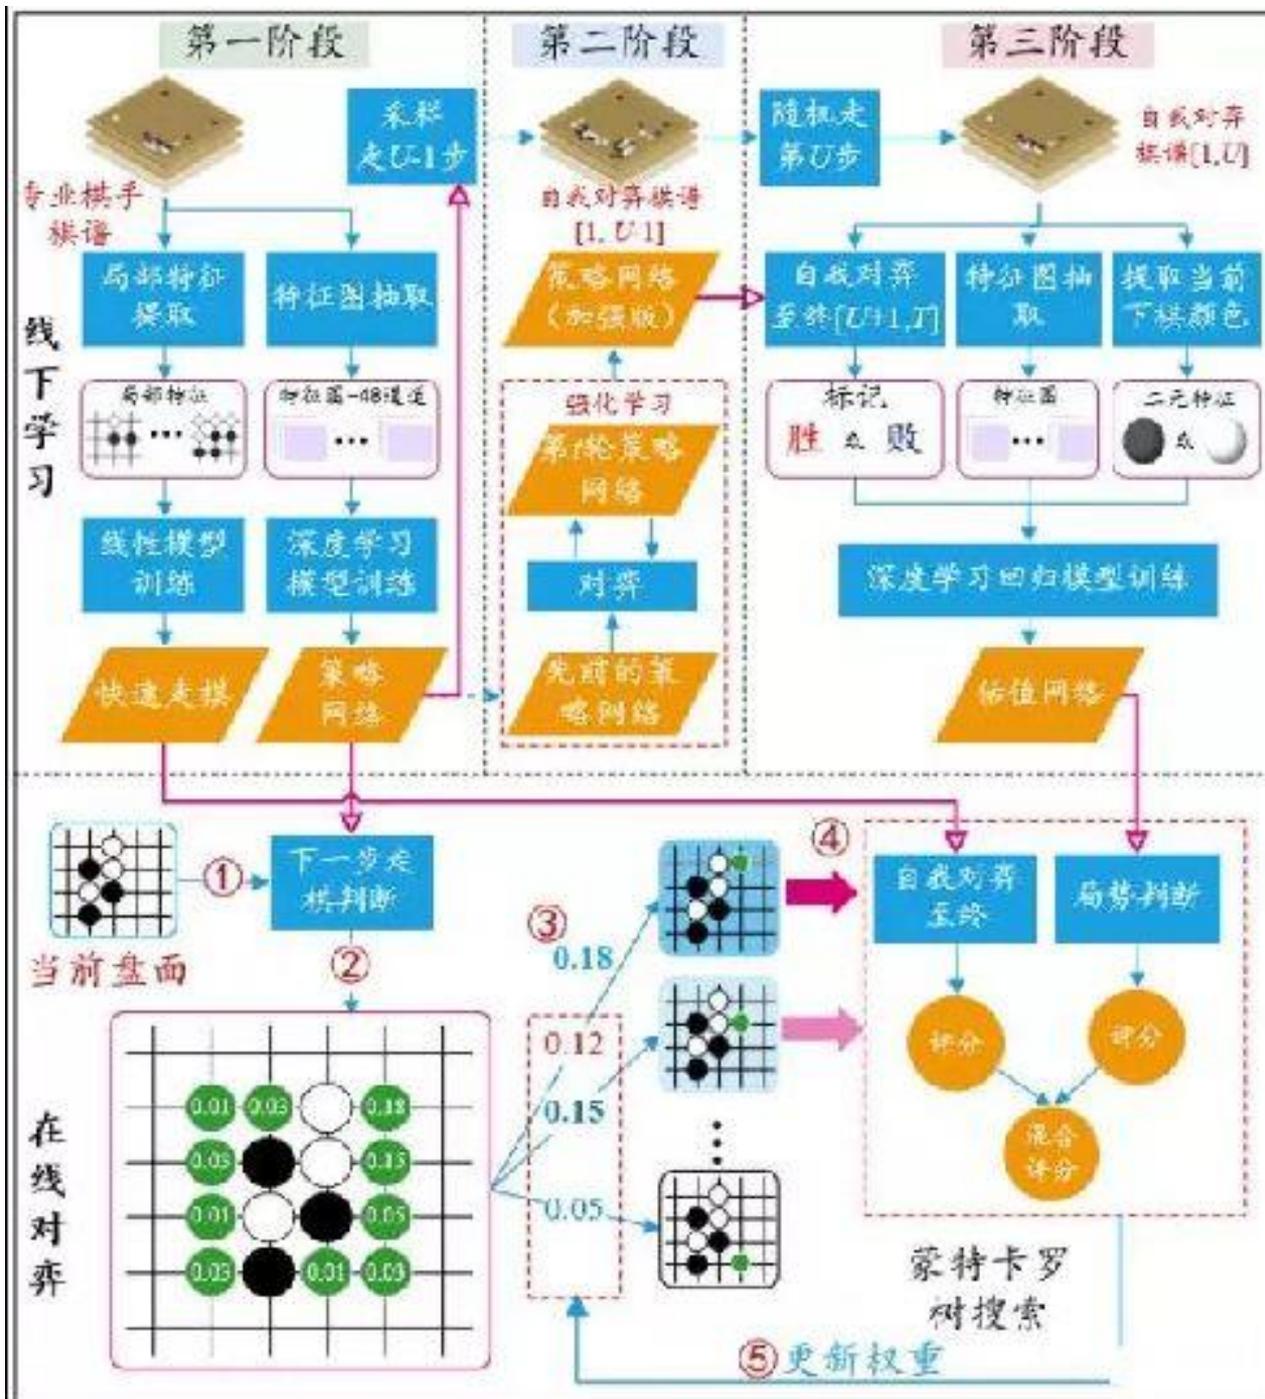
AlphaGo



ARTICLE

Mastering the game of Go with deep neural networks and tree search

David Silver¹*, Aja Huang^{1*}, Chris J. Maddison¹, Arthur Guez¹, Laurent Sifre¹, George van den Driessche¹, Julian Schrittwieser¹, Ioannis Antonoglou¹, Veda Pannemsetharam¹, Marc Lanctot¹, Sander Dieleman¹, Dominik Grewe¹, John Nham¹, Nal Kalchbrenner¹, Ilya Sutskever¹, Timothy Lillicrap¹, Madeleine Leach¹, Koray Kavukcuoglu¹, Thore Graepel¹ & Demis Hassabis¹



AlphaGo vs. AlphaGo Zero



数据很重要!



算法更重要!



AlphaZero

Reinforcement Learning (强化学习)





Transfer Learning (迁移学习)

The international journal of science / 14 November 2019

nature

GAME PLAN

AI program learns to play *StarCraft II* to Grandmaster level

Pharmaceuticals
How to fit a drug factory inside a briefcase

3D printing
Nozzle extrudes multimaterial objects in a single run

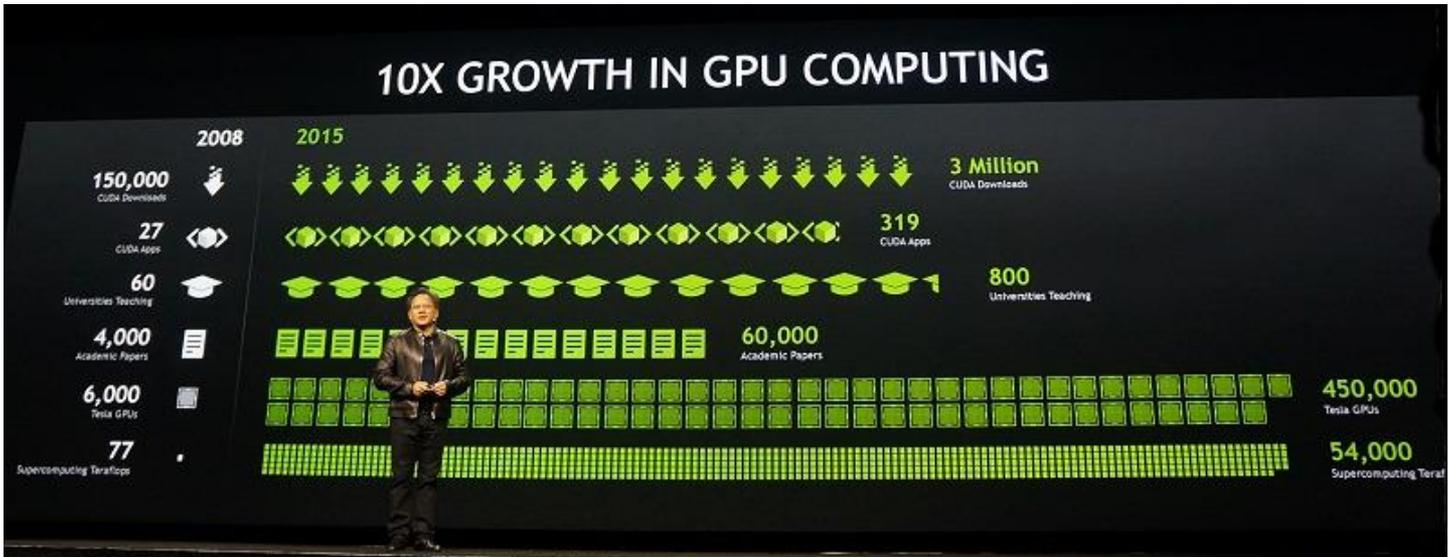
Cancer imaging
Tracer reveals metabolic nature of live lung tumours

Vol. 575, No. 783
61000 nature.com

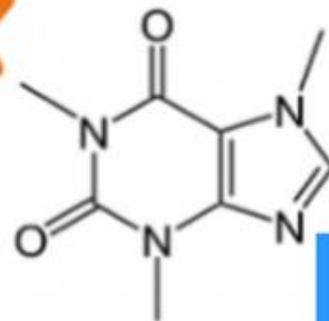


46 >

Hardware: GPU, HPC...



Software: TensorFlow, Caffe...



TensorFlow

Keras

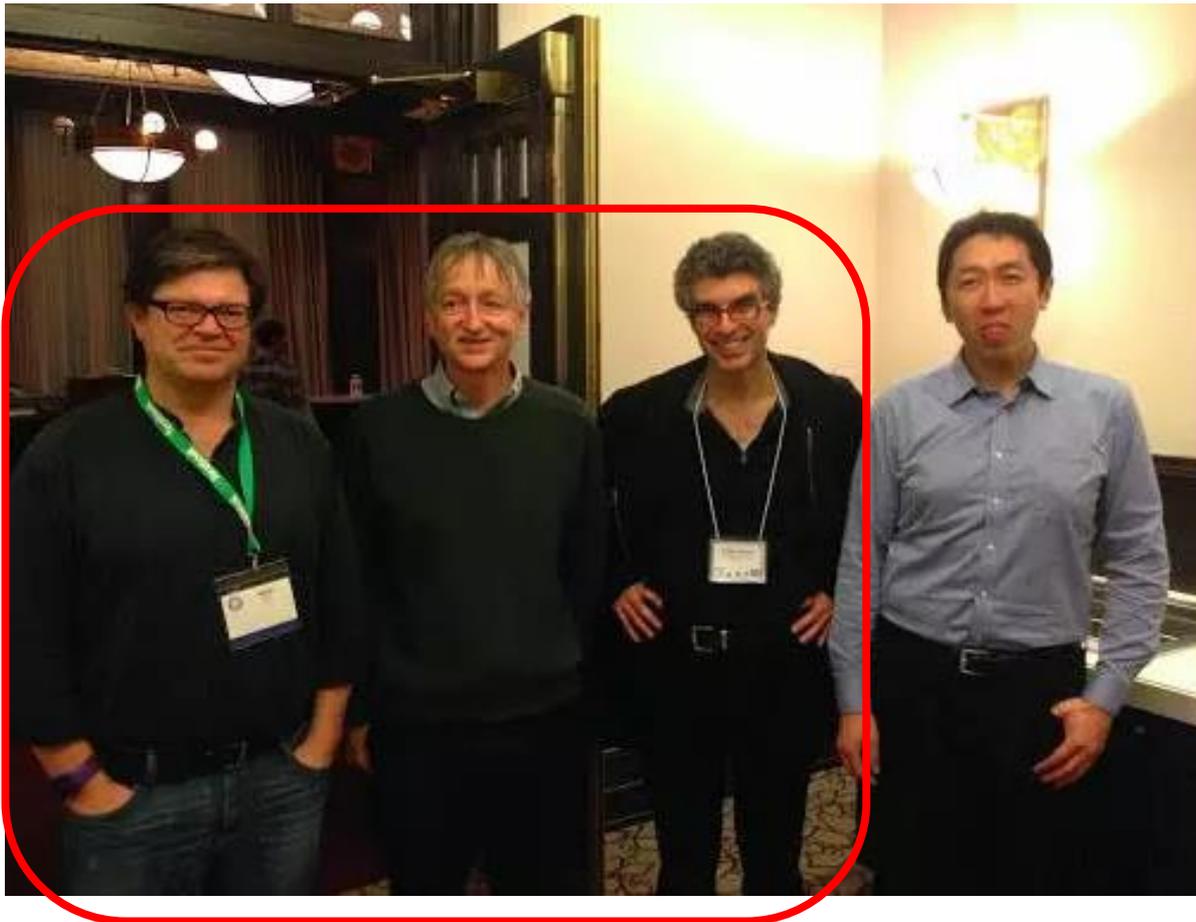
Torch



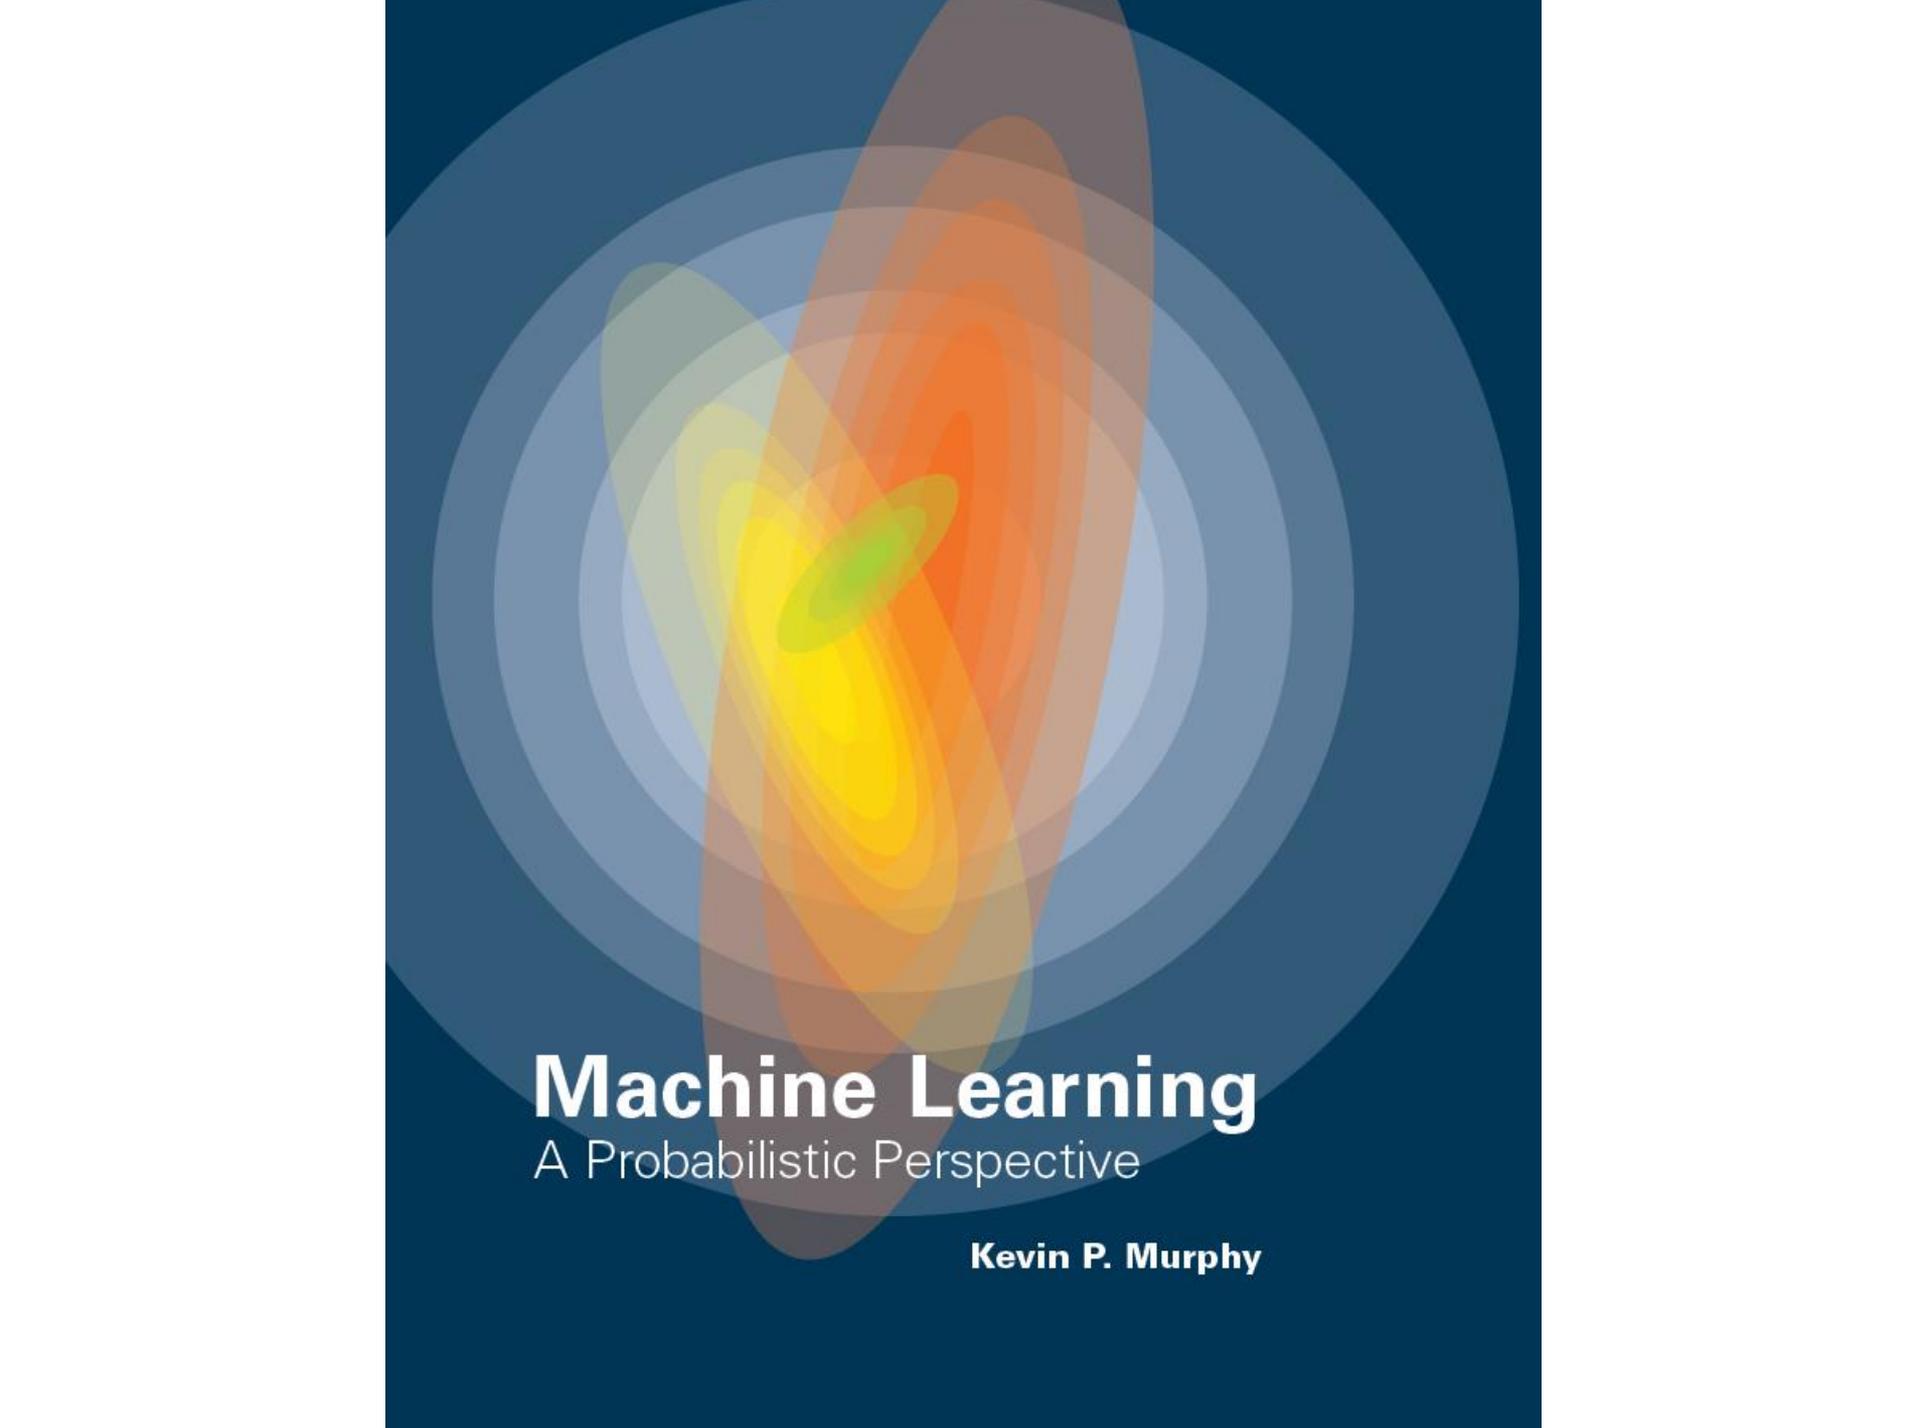
elastic



Turing Award, 2019



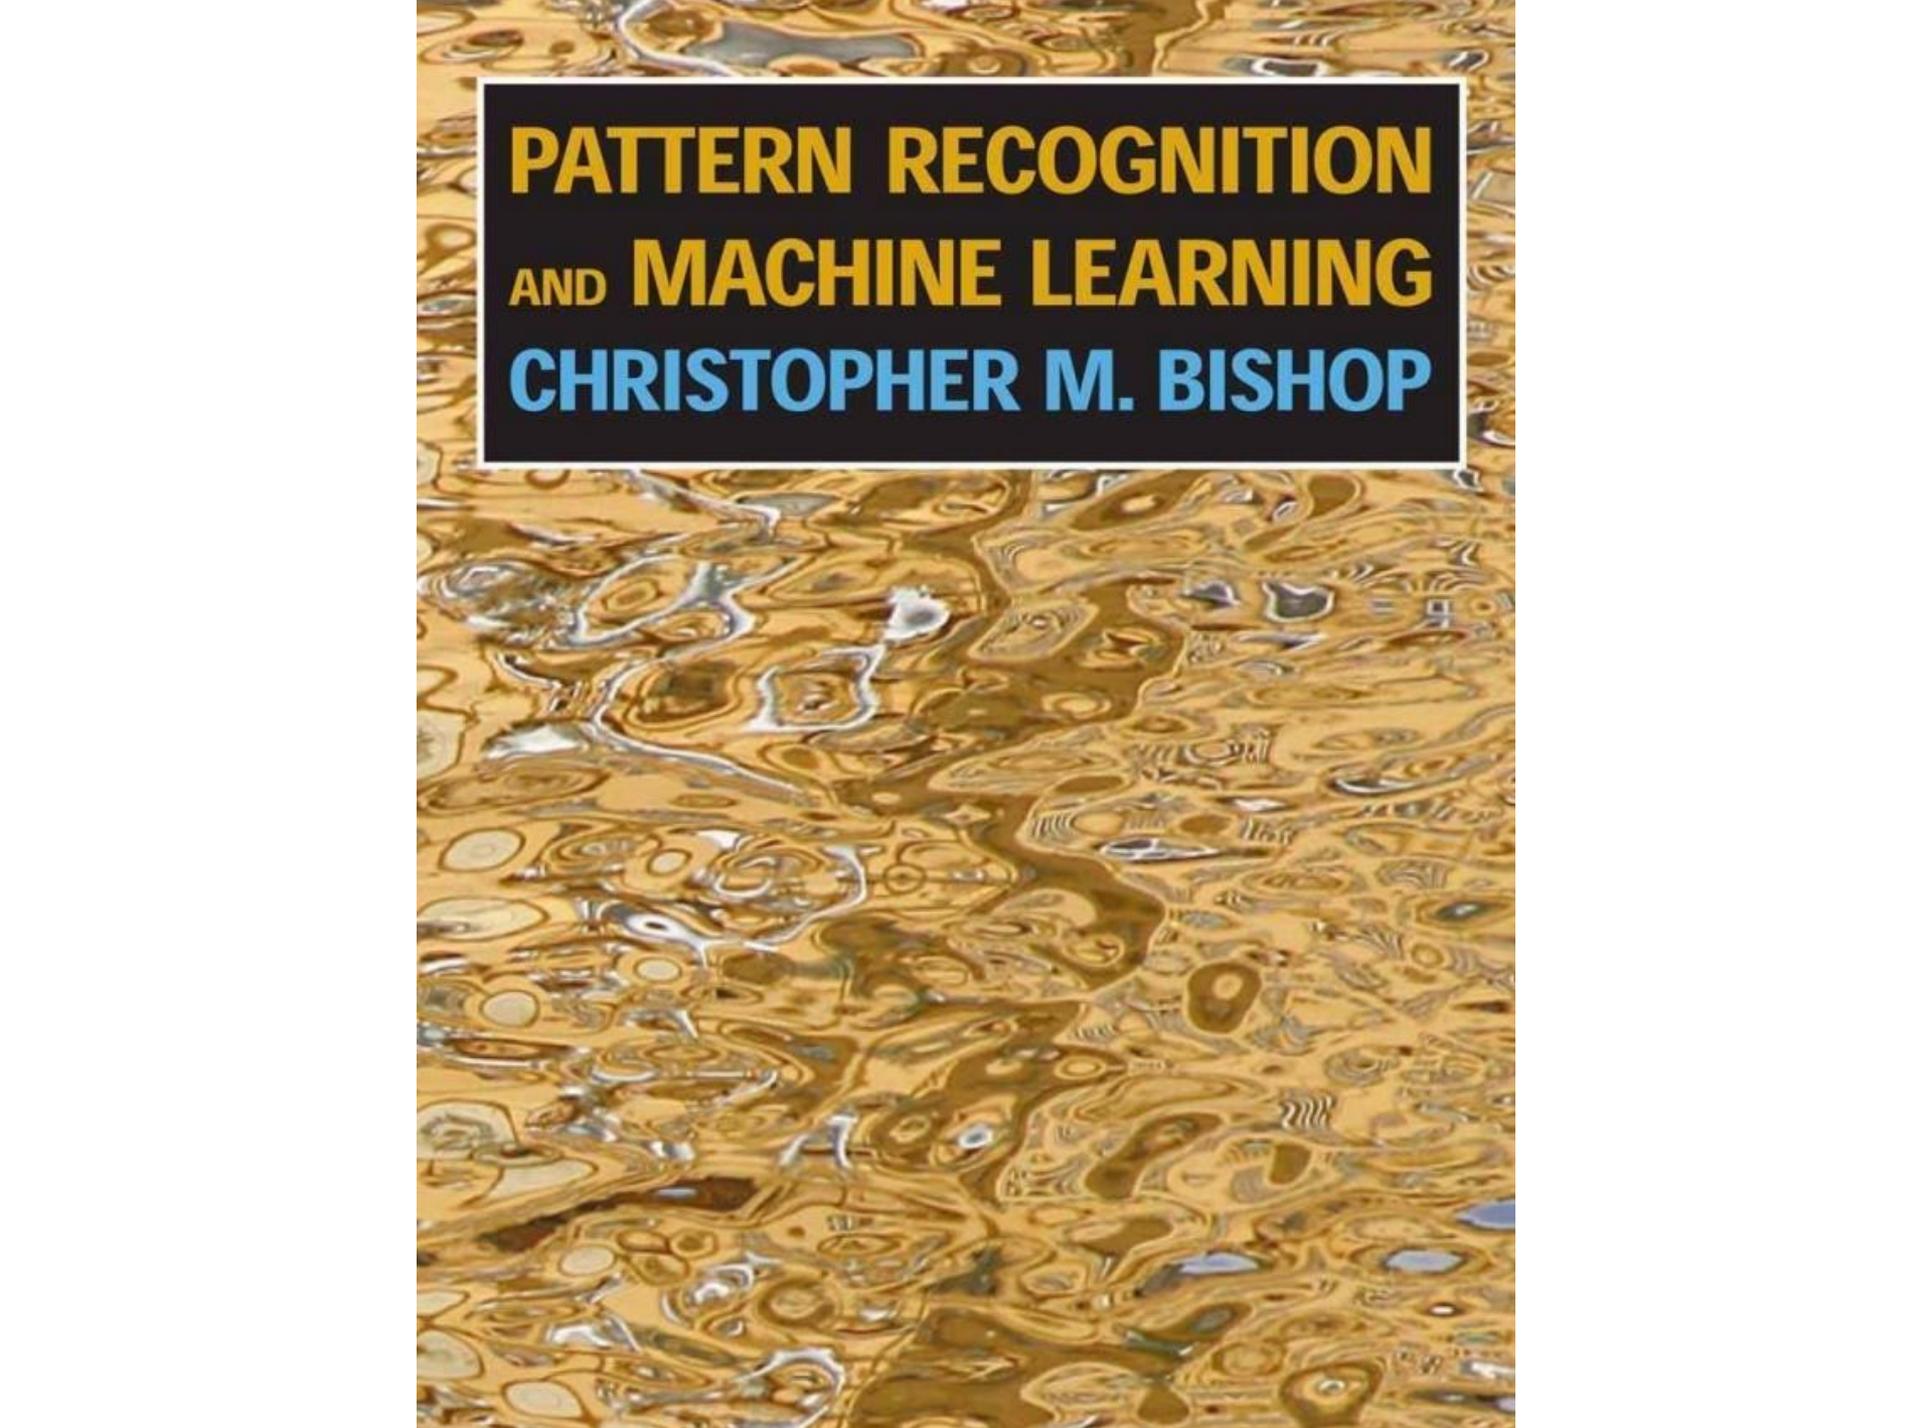
Yann LeCun
Geoffrey Hinton
Yoshua Bengio
Andrew Ng



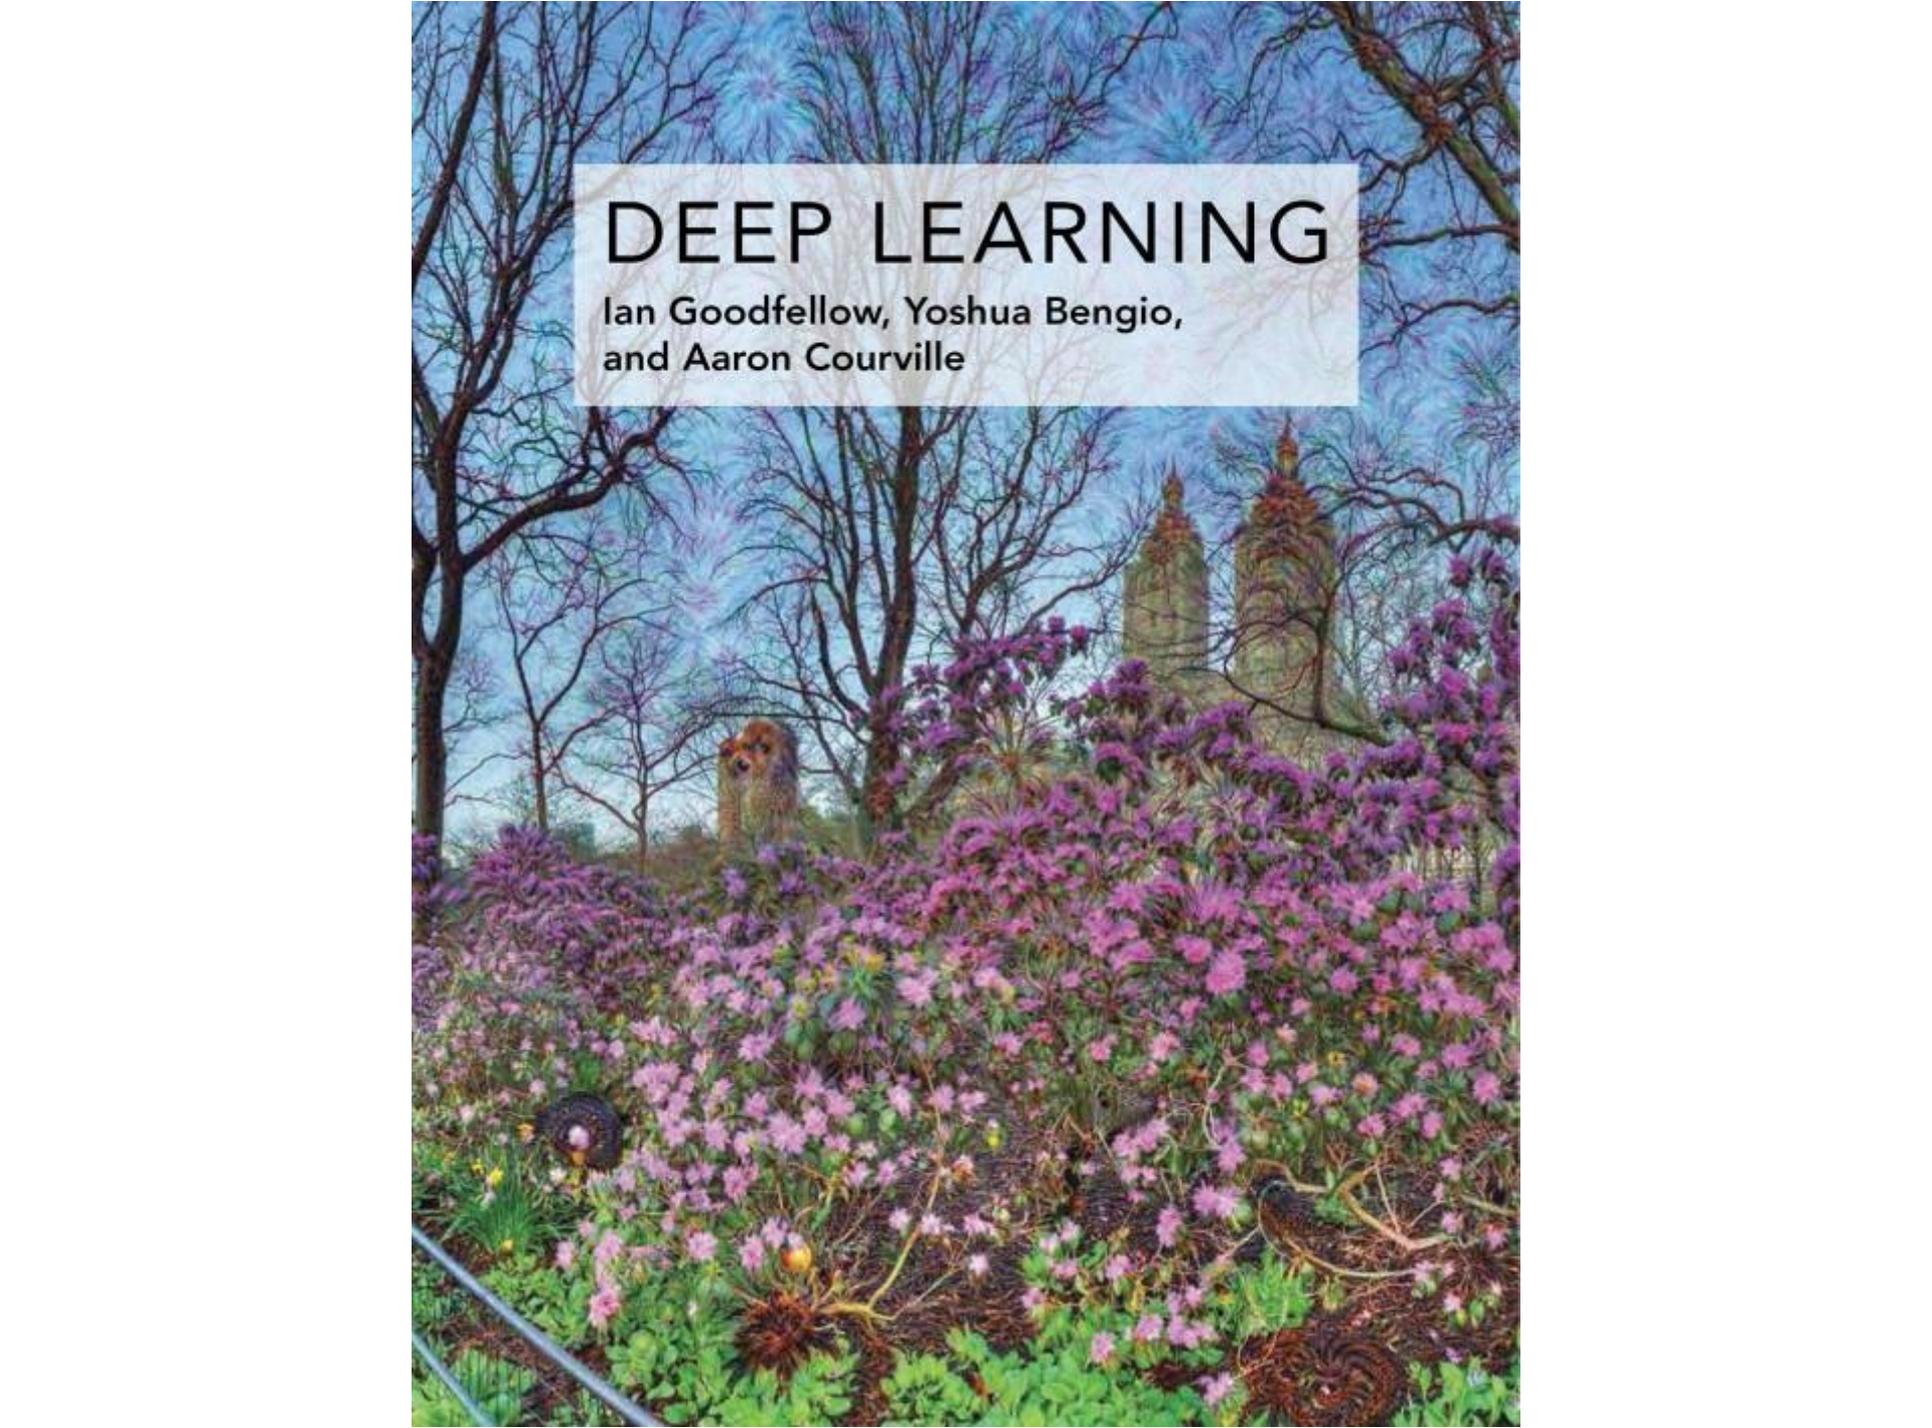
Machine Learning

A Probabilistic Perspective

Kevin P. Murphy



**PATTERN RECOGNITION
AND MACHINE LEARNING
CHRISTOPHER M. BISHOP**

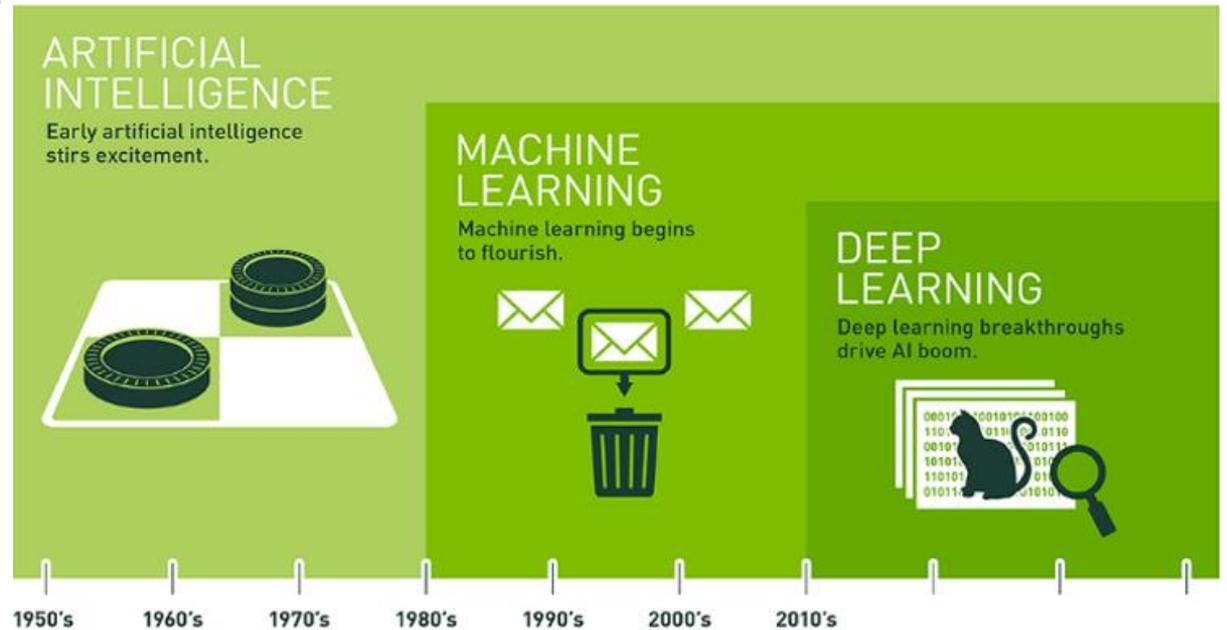
A photograph of a garden with purple flowers and trees, with a white text box overlaid. The text box contains the title 'DEEP LEARNING' and the authors' names: 'Ian Goodfellow, Yoshua Bengio, and Aaron Courville'.

DEEP LEARNING

Ian Goodfellow, Yoshua Bengio,
and Aaron Courville

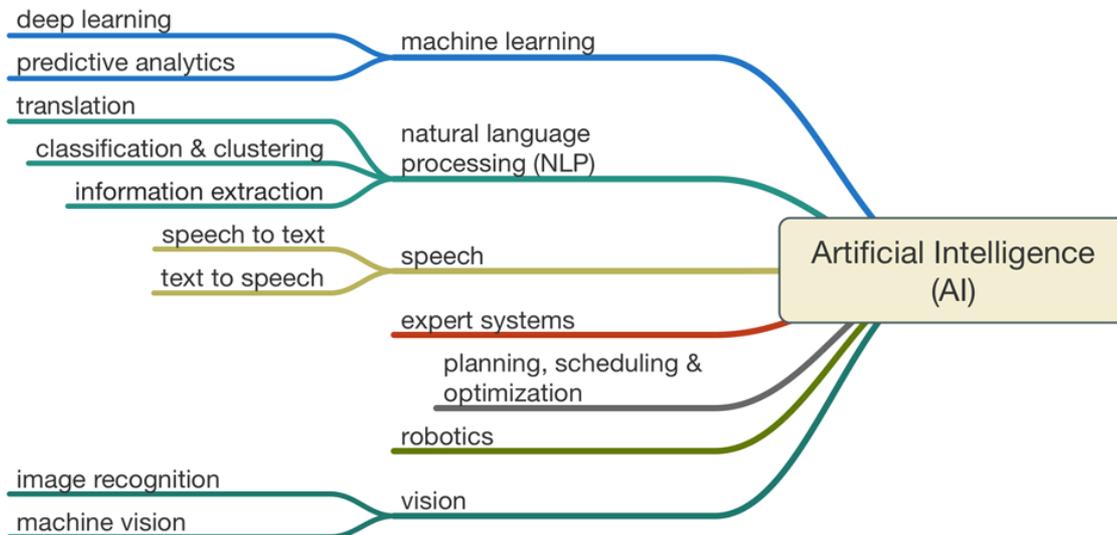
Introduction

- Artificial Intelligence
- Machine learning
- “Deep” learning



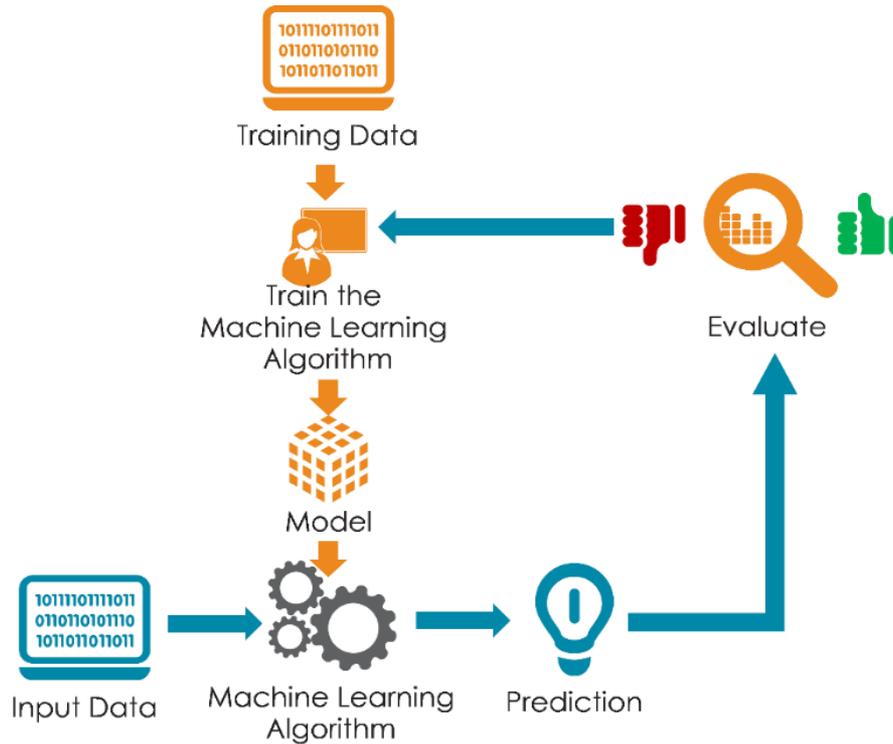
Artificial Intelligence

Artificial intelligence (AI, also machine intelligence, MI) is intelligence displayed by machines, in contrast with the natural intelligence (NI) displayed by humans and other animals.



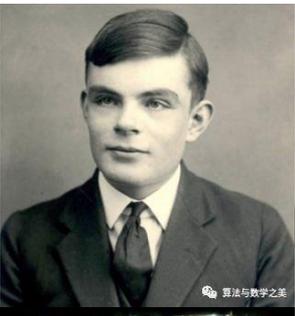
Machine Learning

Machine learning is the science of getting computers to act without being explicitly programmed.



General workflow of Machine Learning

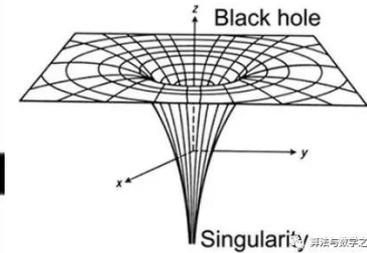
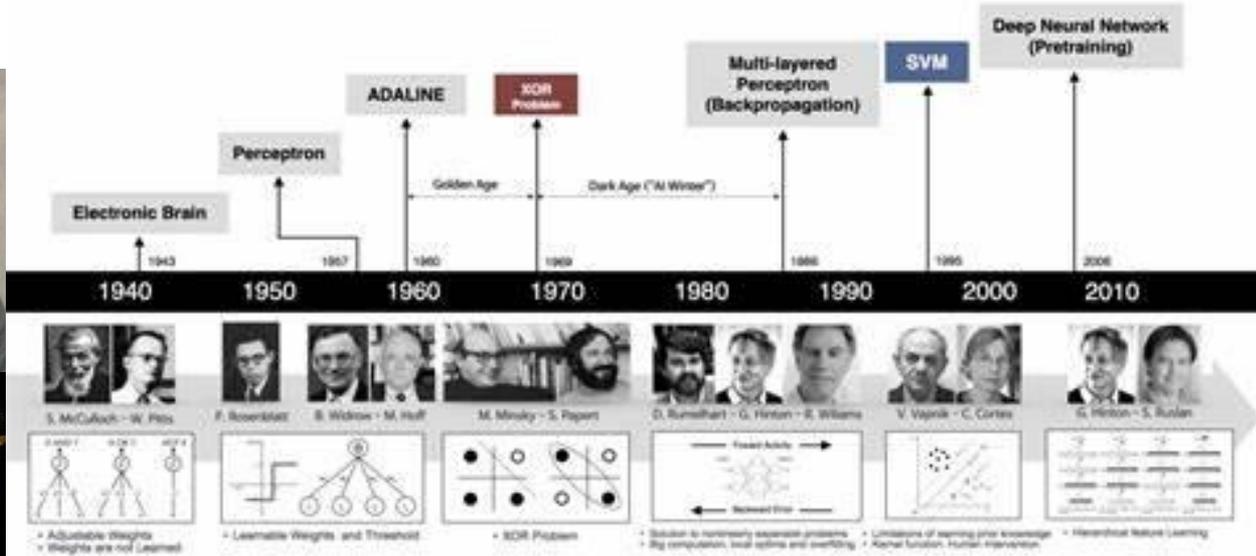
Machine Learning



算法与数学之美



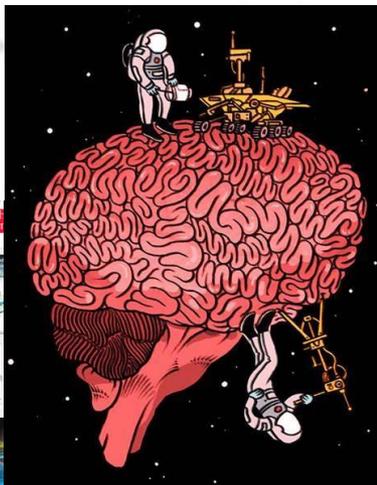
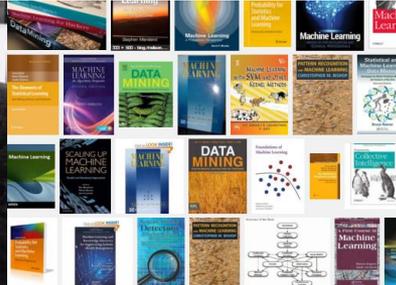
算法与数学之美



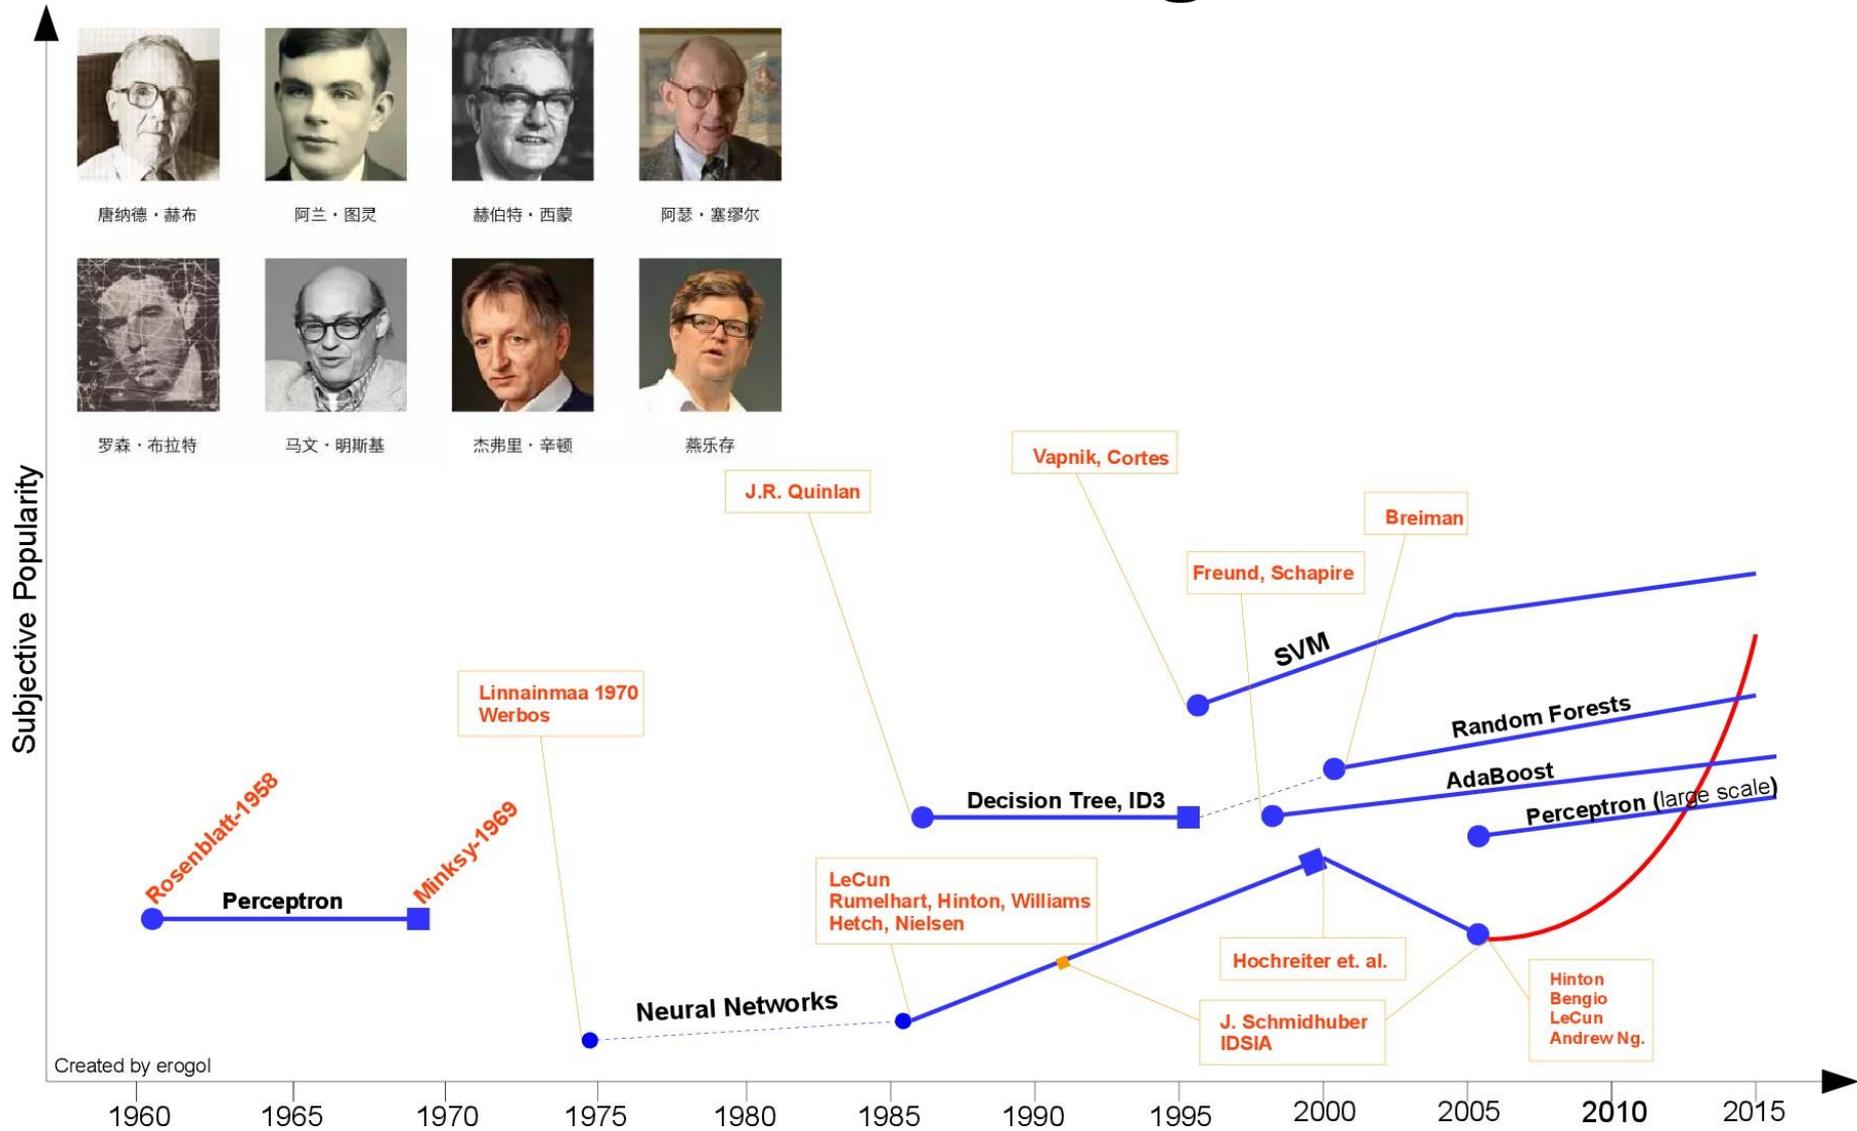
算法与数学之美



算法与数学之美



Machine Learning



唐纳德·赫布



阿兰·图灵



赫伯特·西蒙



阿瑟·塞缪尔



罗森·布拉特



马文·明斯基

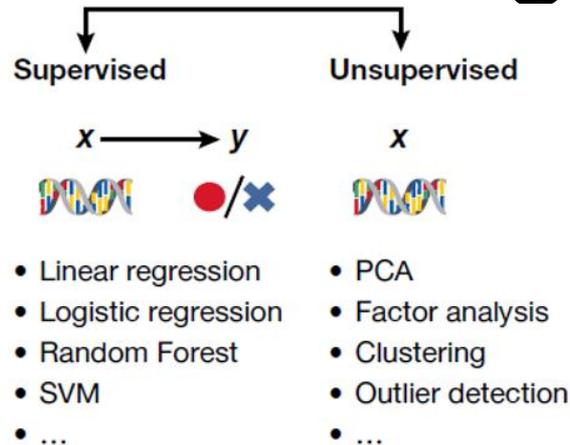


杰弗里·辛顿

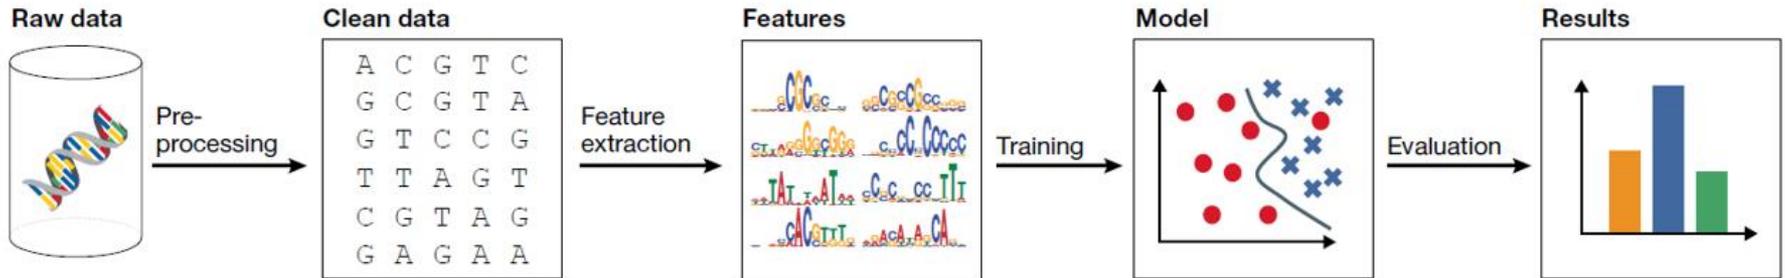


燕乐存

Machine Learning in biology



Machine learning algorithm: supervised and unsupervised



An exemplification of Machine learning in biology : classification model

Deep learning

Deep learning is a part of machine learning.

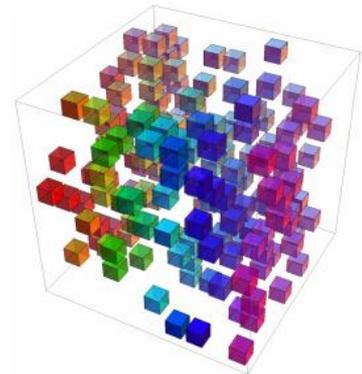
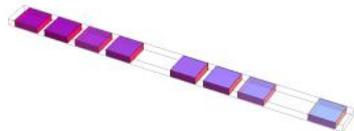
Deep:

Complex Model : Multi-layer characters, Many parameters (**Curse of dimensionality**)

Training data & Testing data : a big volume of data

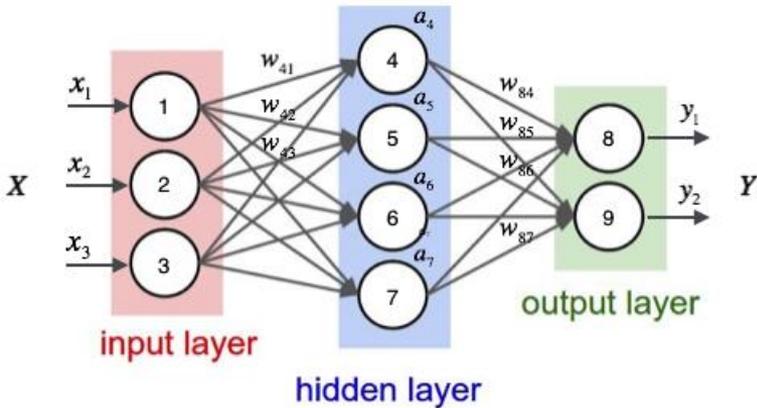
Adjustment: A single model construction could cost one week

几乎所有深度学习算法都可以被描述为一个简单配方：
特定数据集、代价函数、优化过程和模型

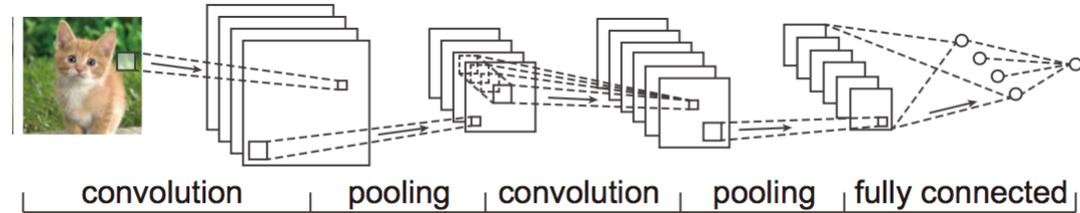


Deep learning

Now, neural network and convolution neural network models that work best



Neural Network

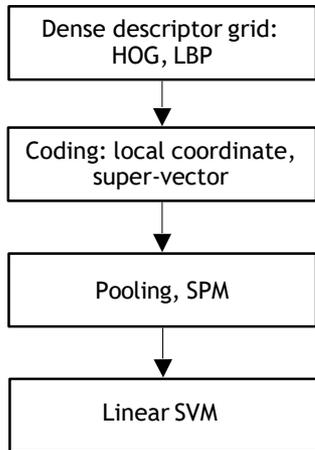


Convolution Neural Network

IMAGENET Large Scale Visual Recognition Challenge

Year 2010

NEC-UIUC

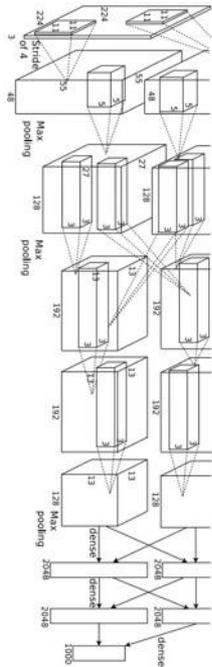


[Lin CVPR 2011]

[Lion image](#) by Swissfrog is licensed under [CC BY 3.0](#)

Year 2012

SuperVision

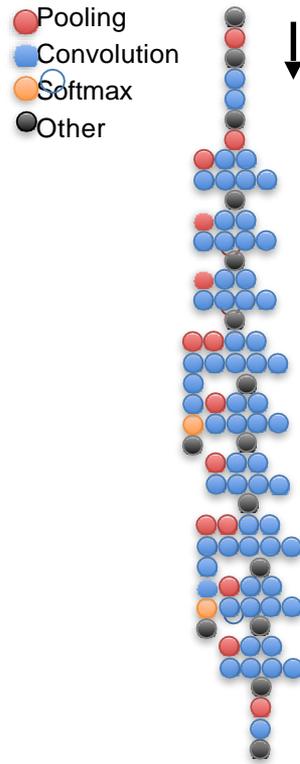


[Krizhevsky NIPS 2012]

Figure copyright Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton, 2012. Reproduced with permission.

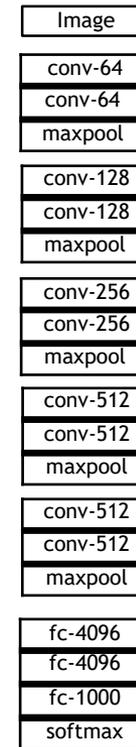
Year 2014

GoogLeNet



[Szegedy arxiv 2014]

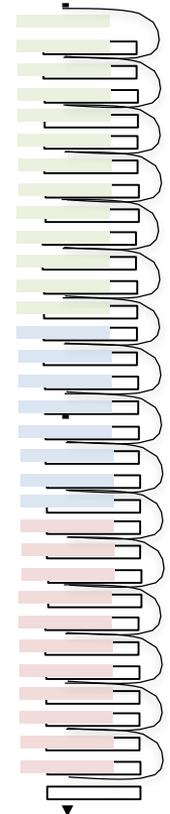
VGG



[Simonyan arxiv 2014]

Year 2015

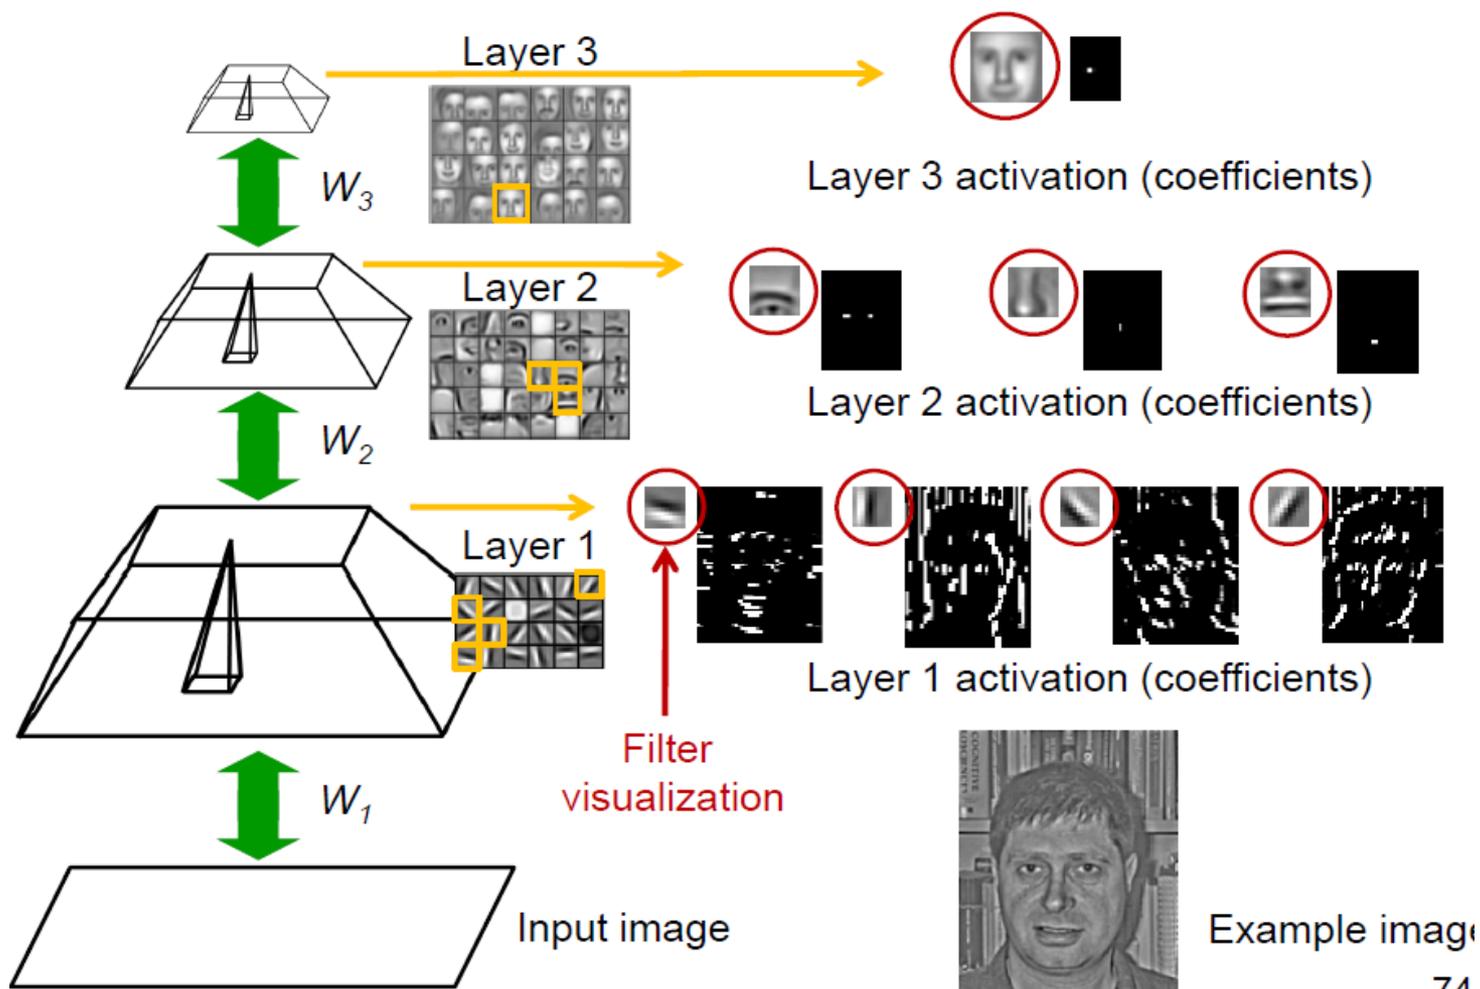
MSRA



[He ICCV 2015]

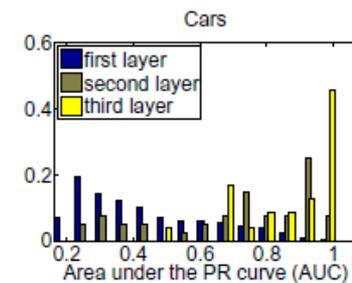
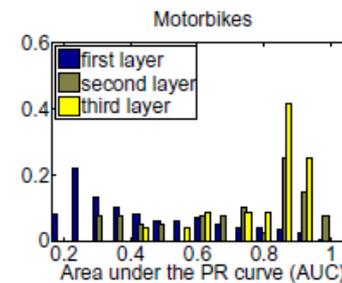
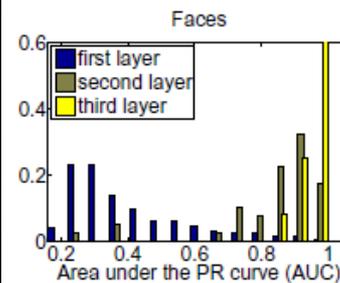
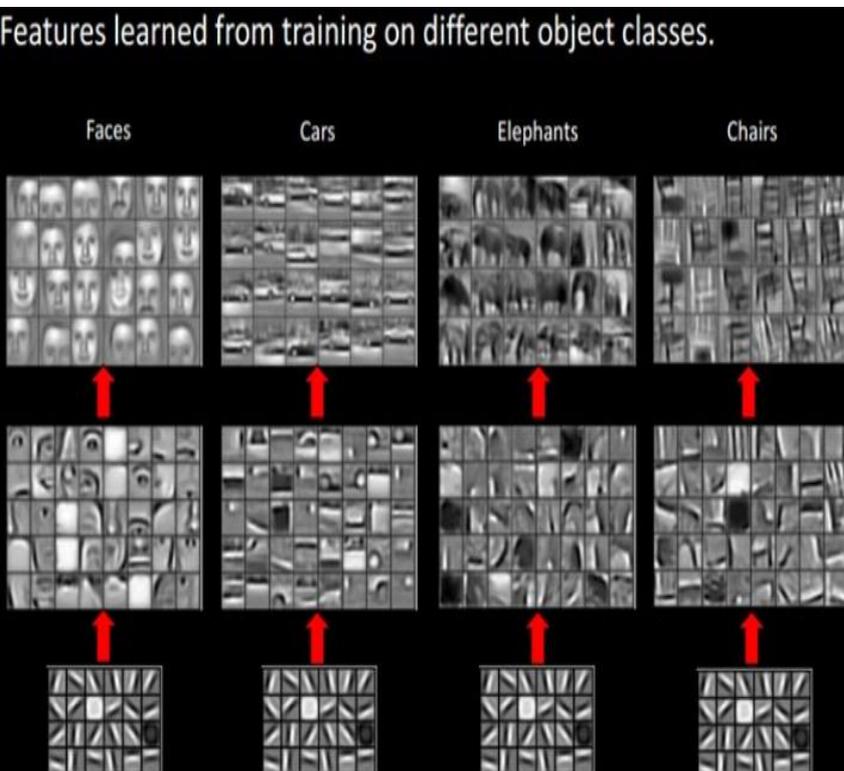
深度学习的应用

- 深度学习在图像识别上的应用



深度学习的应用

- 深度学习在图像识别上的应用

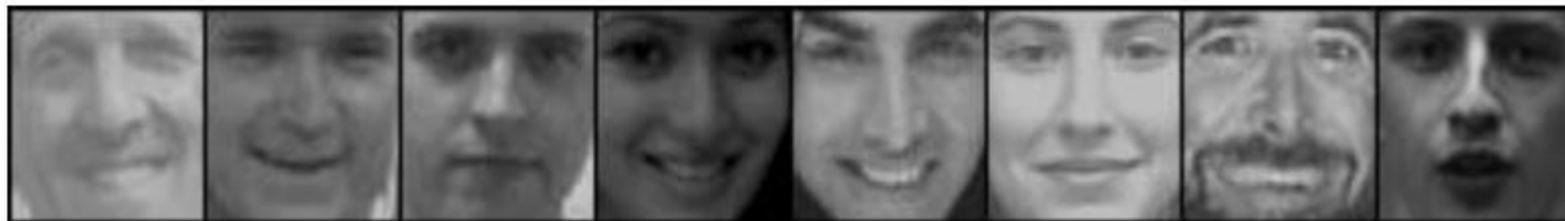


Features	Faces	Motorbikes	Cars
First layer	0.39±0.17	0.44±0.21	0.43±0.19
Second layer	0.86±0.13	0.69±0.22	0.72±0.23
Third layer	0.95±0.03	0.81±0.13	0.87±0.15

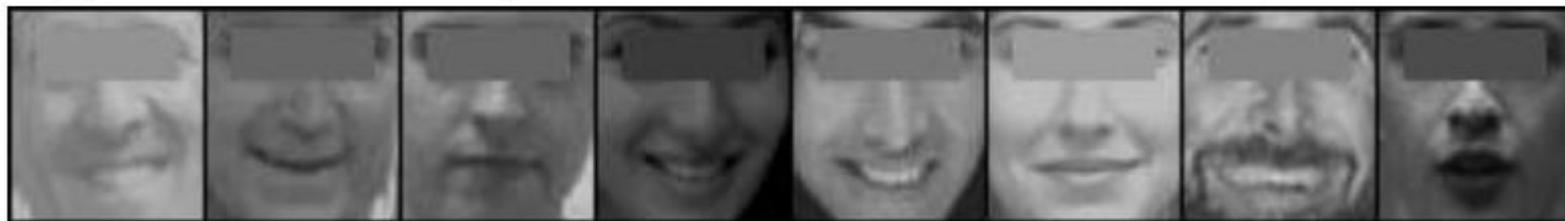
深度学习的应用

- 深度学习在图像识别上的应用

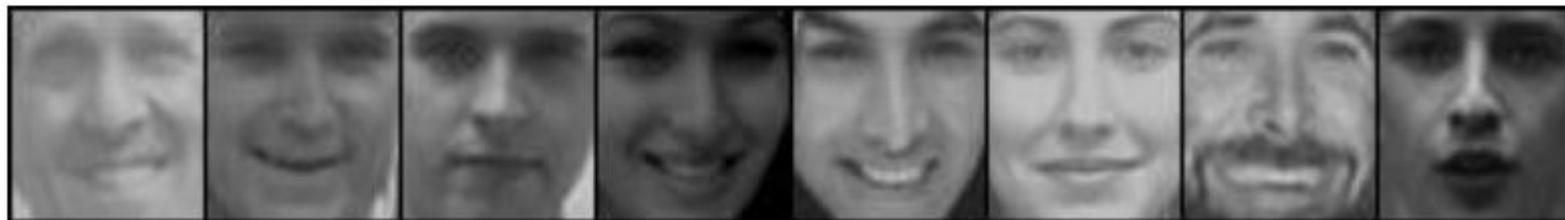
originals



Type 1 occlusion: eyes

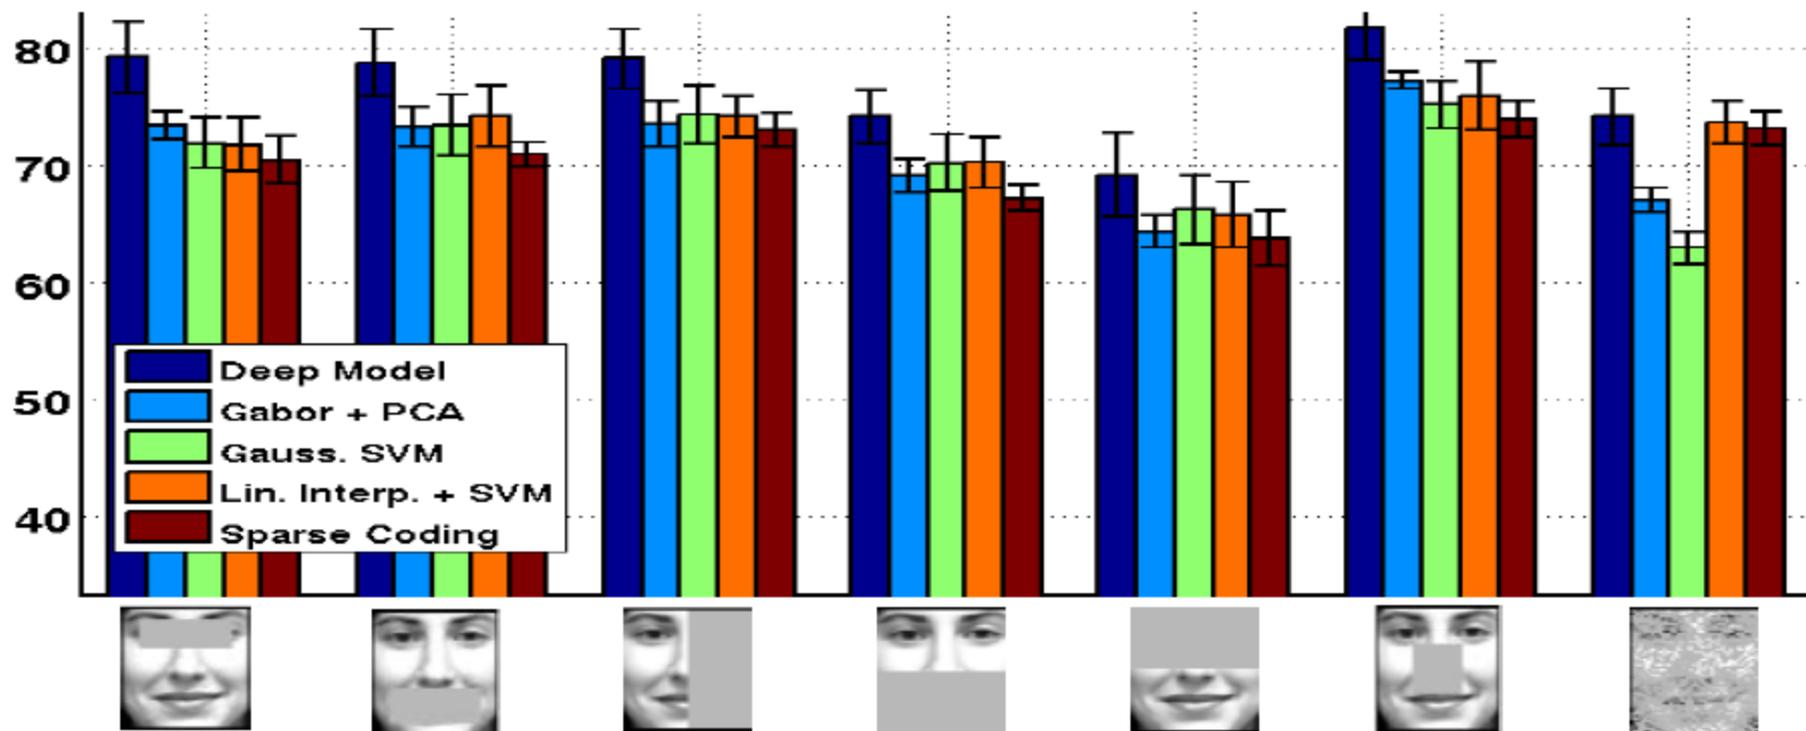


Restored images



深度学习的应用

- 深度学习在图像识别上的应用



深度学习的应用

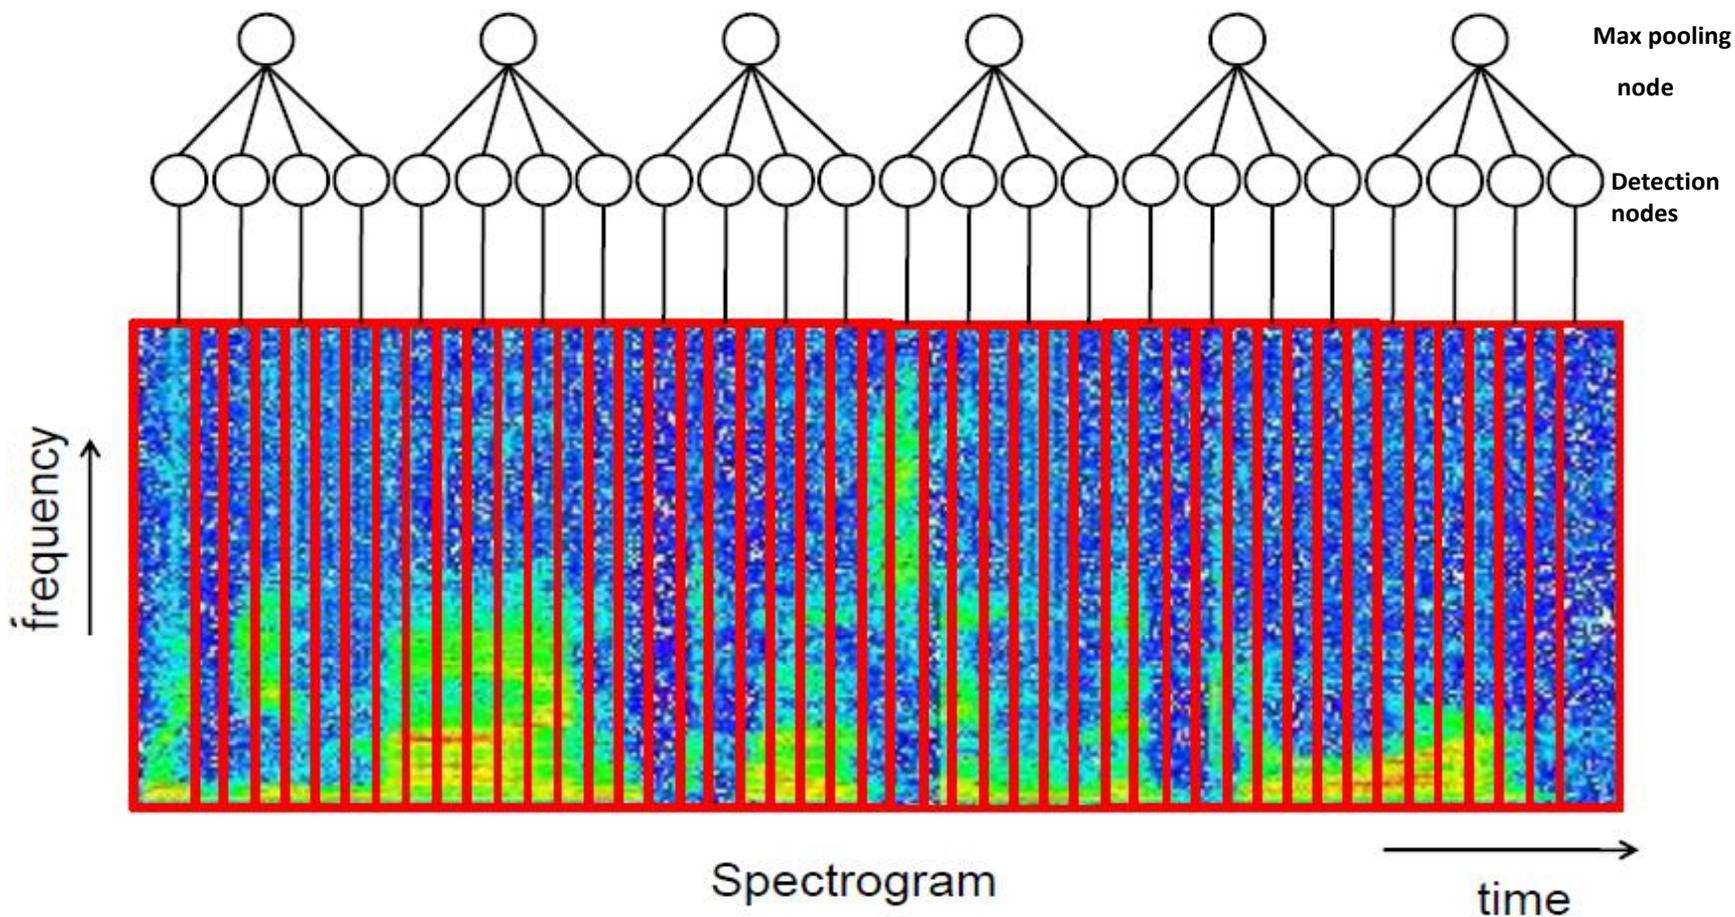
- 深度学习在图像识别上的应用

< Caltech 256 >

# of training images	30	60
Griffin et al. [2]	34.10	-
vanGemert et al., PAMI 2010	27.17	-
ScSPM [Yang et al., CVPR 2009]	34.02	40.14
LLC [Wang et al., CVPR 2010]	41.19	47.68
Sparse CRBM [Sohn et al., ICCV 2011]	42.05	47.94

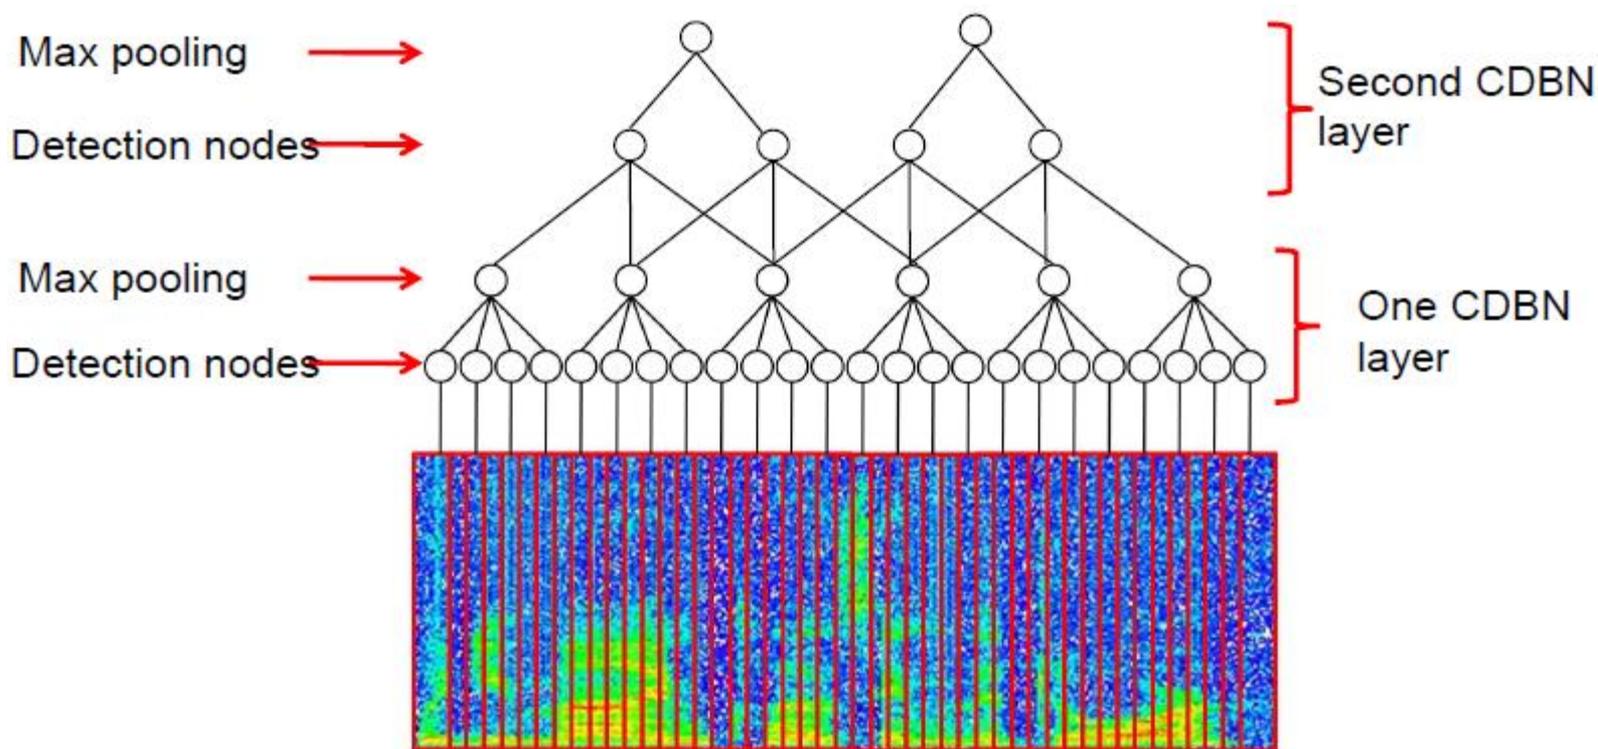
深度学习的应用

- 深度学习在音频识别上的应用



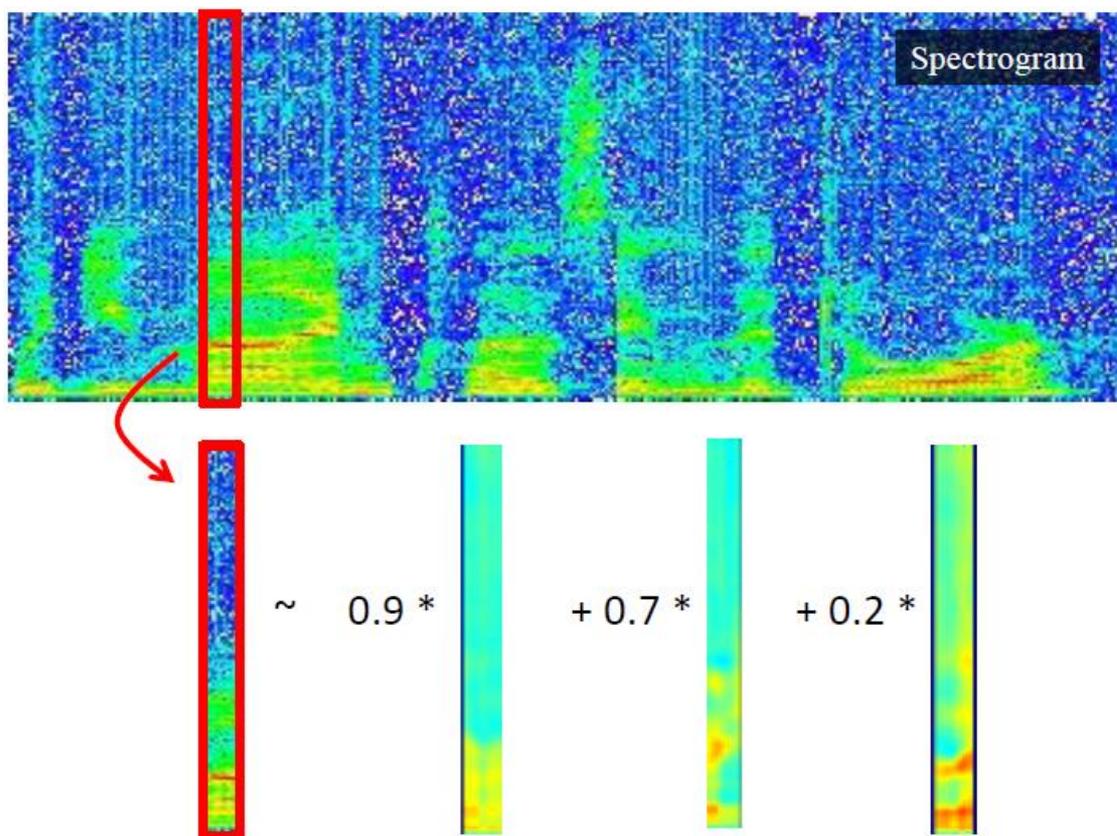
深度学习的应用

- 深度学习在音频识别上的应用



深度学习的应用

- 深度学习在音频识别上的应用



[Lee, Largman, Pham, Ng, NIPS 2009]

深度学习的应用

- 深度学习在音频识别上的应用
- Speaker identification

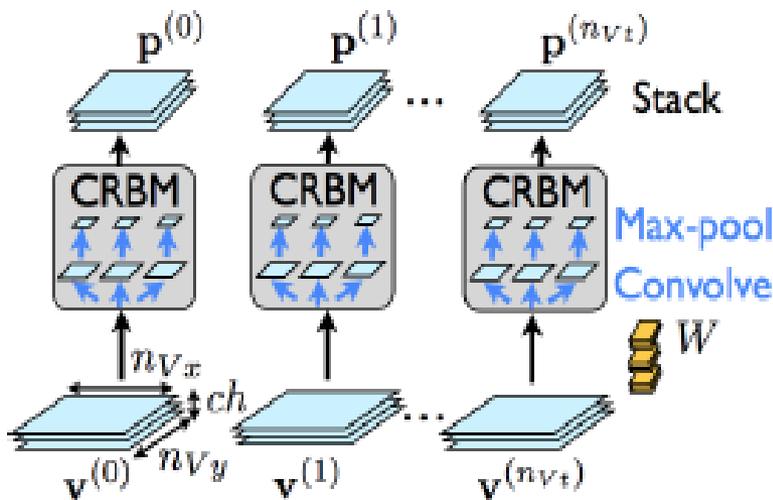
TIMIT Speaker identification	Accuracy
Prior art (Reynolds, 1995)	99.7%
Convolutional DBN	100.0%

- Phone classification

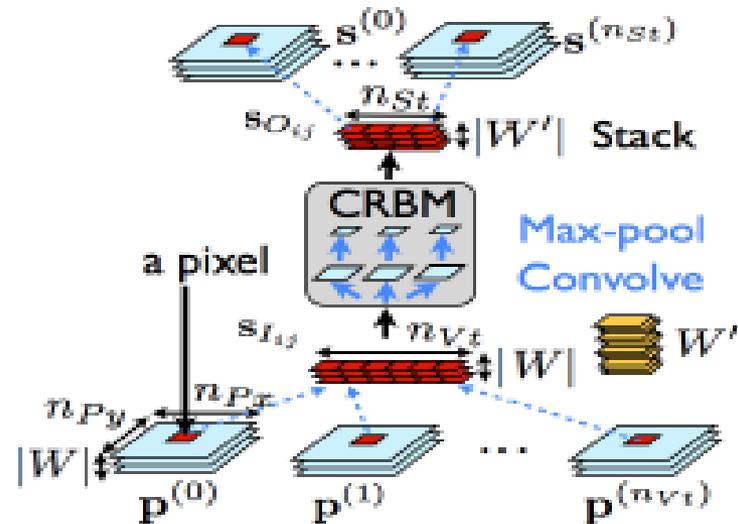
TIMIT Phone classification	Accuracy
Clarkson et al. (1999)	77.6%
Petrov et al. (2007)	78.6%
Sha & Saul (2006)	78.9%
Yu et al. (2009)	79.2%
Convolutional DBN	80.3%
Transformation-invariant RBM (Sohn et al., ICML 2012)	81.5%

深度学习的应用

- 深度学习在视频识别上的应用



Spatial pooling layer



Temporal pooling layer

深度学习的应用

- 深度学习在视频识别上的应用

Video Activity recognition (Hollywood 2 benchmark)

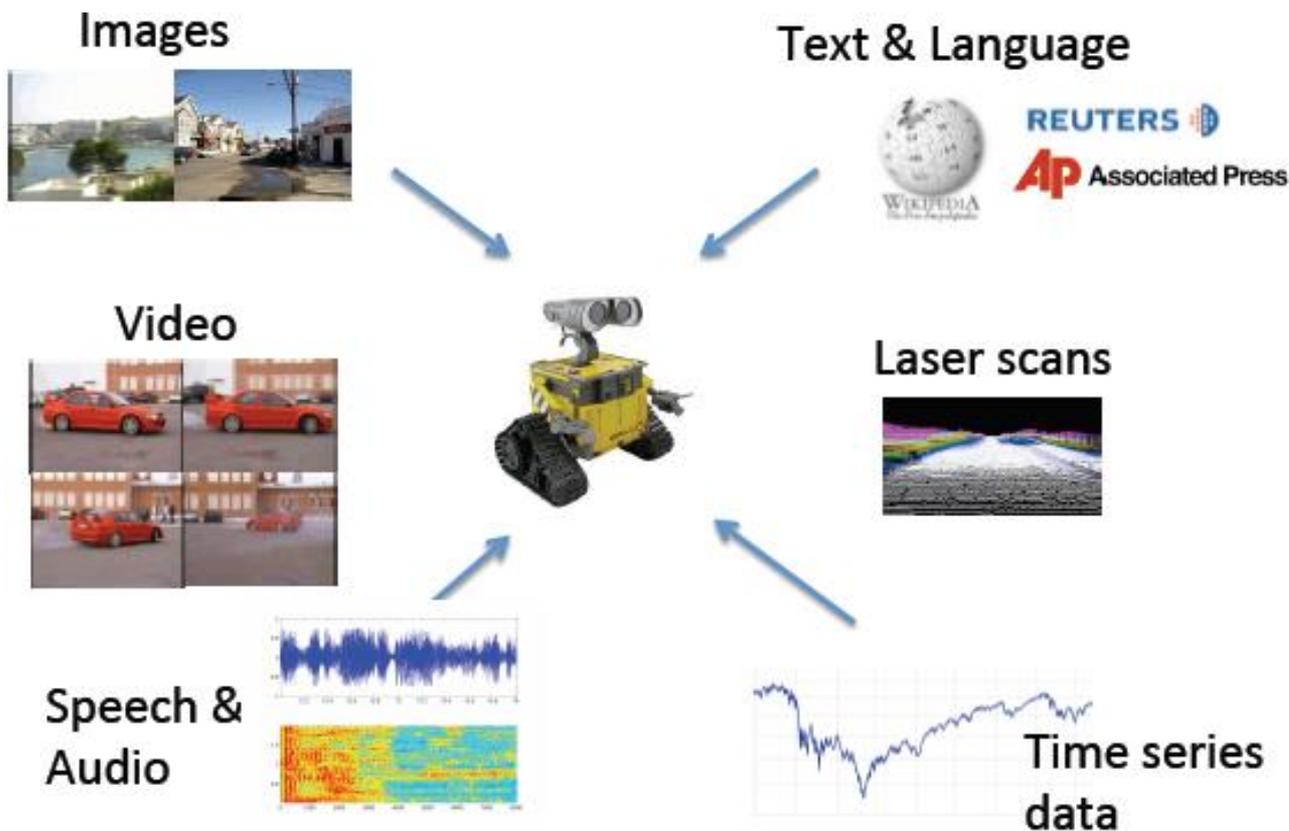


Method	Accuracy
Hessian + ESURF [Williems et al 2008]	38%
Harris3D + HOG/HOF [Laptev et al 2003, 2004]	45%
Cuboids + HOG/HOF [Dollar et al 2005, Laptev 2004]	46%
Hessian + HOG/HOF [Laptev 2004, Williems et al 2008]	46%
Dense + HOG / HOF [Laptev 2004]	47%
Cuboids + HOG3D [Klaser 2008, Dollar et al 2005]	46%
Unsupervised feature learning (our method)	52%

Unsupervised feature learning significantly improves on the previous state-of-the-art.

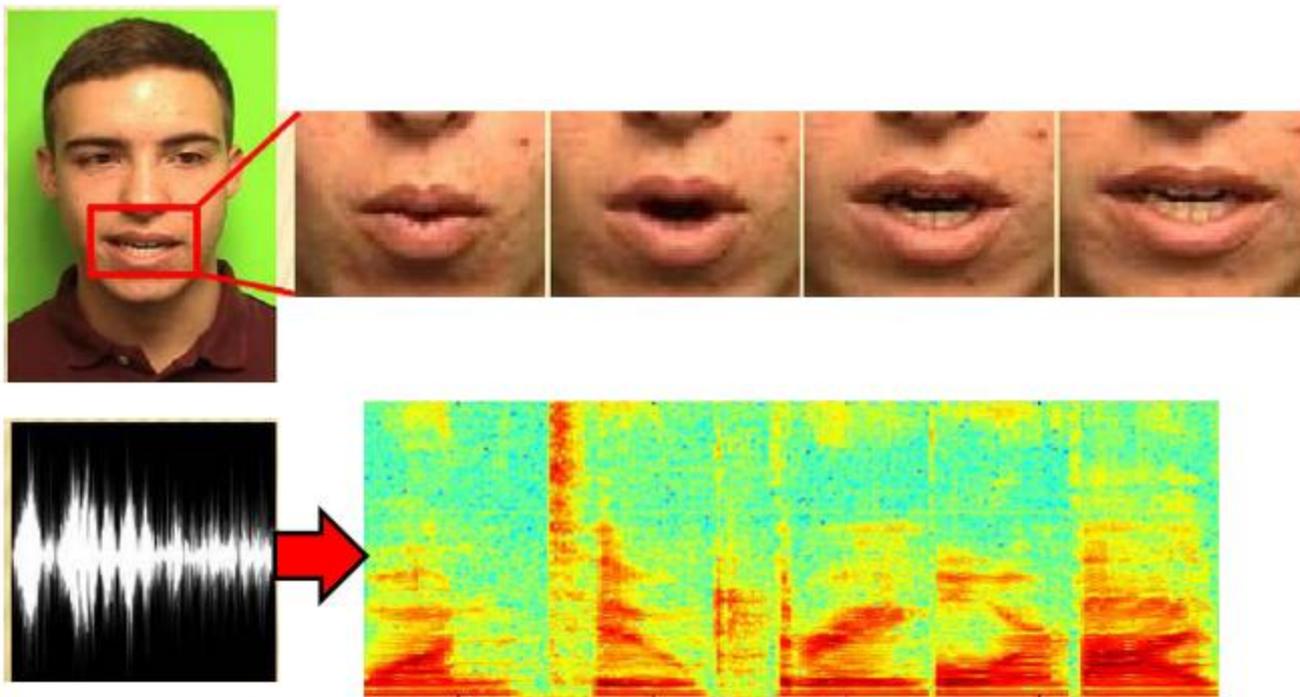
深度学习的应用

- 深度学习在多模态学习中的应用



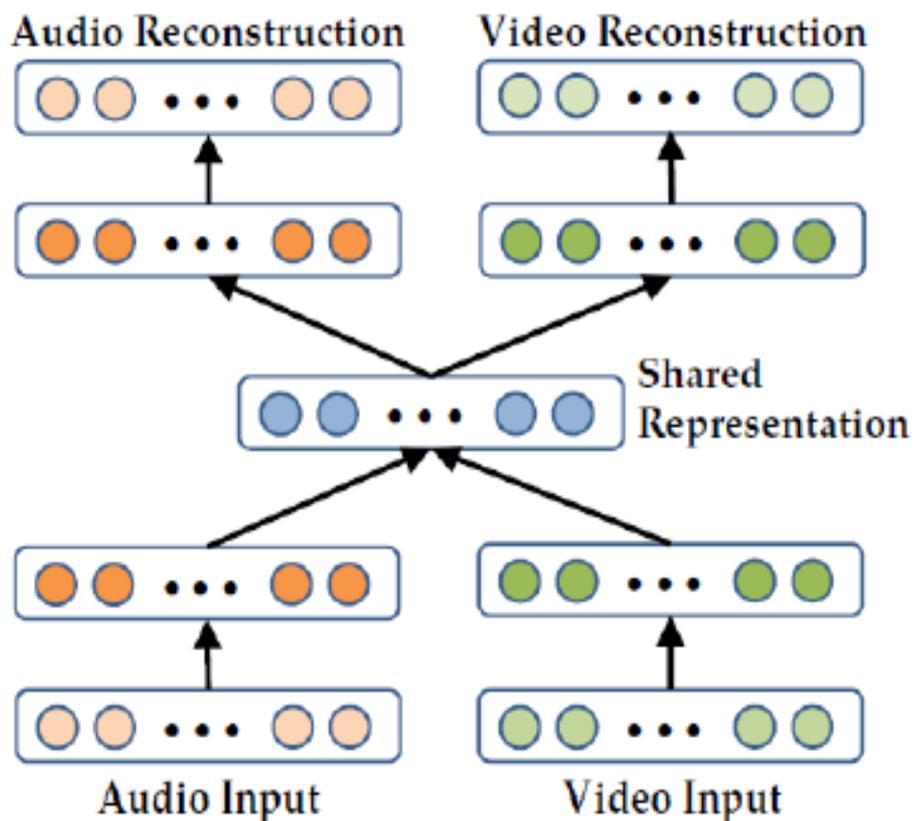
深度学习的应用

- 深度学习在多模态学习中的应用



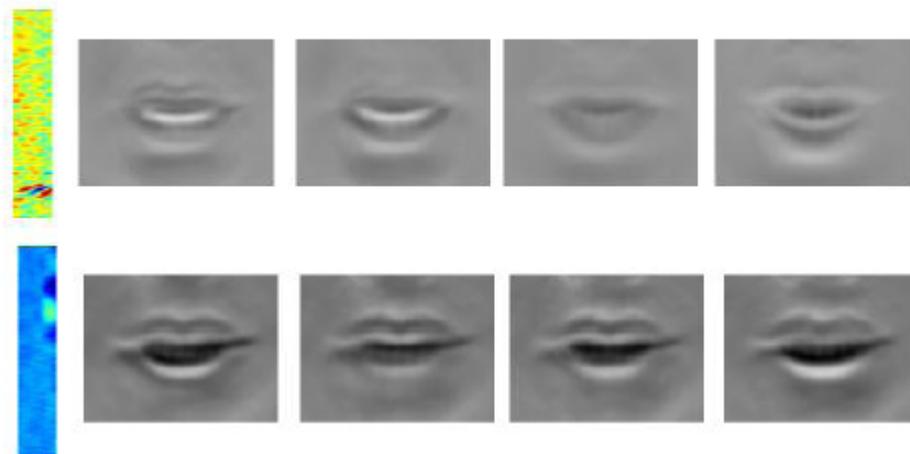
深度学习的应用

- 深度学习在多模态学习中的应用



深度学习的应用

- 深度学习在多模态学习中的应用
 - Visualization of learned filters



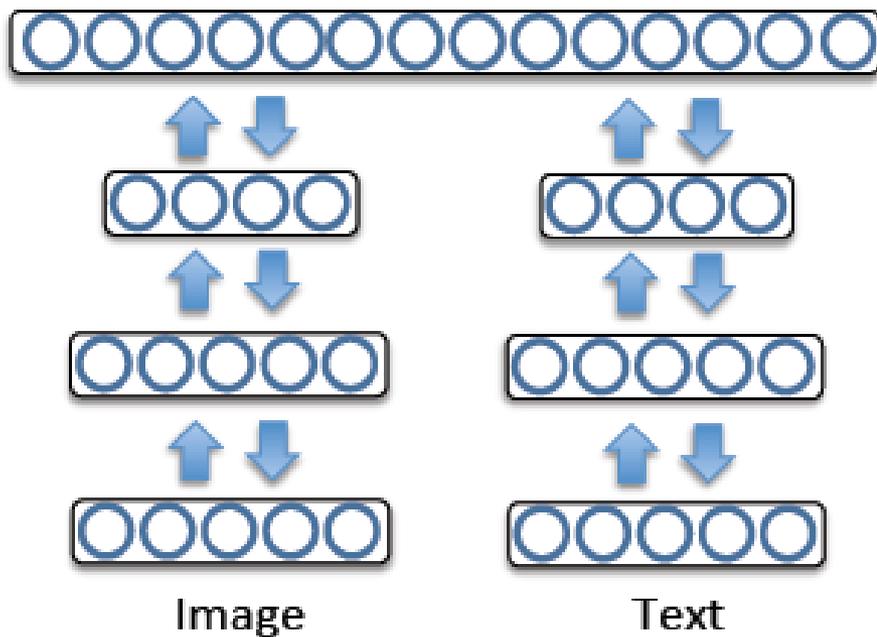
Audio(spectrogram) and Video features learned over 100ms windows

- Results: AVLetters Lip reading dataset

Method	Accuracy
Prior art (Zhao et al., 2009)	58.9%
Multimodal deep autoencoder (Ngiam et al., 2011)	65.8%

深度学习的应用

- 深度学习在多模态学习中的应用

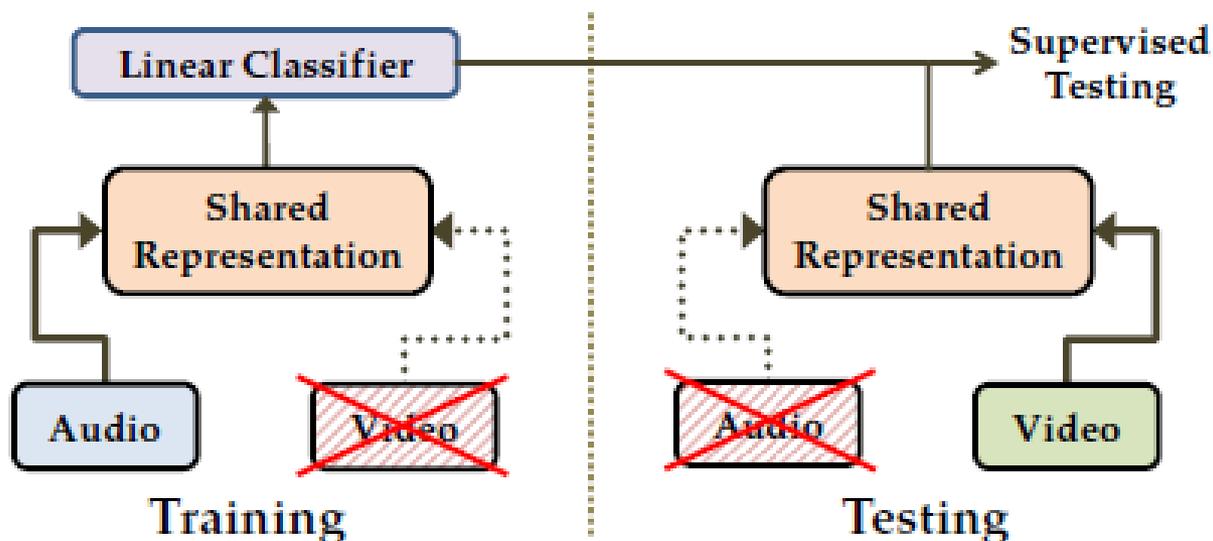


- Multimodal Inputs (images + text), 38 classes.

Learning Algorithm	Mean Average Precision
Image-text SVM	0.475
Image-text LDA	0.492
Multimodal DBN	0.566

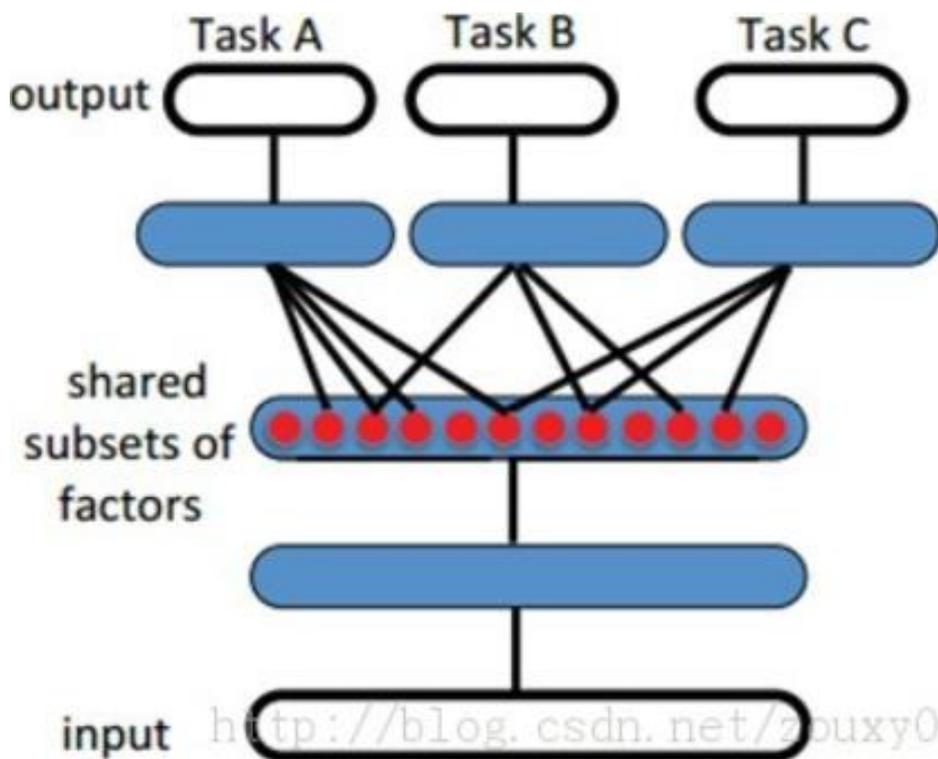
深度学习的应用

- 深度学习在多模态学习中的应用



深度学习的应用

- 深度学习在多任务学习中的应用



input <http://blog.csdn.net/zouxy0>

深度学习的应用

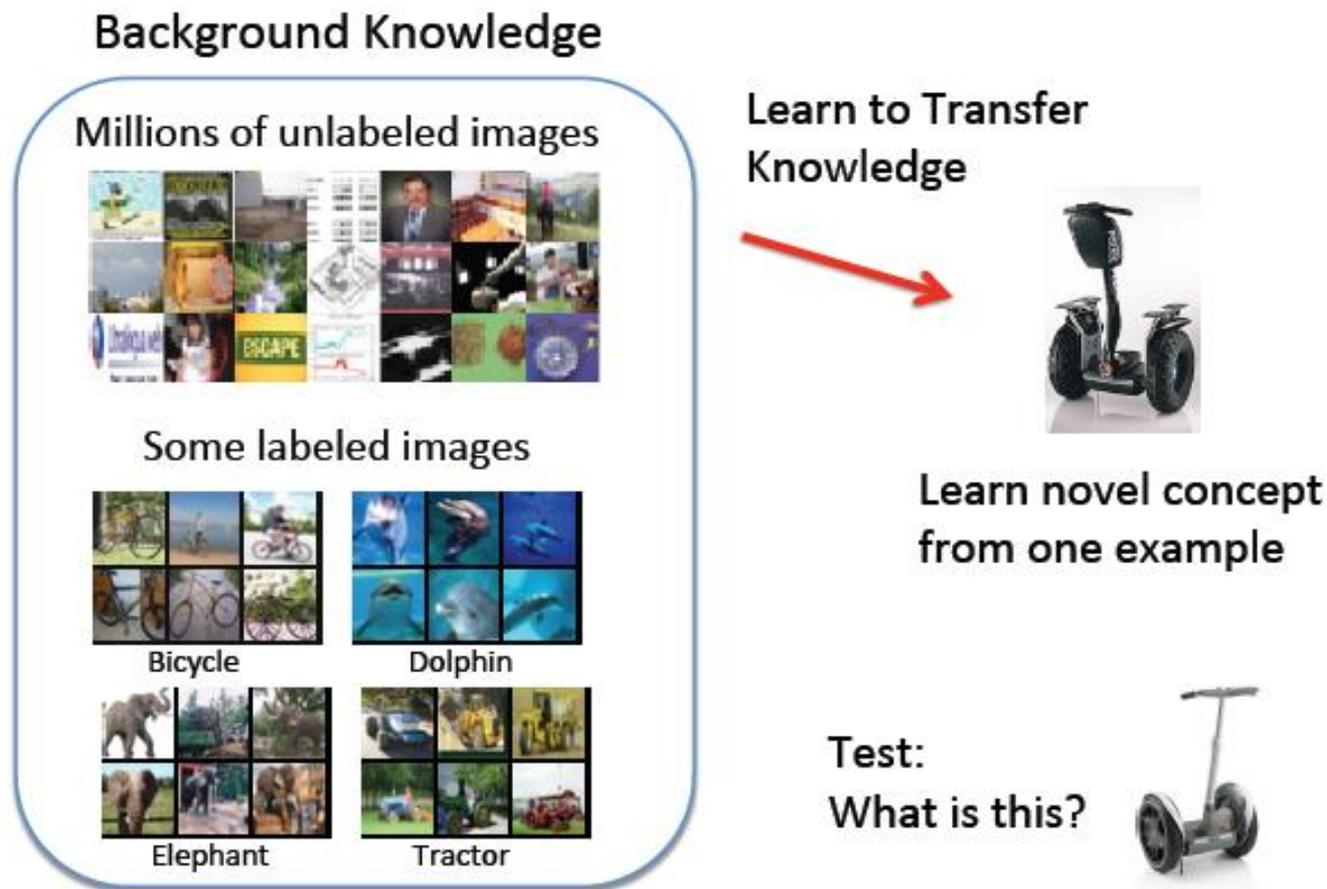
- 深度学习在多任务学习中的应用

- ✓ 在深度学习模型中，对于相关任务的联合学习，往往会取得较好的特征表达；
- ✓ 多任务联合学习，能够增强损失函数的作用效能；

比如：单独进行人脸检测会比较难（光照、遮挡等因素），但是当人脸检测与人脸识别这两个相关的任务联合学习时，人脸检测的难度反而降低了。

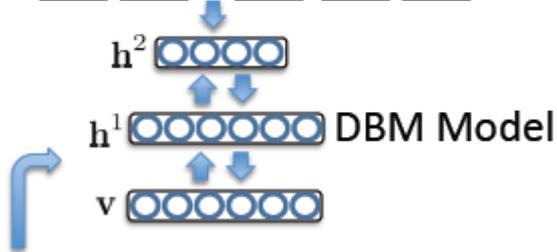
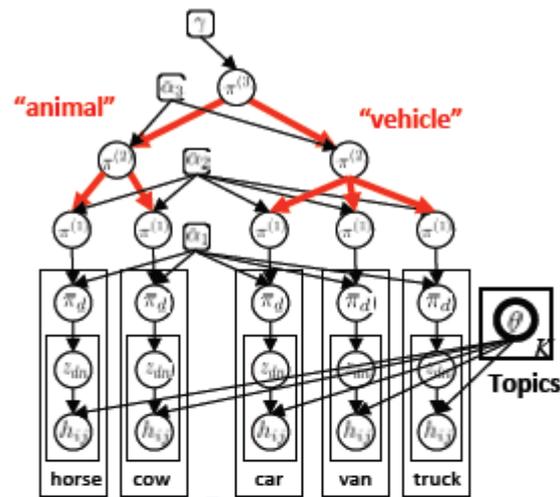
深度学习的应用

- 基于深度学习的迁移学习应用

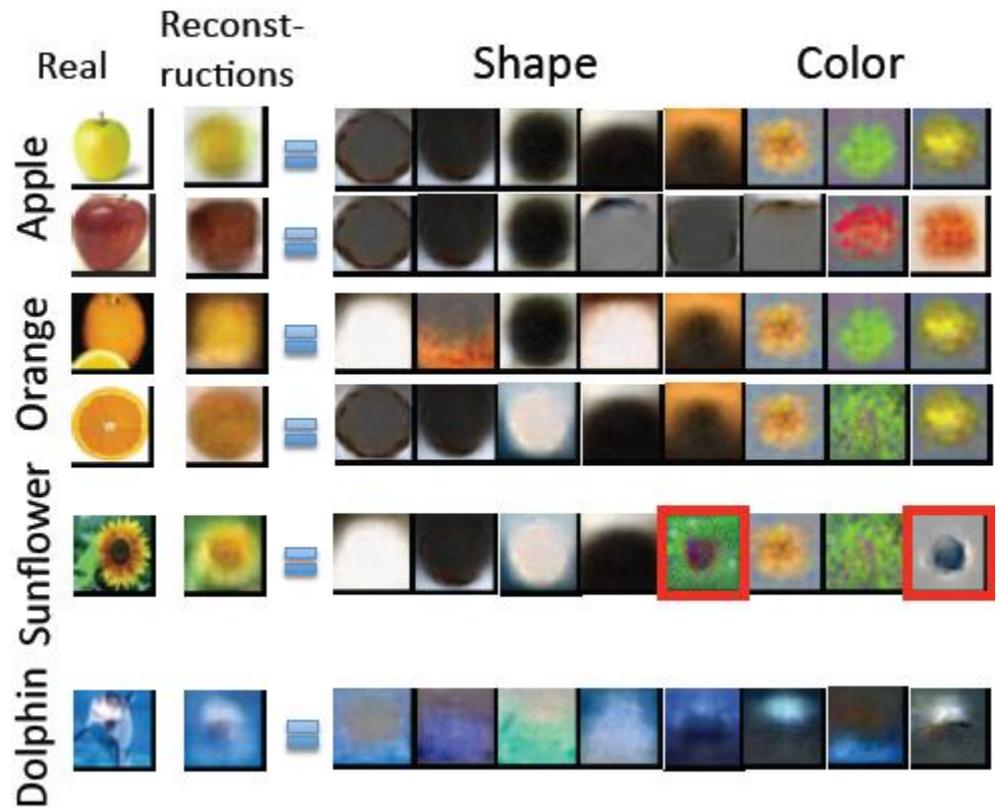


深度学习的应用

- 基于深度学习的迁移学习应用

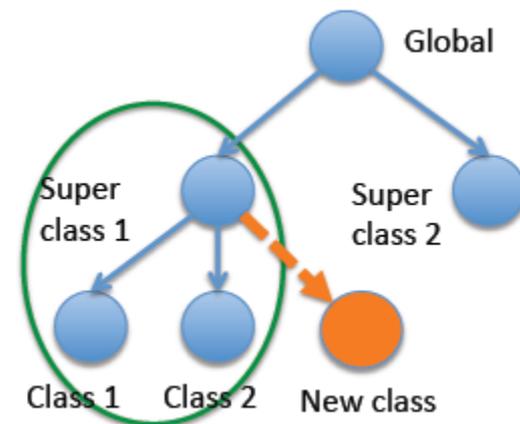
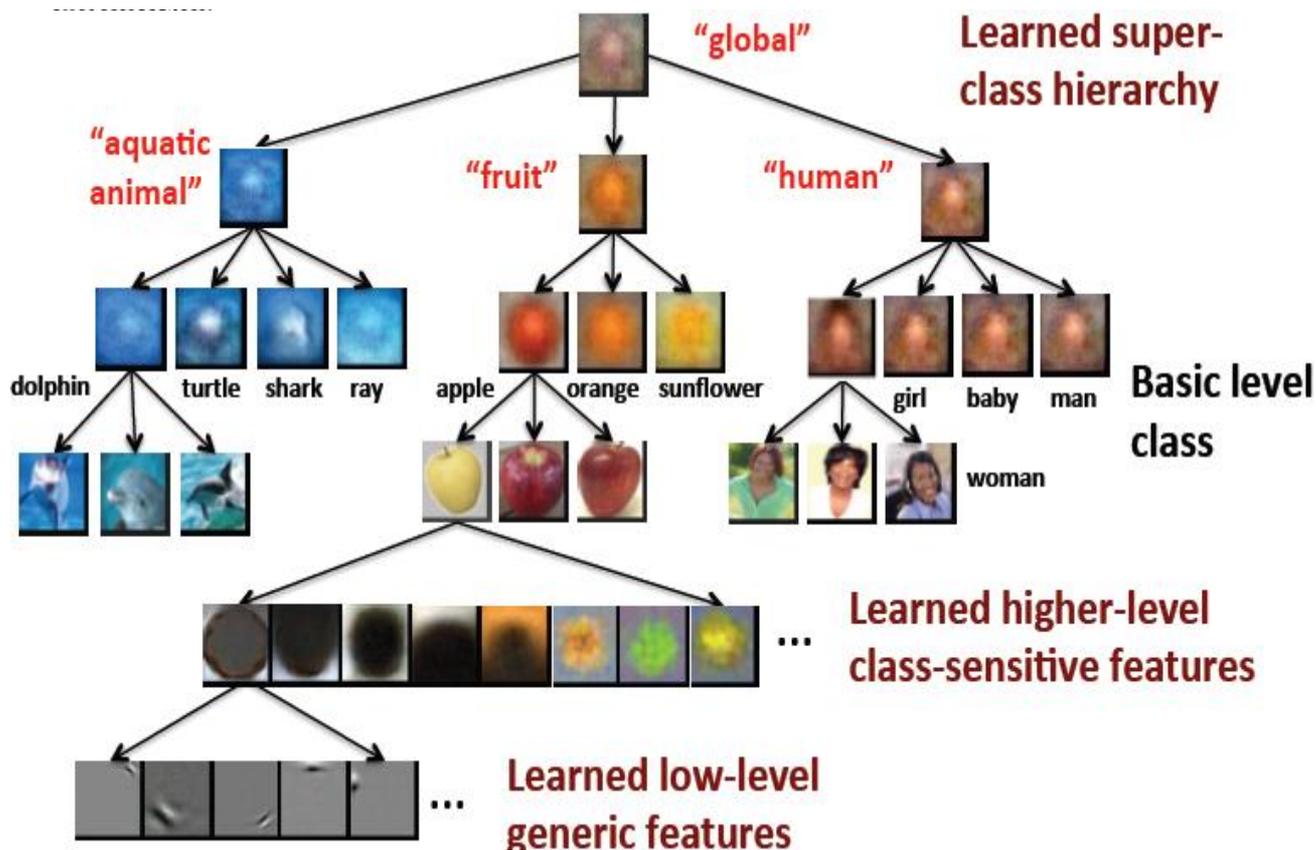


Low-level features:
replace GIST, SIFT



深度学习的应用

- 基于深度学习的迁移学习应用



深度学习的应用

- 深度学习在大尺度数据集上的应用

- 大尺度数据集:

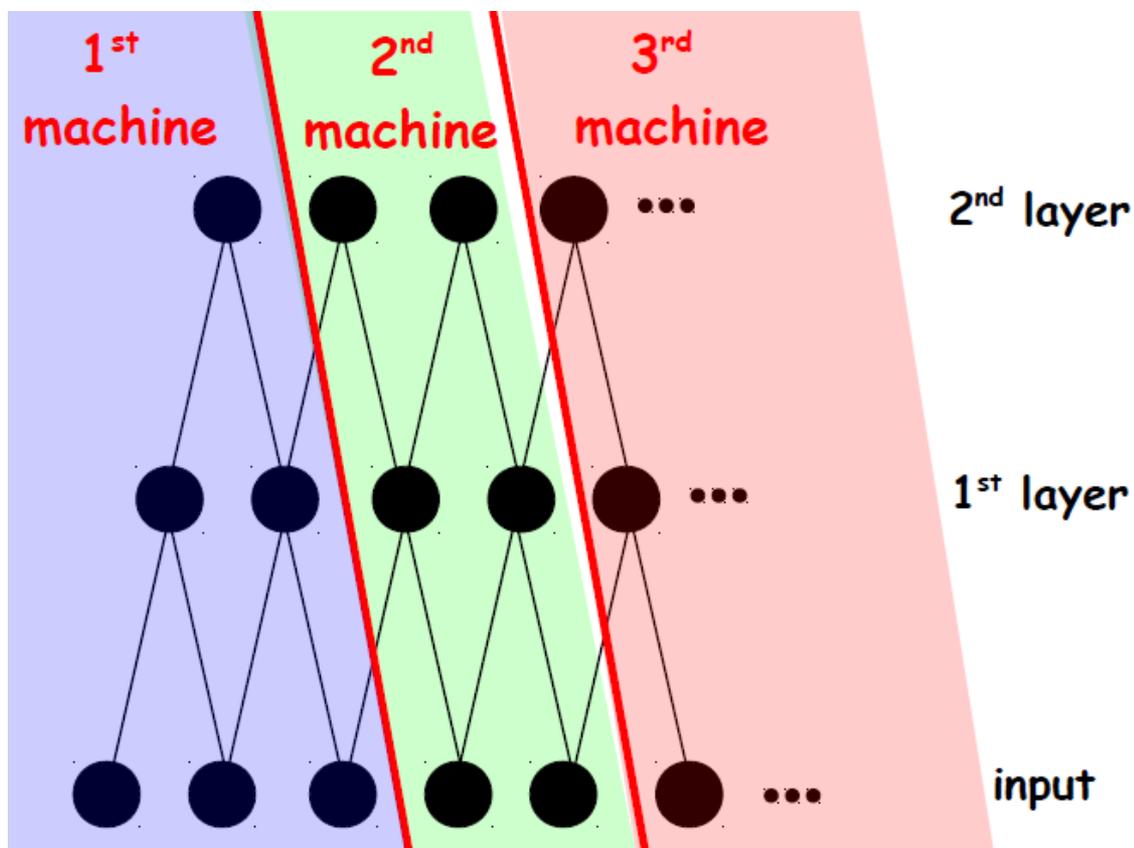
- ✓ 样本总数 > 100M,

- ✓ 类别总数 > 10K,

- ✓ 特征维度 > 10K

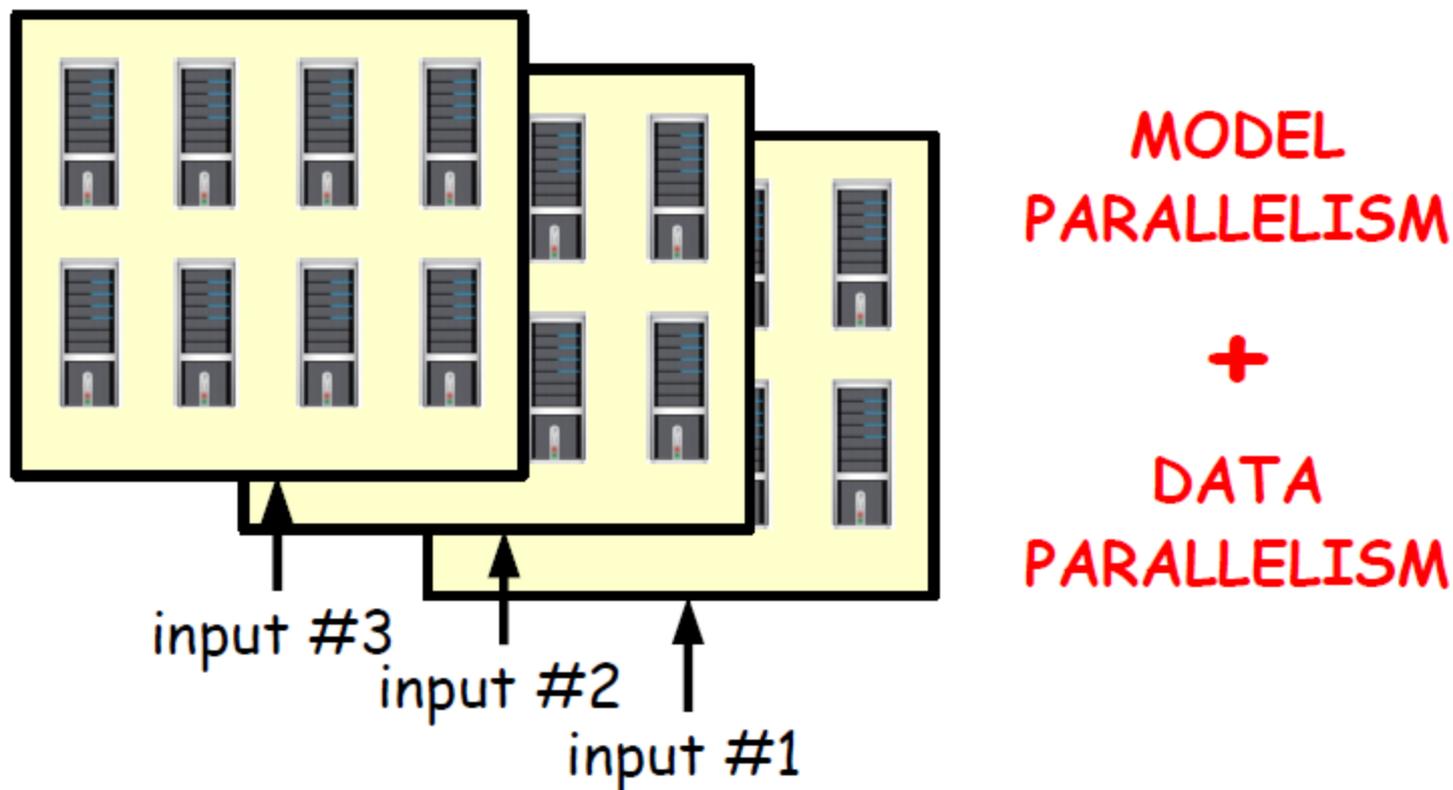
深度学习的应用

- 深度学习在大尺度数据集上的应用



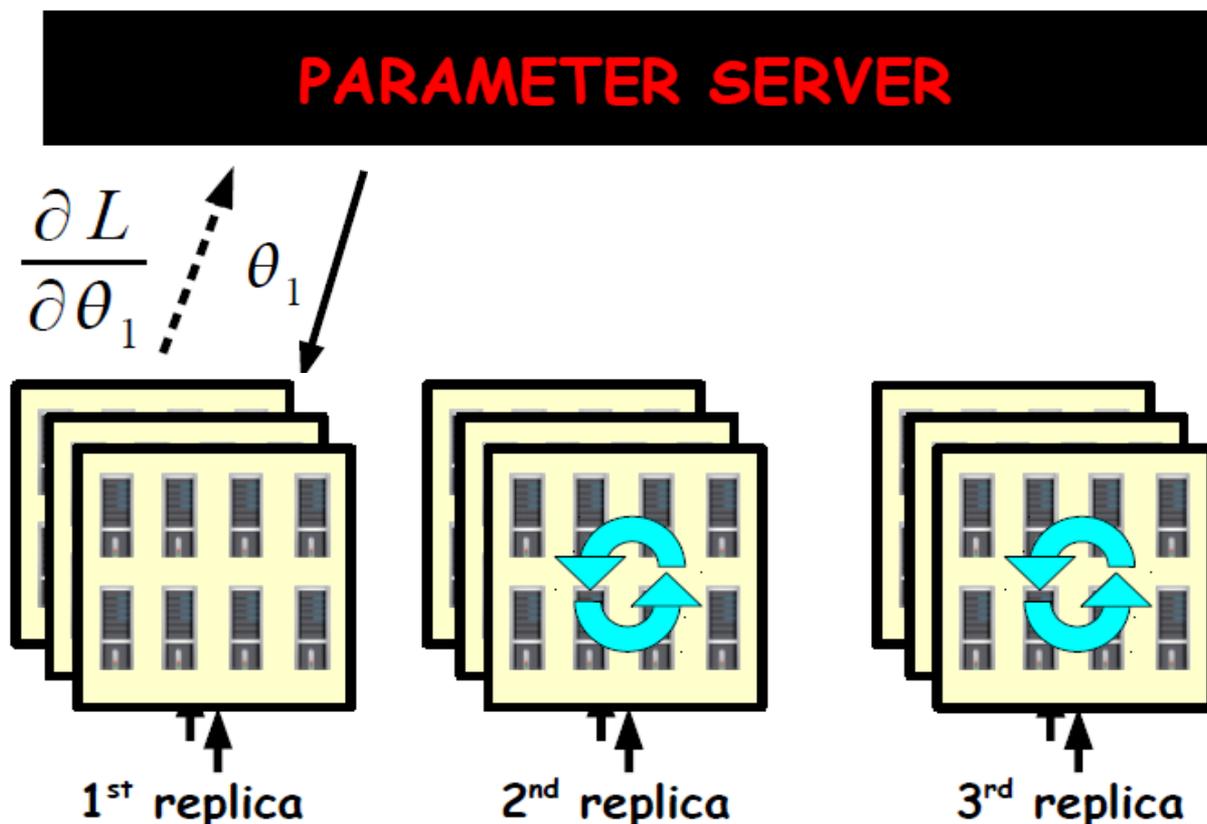
深度学习的应用

- 深度学习在大尺度数据集上的应用



深度学习的应用

- 深度学习在大尺度数据集上的应用



深度学习的应用

- 深度学习在大尺度数据集上的应用



IMAGENET v.2011 (16M images, 20K categories)

METHOD	ACCURACY %
Weston & Bengio 2011	9.3
Linear Classifier on deep features	13.1
Deep Net (from random)	13.6
Deep Net (from unsup.)	15.8

深度学习的应用

- 深度学习的State-of-the-art

Images		NORB Object classification	
CIFAR Object classification	Accuracy	Accuracy	Accuracy
Prior art (Ciresan et al., 2011)	80.5%	Prior art (Scherer et al., 2010)	94.4%
Stanford Feature learning	82.0%	Stanford Feature learning	95.0%

Video		YouTube	
Hollywood2 Classification	Accuracy	Accuracy	Accuracy
Prior art (Laptev et al., 2004)	48%	Prior art (Liu et al., 2009)	71.2%
Stanford Feature learning	53%	Stanford Feature learning	75.8%

KTH		UCF	
Accuracy	Accuracy	Accuracy	Accuracy
Prior art (Wang et al., 2010)	92.1%	Prior art (Wang et al., 2010)	85.6%
Stanford Feature learning	93.9%	Stanford Feature learning	86.5%

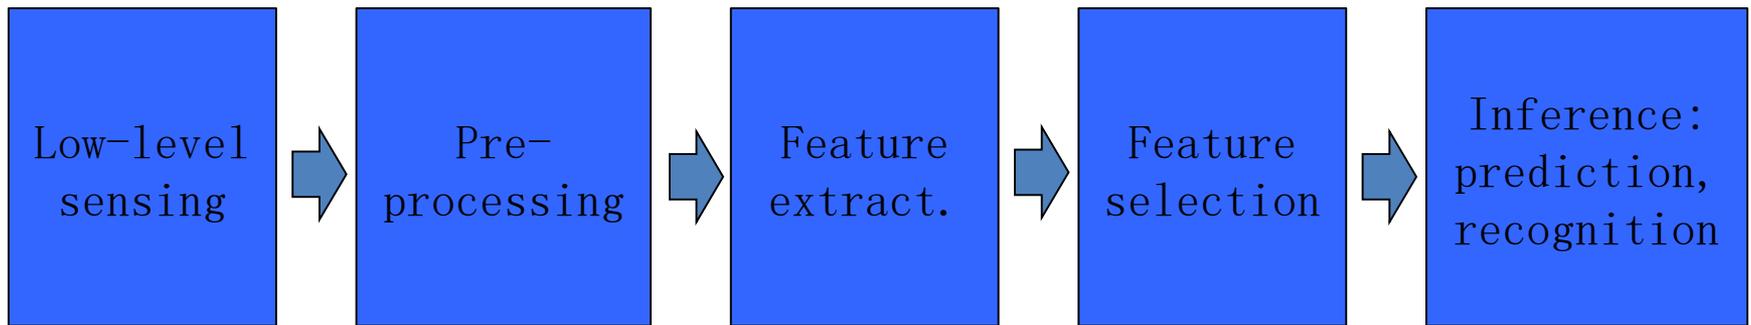
Text/NLP		Sentiment (MR/MPQA data)	
Paraphrase detection	Accuracy	Accuracy	Accuracy
Prior art (Das & Smith, 2009)	76.1%	Prior art (Nakagawa et al., 2010)	77.3%
Stanford Feature learning	76.4%	Stanford Feature learning	77.7%

Multimodal (audio/video)	
AVLetters Lip reading	Accuracy
Prior art (Zhao et al., 2009)	58.9%
Stanford Feature learning	65.8%

Other unsupervised feature learning records:
Pedestrian detection (Yann LeCun)
Speech recognition (Geoff Hinton)
PASCAL VOC object classification (Kai Yu)

动机

传统的模式识别方法：

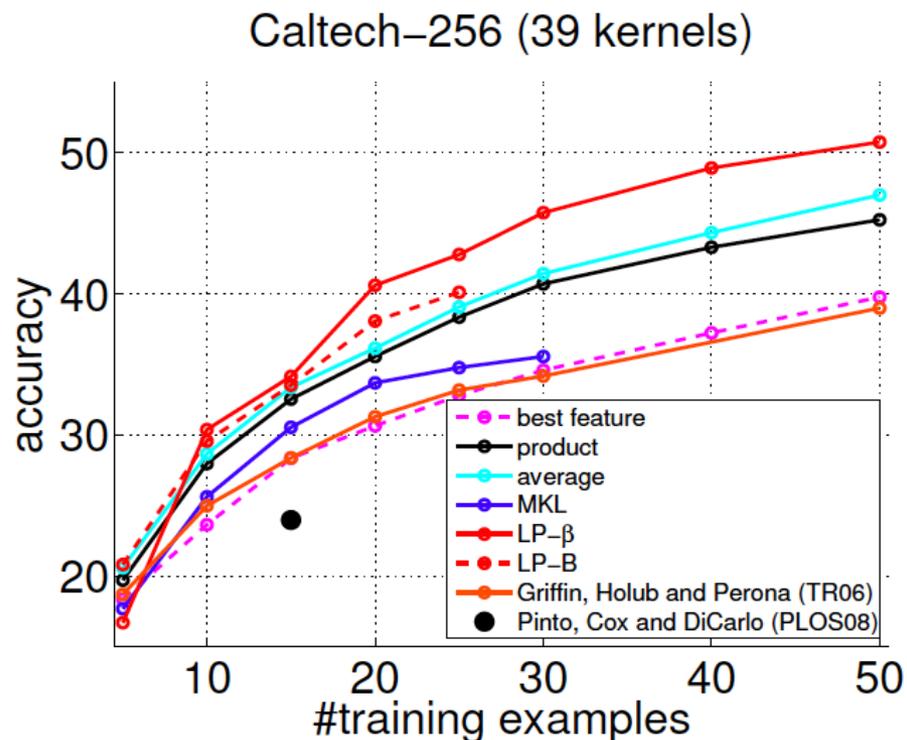


- 良好的特征表达，对最终算法的准确性起了非常关键的作用；
- 识别系统主要的计算和测试工作耗时主要集中在特征提取部分；
- 特征的样式目前一般都是人工设计的，靠人工提取特征。

动机——为什么要自动学习特征

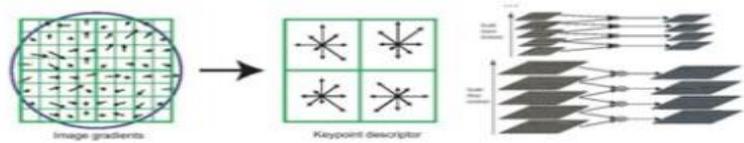
- 实验: LP- β Multiple Kernel Learning
 - Gehler and Nowozin, On Feature Combination for Multiclass Object Classification, ICCV' 09
- 采用39 个不同的特征
 - PHOG, SIFT, V1S+, Region Cov. Etc.
- 在普通特征上MKL表现有限

结论: 特征更重要

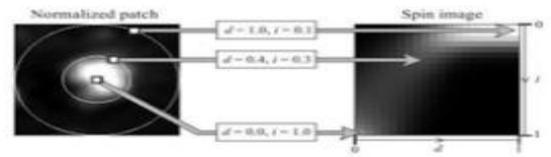


动机——为什么要自动学习特征

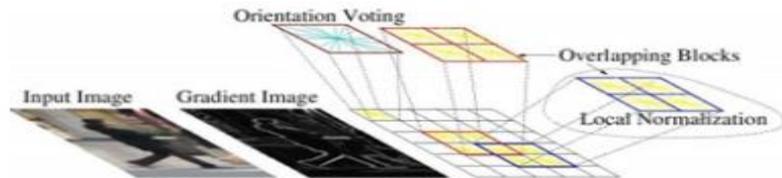
- 机器学习中，获得好的特征是识别成功的关键
- 目前存在大量人工设计的特征，不同研究对象特征不同，特征具有多样性，如：SIFT, HOG, LBP等
- 手工选取特征费时费力，需要启发式专业知识，很大程度上靠经验和运气
- 是否能自动地学习特征？



SIFT



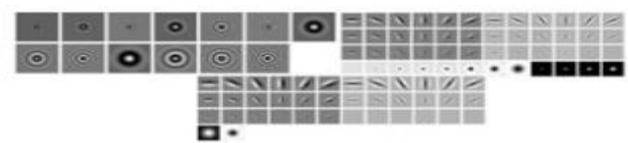
Spin image



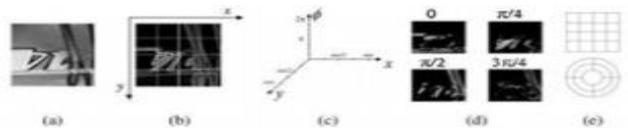
HoG



RIFT



Textons

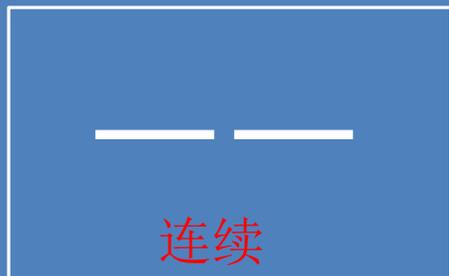


GLOH

动机——为什么要自动学习特征

- 中层特征

✓ 中层信号:



“Tokens” from Vision by D.Marr:

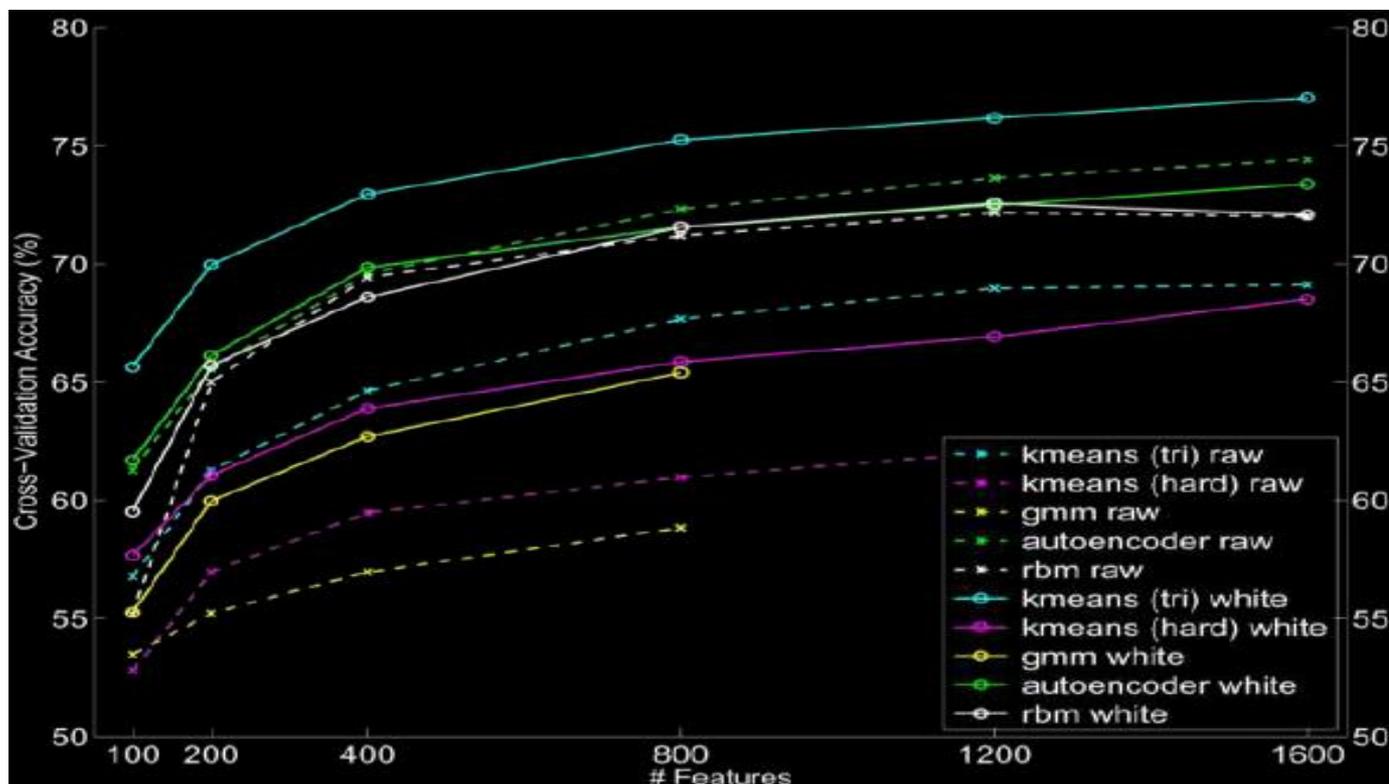


✓ 物体部件:



- 他们对于人工而言是十分困难的，那么如何学习呢？

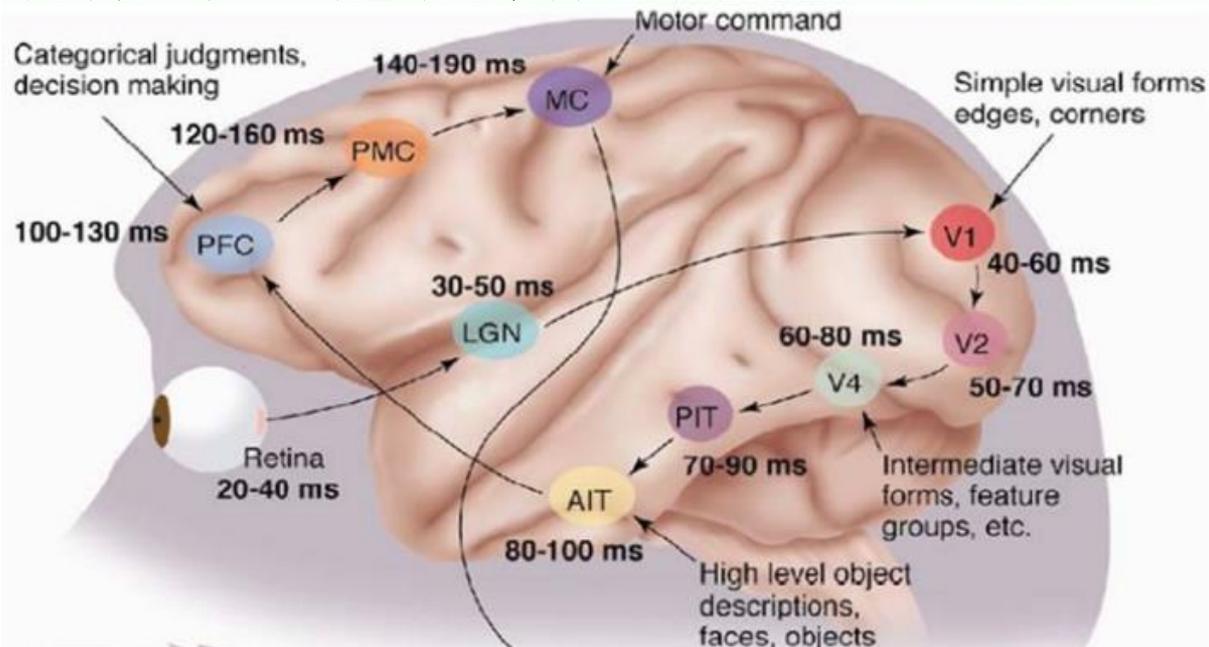
动机——为什么要自动学习特征



- 一般而言，特征越多，给出信息就越多，识别准确性会得到提升；
- 但特征多，计算复杂度增加，探索的空间大，可以用来训练的数据在每个特征上就会稀疏。
- **结论：不一定特征越多越好！需要有多少个特征，需要学习确定。**

动机——为什么采用层次网络结构

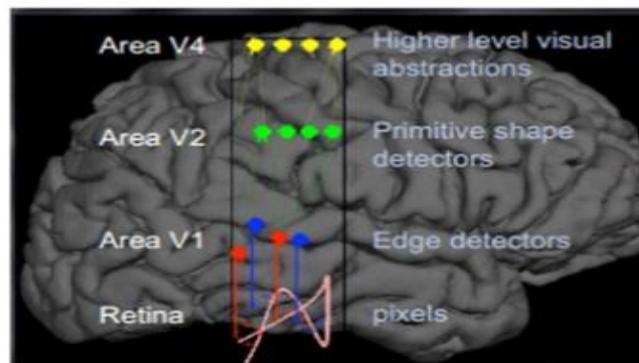
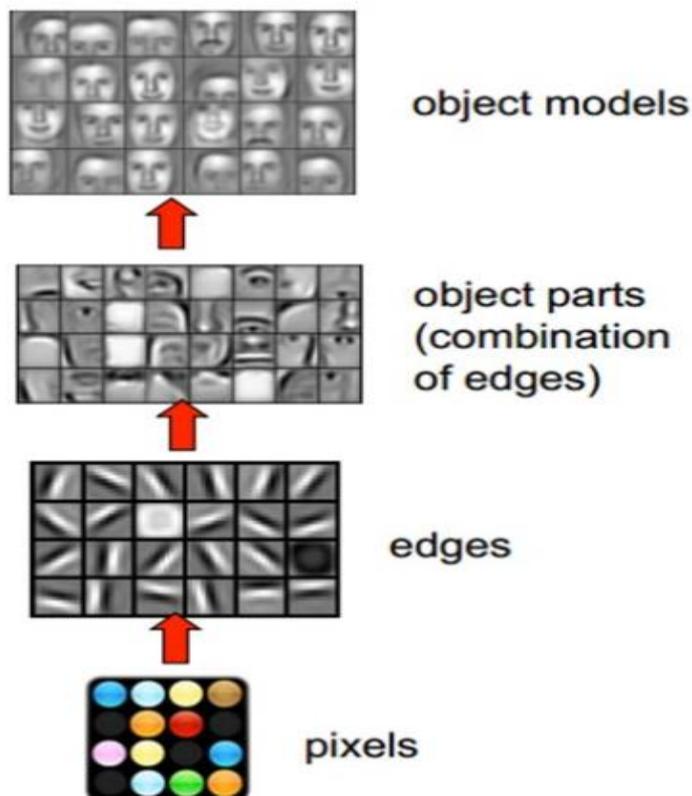
- 人脑视觉机理
 - ✓ 1981年的诺贝尔医学奖获得者 David Hubel和TorstenWiesel发现了视觉系统的信息处理机制
 - ✓ 发现了一种被称为“方向选择性细胞的神经元细胞，当瞳孔发现了眼前的物体的边缘，而且这个边缘指向某个方向时，这种神经元细胞就会活跃



动机——为什么采用层次网络结构

• 人脑视觉机理

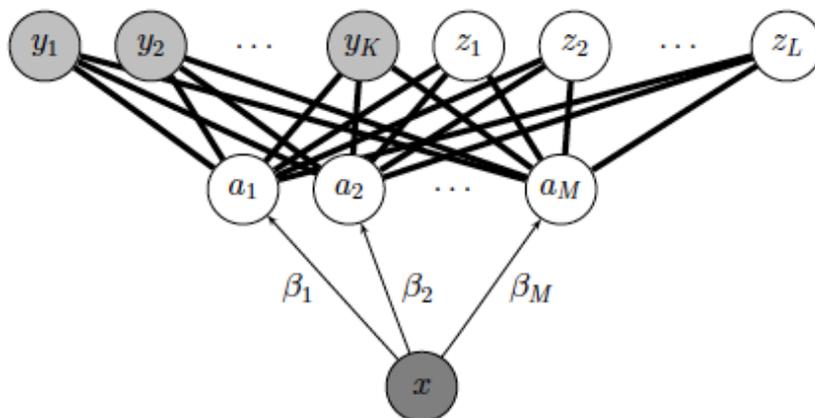
- ✓ 人的视觉系统的信息处理是分级的
- ✓ 高层的特征是低层特征的组合，从低层到高层的特征表示越来越抽象，越来越能表现语义或者意图
- ✓ 抽象层面越高，存在的可能猜测就越少，就越利于分类



动机——为什么采用层次网络结构

- 视觉的层次性
- ✓ 属性学习，类别作为属性的一种组合映射

Lampert et al. CVPR' 09



otter

black: yes
white: no
brown: yes
stripes: no
water: yes
eats fish: yes



polar bear

black: no
white: yes
brown: no
stripes: no
water: yes
eats fish: yes



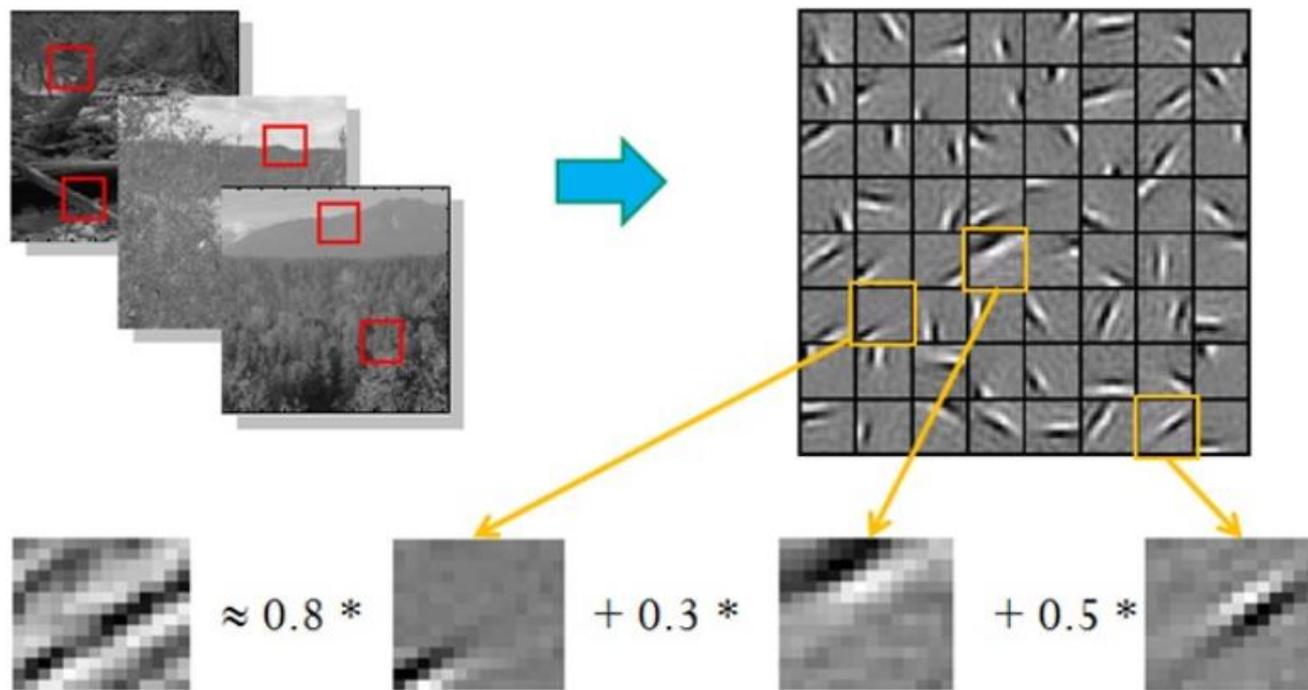
zebra

black: yes
white: yes
brown: no
stripes: yes
water: no
eats fish: no



动机——为什么采用层次网络结构

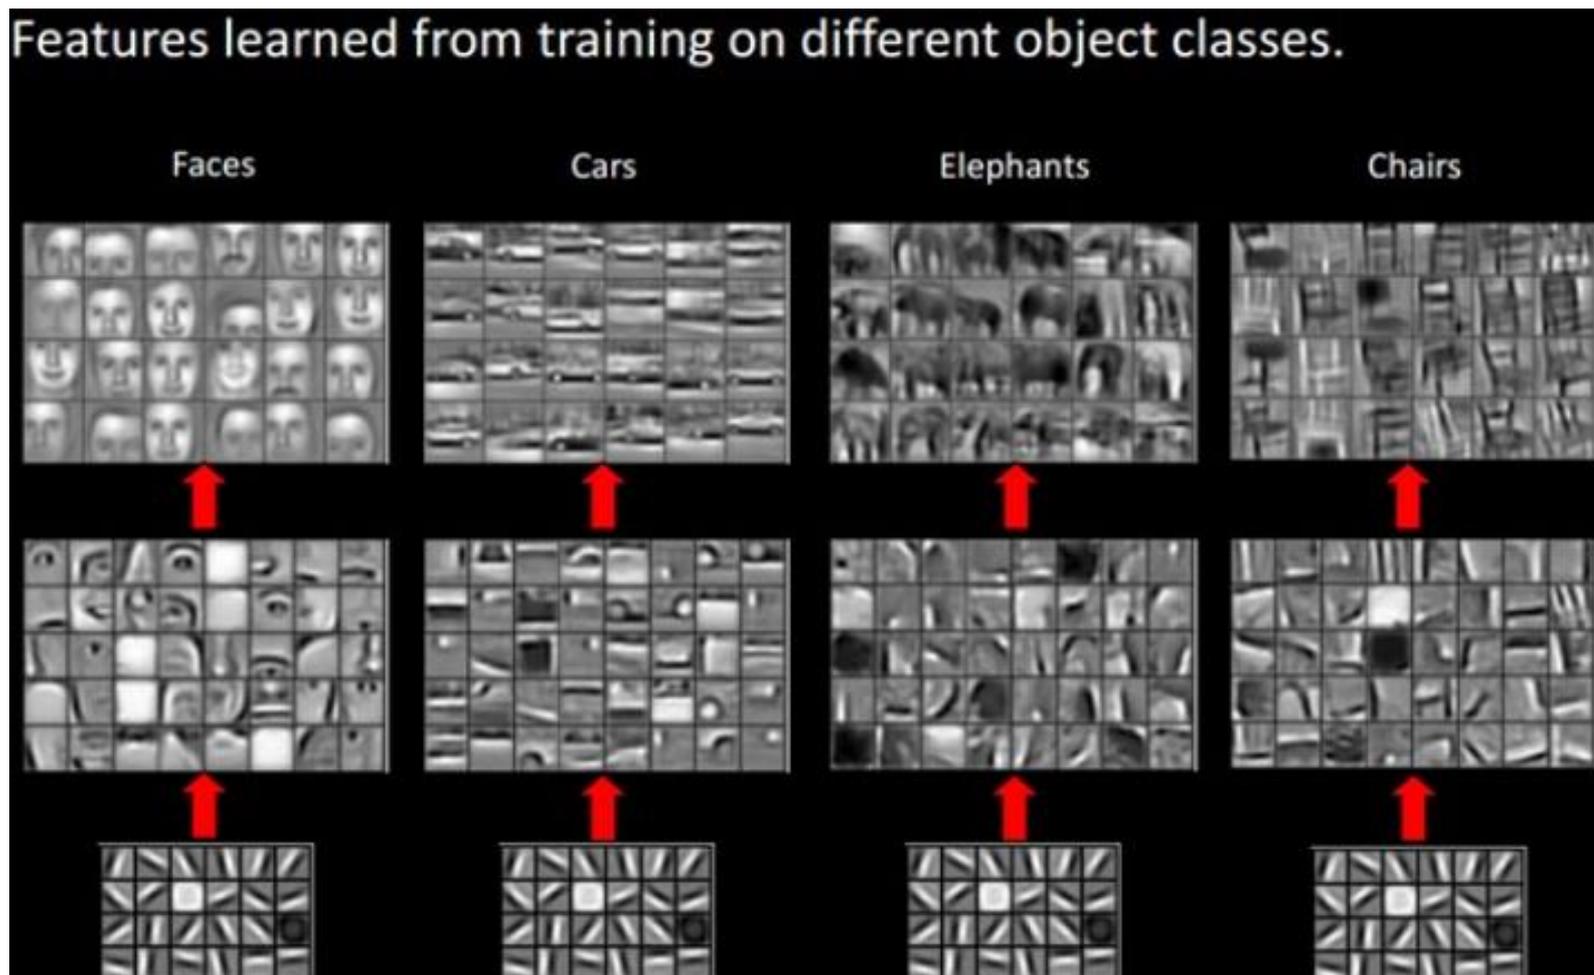
- 初级（浅层）特征表示



$[a_1, \dots, a_{64}] = [0, 0, \dots, 0, \mathbf{0.8}, 0, \dots, 0, \mathbf{0.3}, 0, \dots, 0, \mathbf{0.5}, 0]$
(feature representation)

动机——为什么采用层次网络结构

- 结构性特征表示



动机——为什么采用层次网络结构

- 浅层学习的局限

- ✓ 人工神经网络（BP算法）

- 虽被称作多层感知机，但实际是种只含有一层隐层节点的浅层模型

- ✓ SVM、Boosting、最大熵方法（如LR, Logistic Regression）

- 带有一层隐层节点（如SVM、Boosting），或没有隐层节点（如LR）的浅层模型

局限性：有限样本和计算单元情况下对复杂函数的表示能力有限，针对复杂分类问题其泛化能力受限。

深度学习

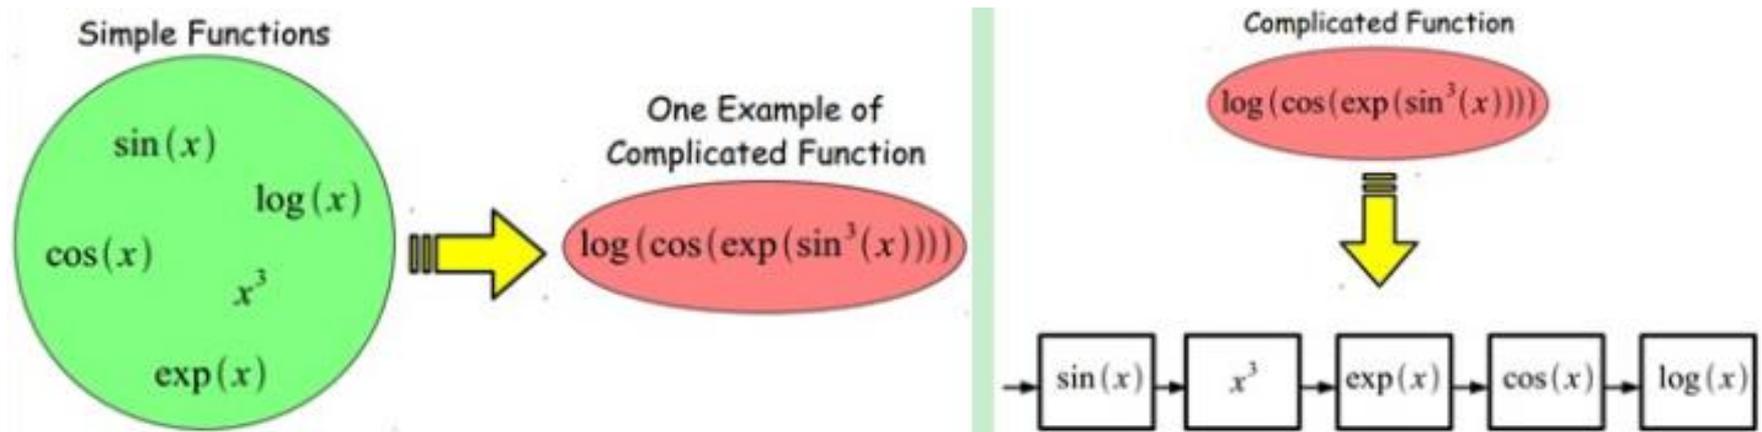
- 2006年，加拿大多伦多大学教授、机器学习领域的泰斗Geoffrey Hinton在《科学》上发表论文提出深度学习主要观点：
 - 1) 多隐层的人工神经网络具有优异的特征学习能力，学习得到的特征对数据有更本质的刻画，从而有利于可视化或分类；
 - 2) 深度神经网络在训练上的难度，可以通过“逐层初始化”（layer-wise pre-training）来有效克服，逐层初始化可通过无监督学习实现的。

深度学习

- **本质：**通过构建多隐层的模型和海量训练数据（可为无标签数据），来学习更有用的特征，从而最终提升分类或预测的准确性。“深度模型”是手段，“特征学习”是目的。
- **与浅层学习区别：**
 - 1) 强调了模型结构的深度，通常有5-10多层的隐层节点；
 - 2) 明确突出了特征学习的重要性，通过逐层特征变换，将样本在原空间的特征表示变换到一个新特征空间，从而使分类或预测更加容易。与人工规则构造特征的方法相比，利用大数据来学习特征，更能够刻画数据的丰富内在信息。

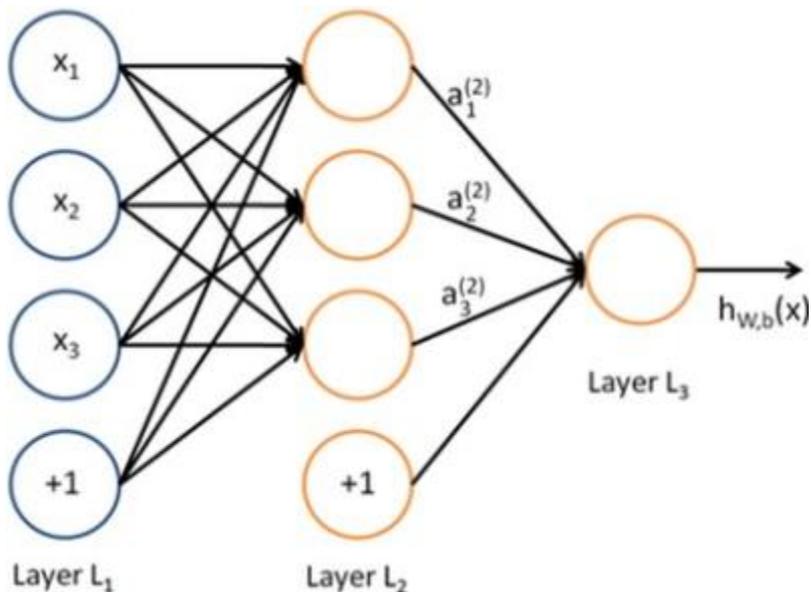
深度学习

- 好处：可通过学习一种深层非线性网络结构，实现复杂函数逼近，表征输入数据分布式表示。

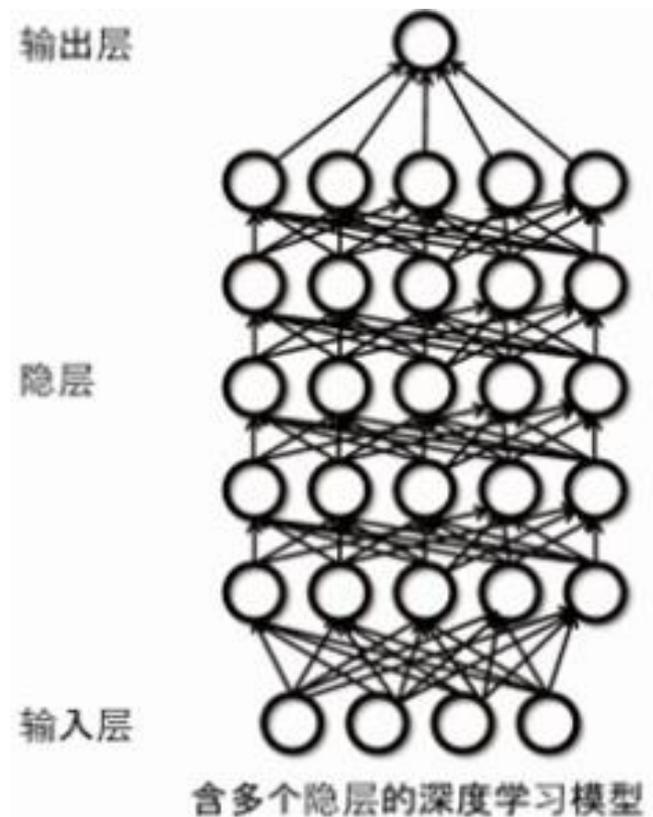


深度学习 vs. 神经网络

神经网络 :



深度学习 :



深度学习 vs. 神经网络

相同点：二者均采用分层结构，系统包括输入层、隐层（多层）、输出层组成的多层网络，只有相邻层节点之间有连接，同一层以及跨层节点之间相互无连接，每一层可以看作是一个logistic 回归模型。

不同点：

神经网络：采用BP算法调整参数，即采用迭代算法来训练整个网络。随机设定初值，计算当前网络的输出，然后根据当前输出和样本真实标签之间的差去改变前面各层的参数，直到收敛；

深度学习：采用逐层训练机制。采用该机制的原因在于如果采用BP机制，对于一个deep network（7层以上），残差传播到最前面的层将变得很小，出现所谓的gradient diffusion（梯度扩散）。

深度学习 vs. 神经网络

- 神经网络的局限性:
 - 1) 比较容易过拟合，参数比较难调整，而且需要不少技巧；
 - 2) 训练速度比较慢，在层次比较少（小于等于3）的情况下效果并不比其它方法更优；

Deep learning

Yoshua Bengio:

Science is NOT a battle, it is a collaboration. We all build on each other's ideas. Science is an act of love, not war. Love for the beauty in the world that surrounds us and love to share and build something together. That makes science a highly satisfying activity, emotionally speaking!

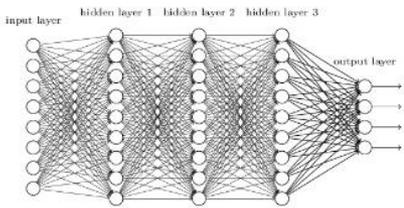
科学不是战争而是合作，任何学科的发展从来都不是一条路走到黑，而是同行之间互相学习、互相借鉴、博采众长、相得益彰，站在巨人的肩膀上不断前行。机器学习的研究也是一样，你死我活那是邪教，开放包容才是正道。

Part II

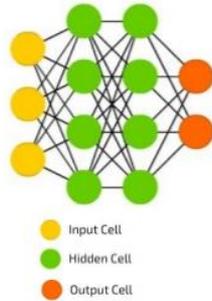
深度学习的基本方法
(以及神经网络算法)

Neural network

Deep Neural Network (DNN)

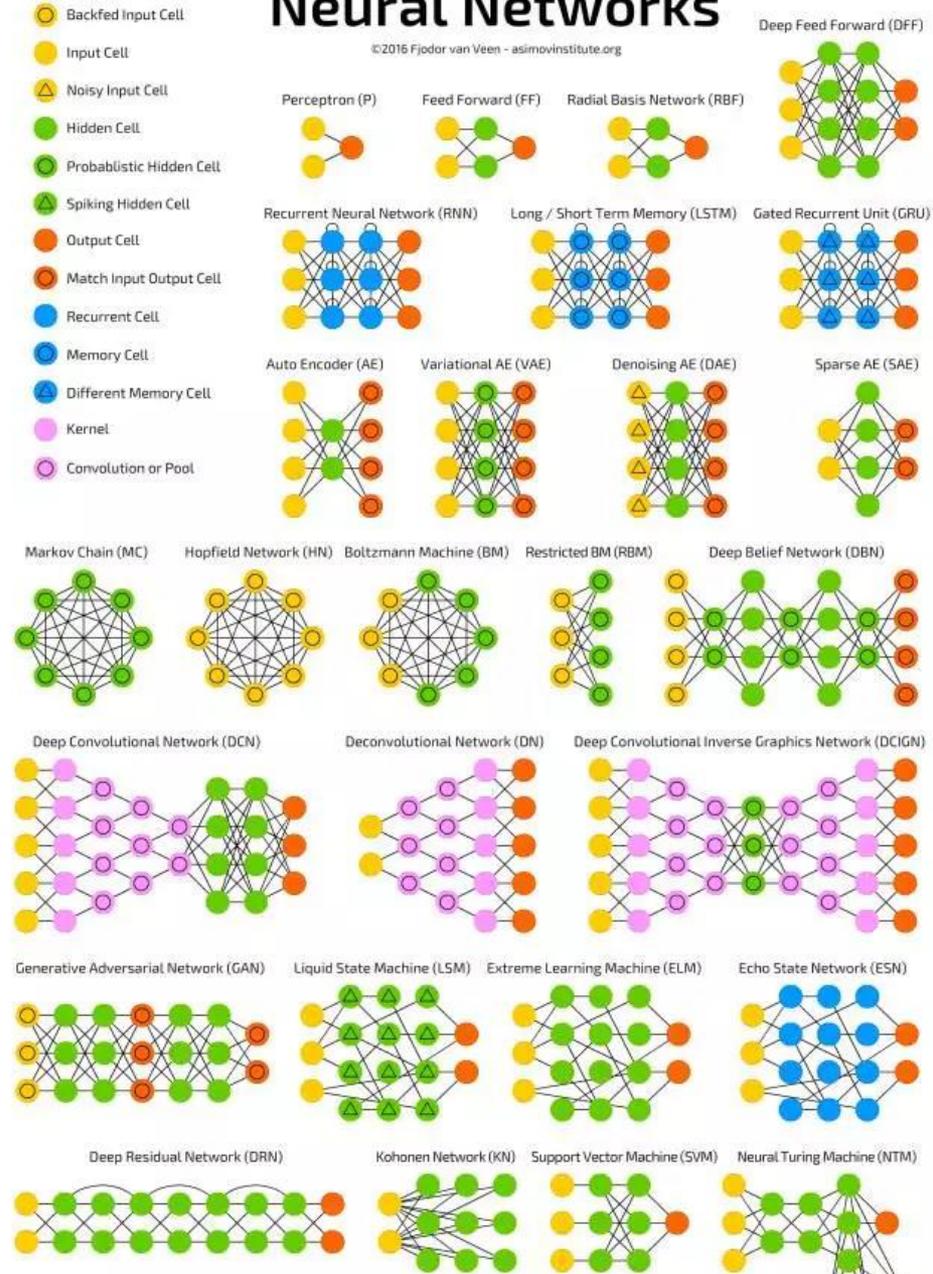


Deep Feed Forward (DFF)

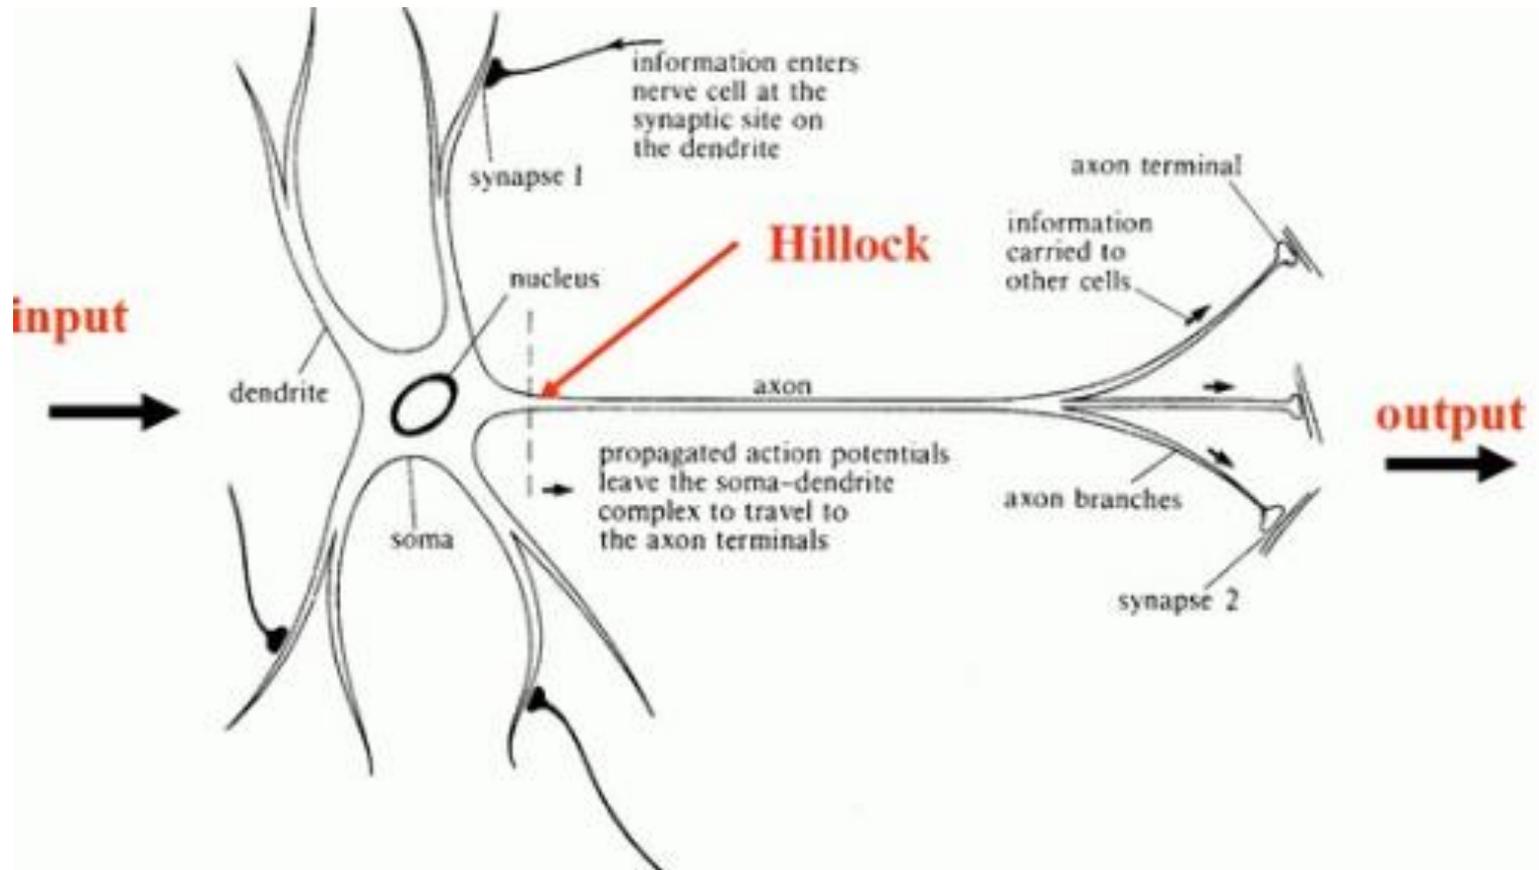


A mostly complete chart of Neural Networks

©2016 Fjodor van Veen - asimovinstitute.org



Neural network

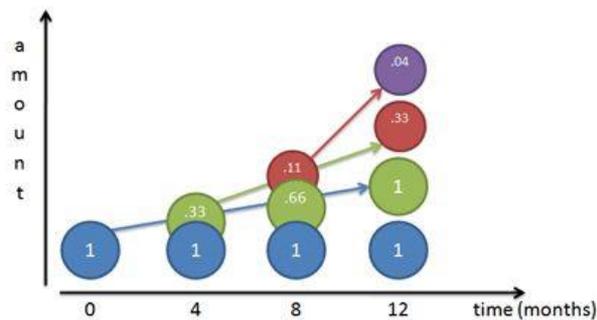
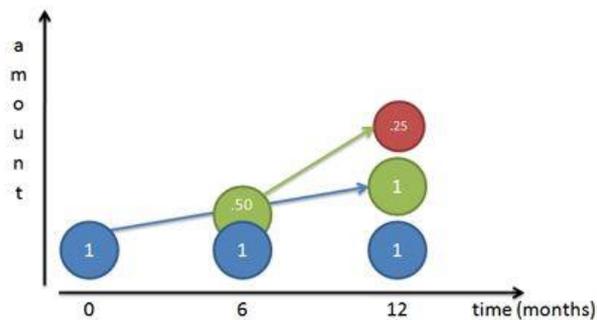
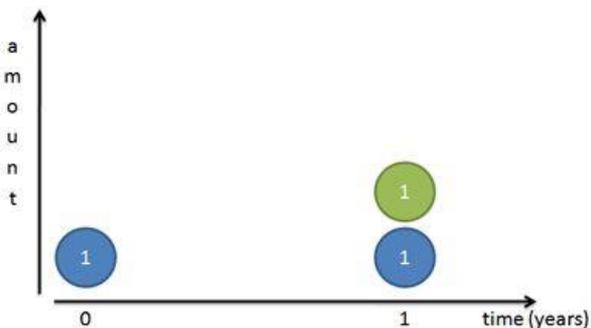


Neural network

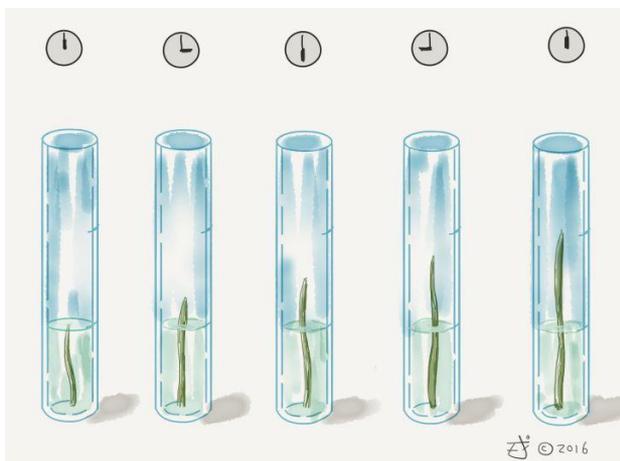
- 自然对数 e
- 逻辑回归
- 神经网络
- 反向传播算法(BP算法)
- 示例

Neural network

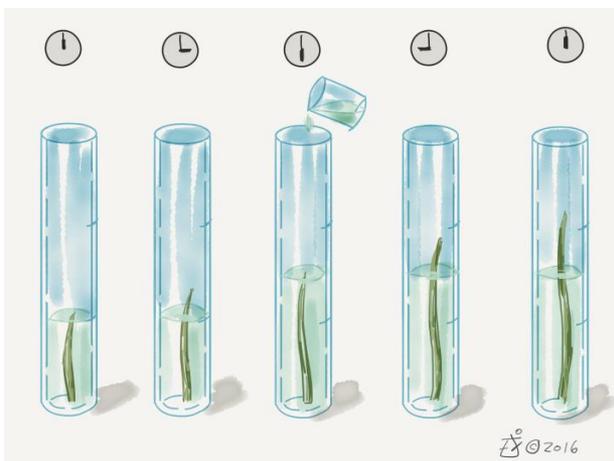
自然对数e



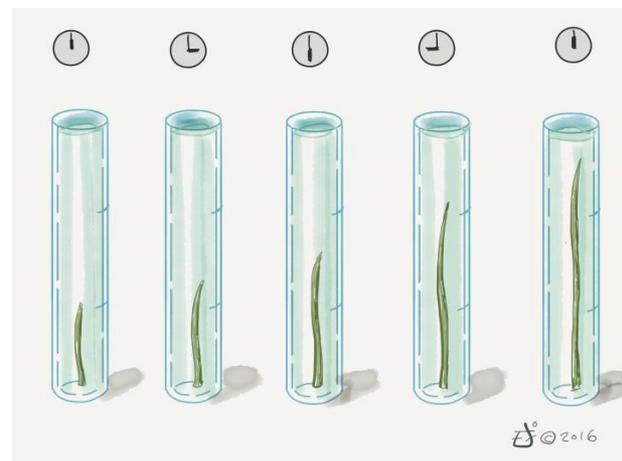
1元存1年，在年利率100%下，无论怎么利滚利，其余额总有一个上限，这个天花板就是e



有一种草齐头泡水半天变**2倍**



中间加一次水能长到**2.25倍**



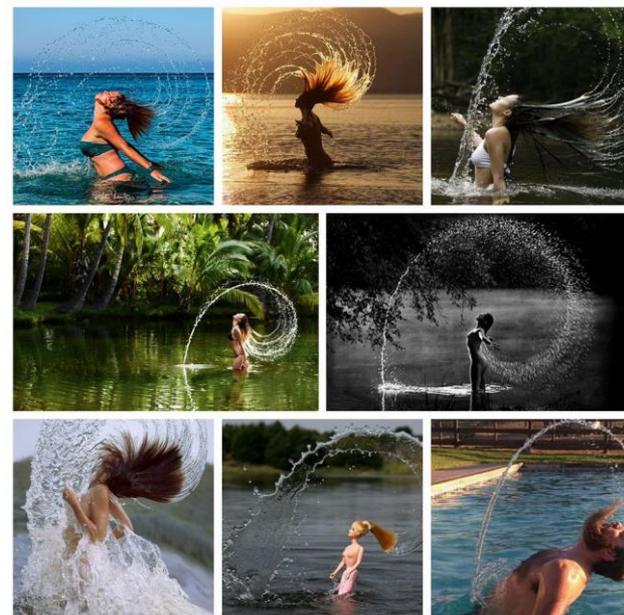
加足水最多能长到**2.718...倍**

Neural network

自然对数e

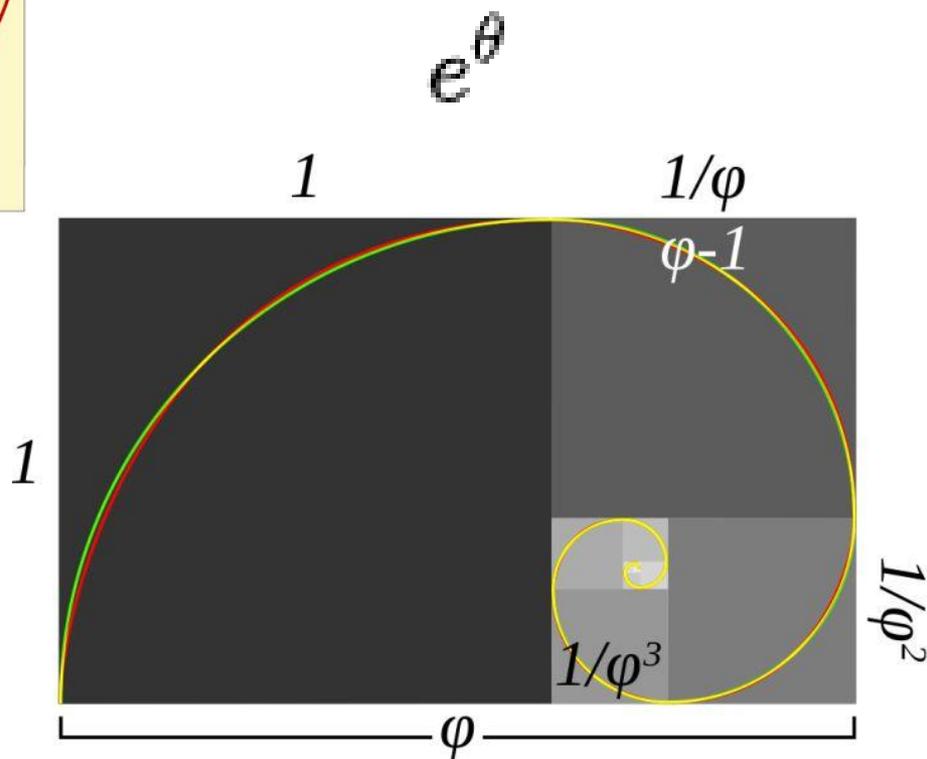
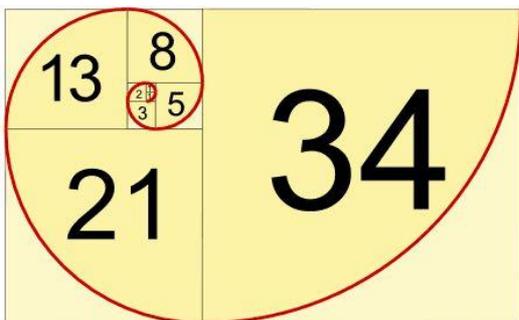
$$e = 1 + 1/1! + 1/2! + 1/3! + 1/4! + 1/5! + 1/6! + 1/7! + \dots = 1 + 1 + 1/2 + 1/6 + 1/24 + 1/120 + \dots \approx 2.71828$$

$$e = (1 + 1/x)^x \approx 2.71828$$



Neural network

自然对数 e vs. 斐波拉切数列

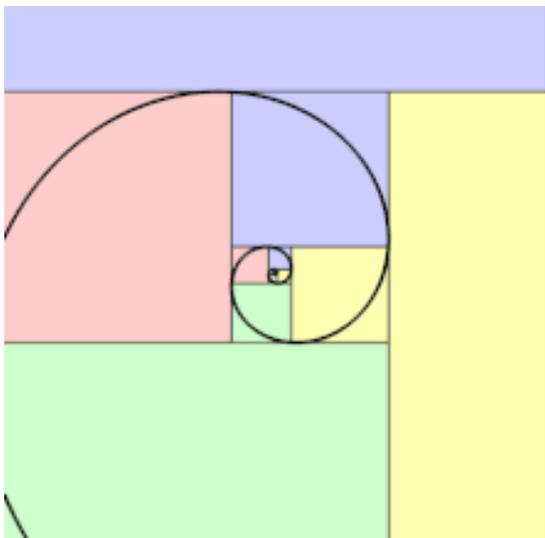


斐波那契螺旋线仅仅是对一种叫黄金螺旋线（Golden spiral）的近似，黄金螺旋线是一种内涵黄金分割比例的对数螺旋线。

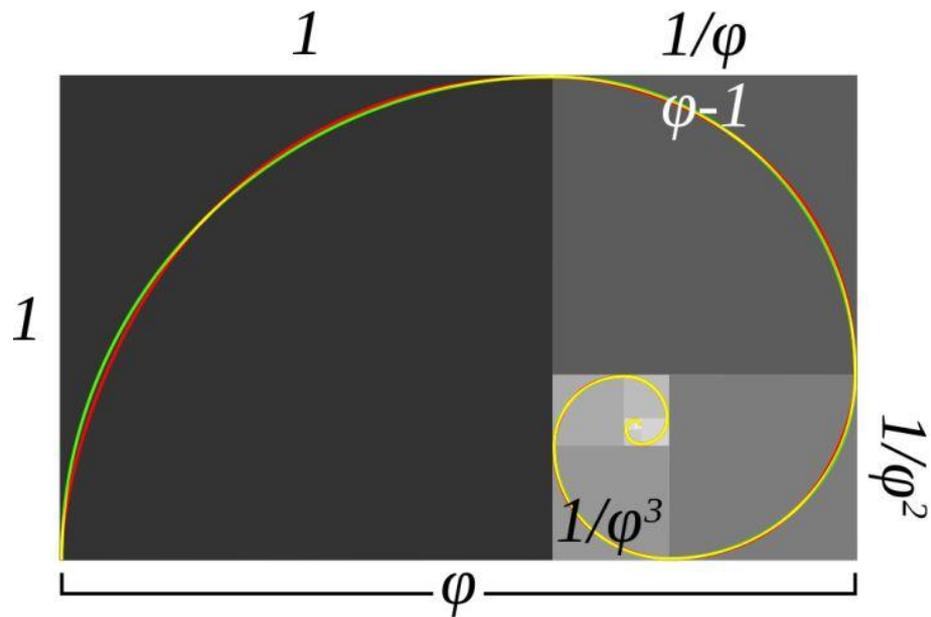
红色的才是黄金曲线，绿色的是“假黄金螺旋线”（斐波那契螺旋线），近似却不重合。

Neural network

自然对数 e & 黄金分割



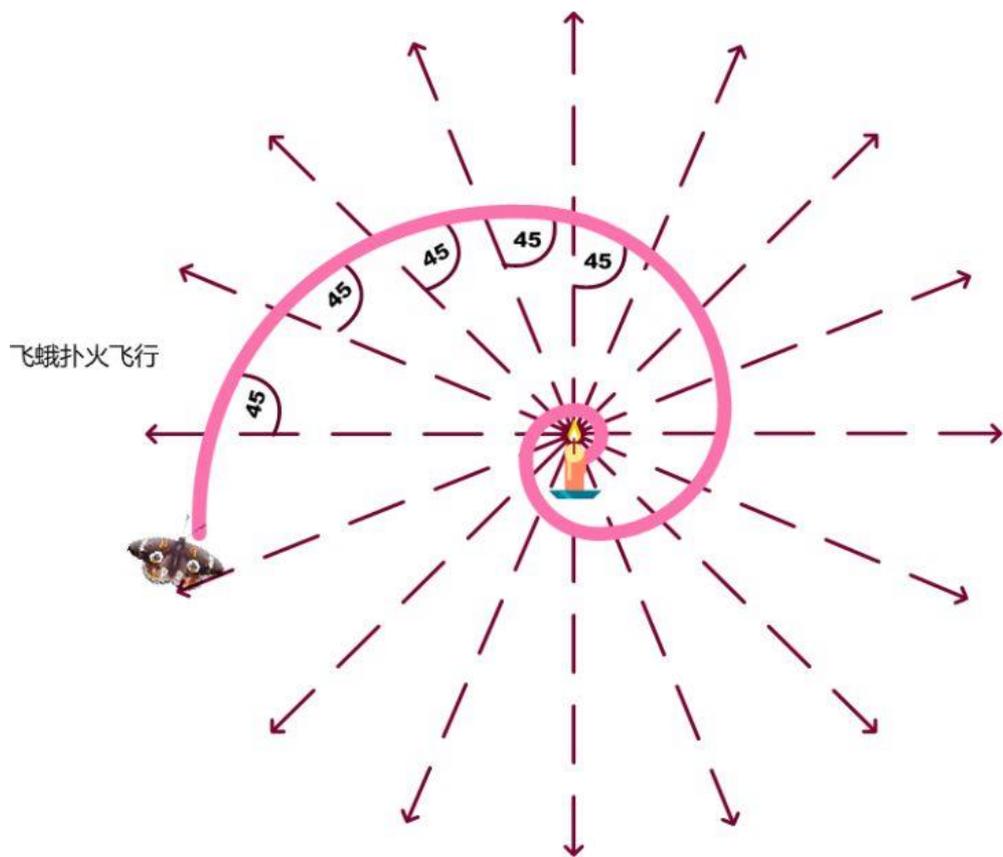
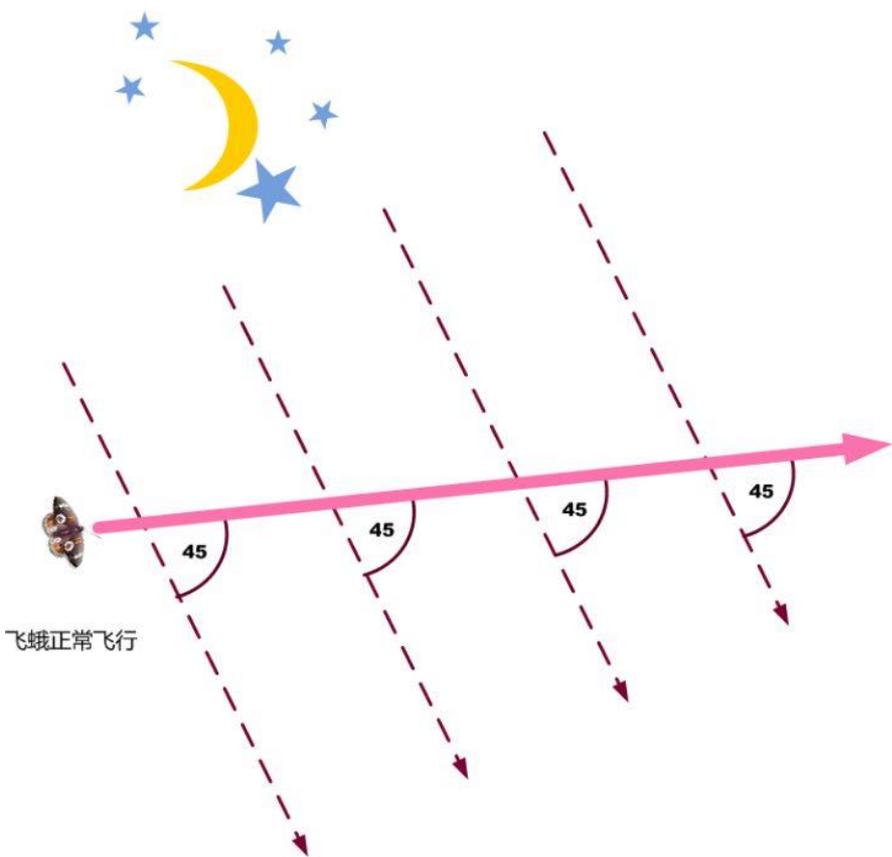
e^{θ}



$\varphi=1.6180339887498948482\dots$

Neural network

自然对数 e

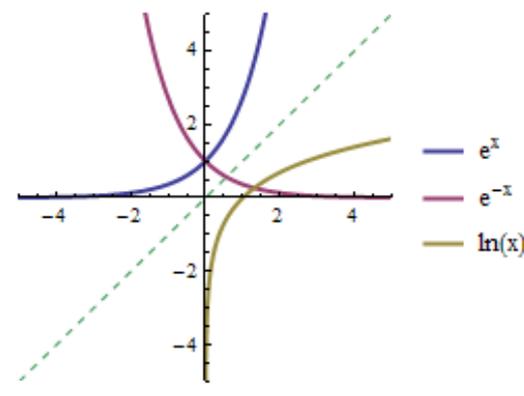


Neural network

自然对数e

幂(power) $\rightarrow y = e^x$ 指数(exponent)
底数(base number)

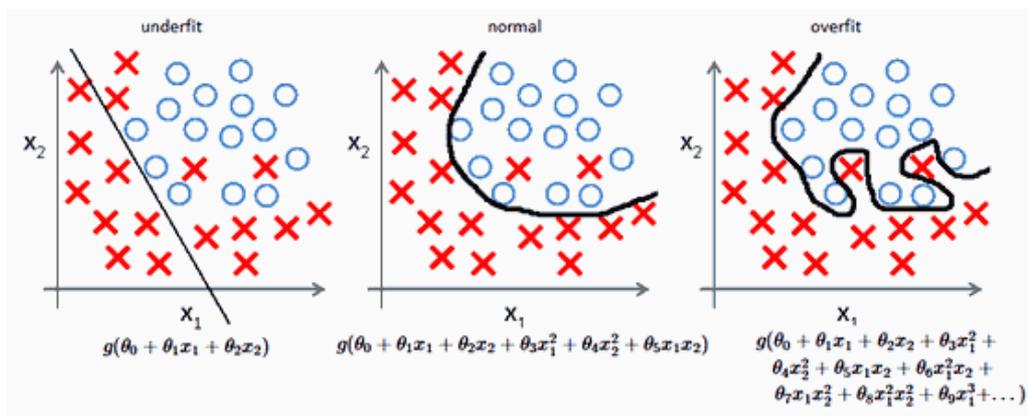
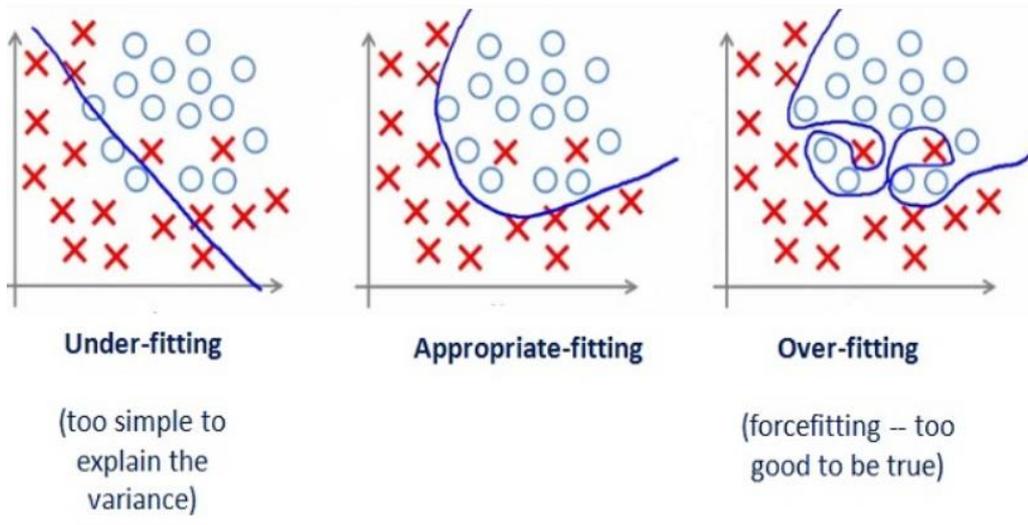
对数(logarithm) $\rightarrow x = \log_e(y) = \ln(y)$ 真数(正数) 自然对数(natural logarithm)



e^x 和 e^{-x} 的图形是对称的； $\ln(x)$ 是 e^x 的逆函数，它们呈45度对称。

Neural network

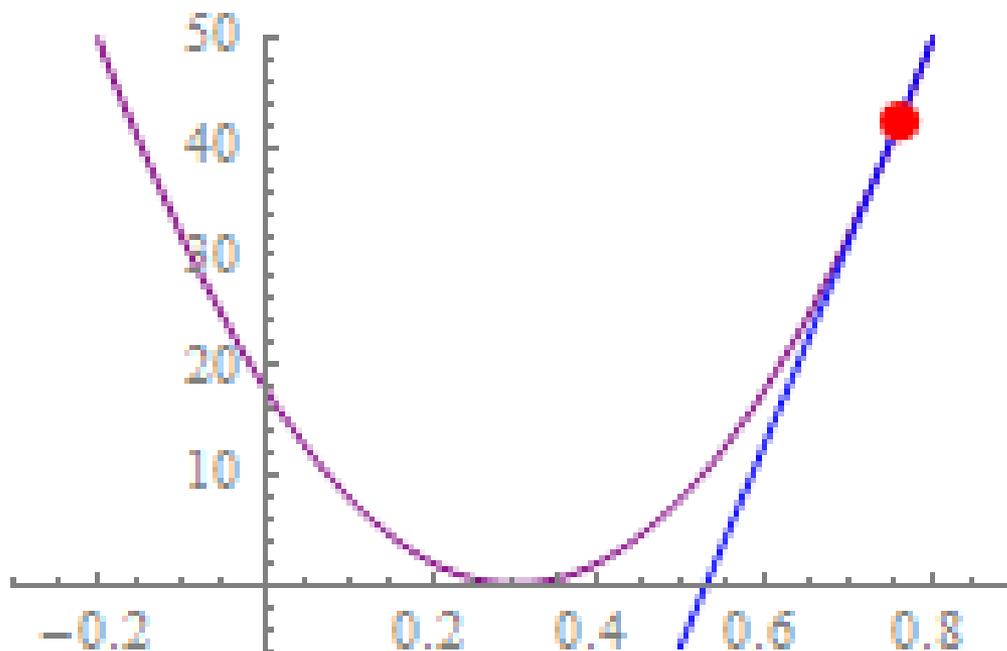
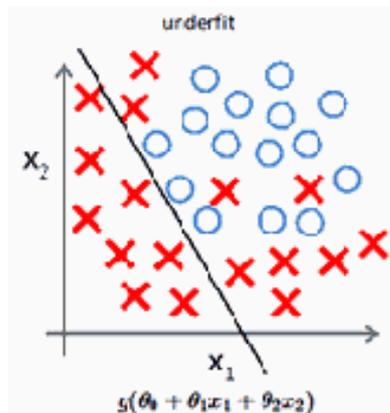
逻辑回归



Neural network

逻辑回归

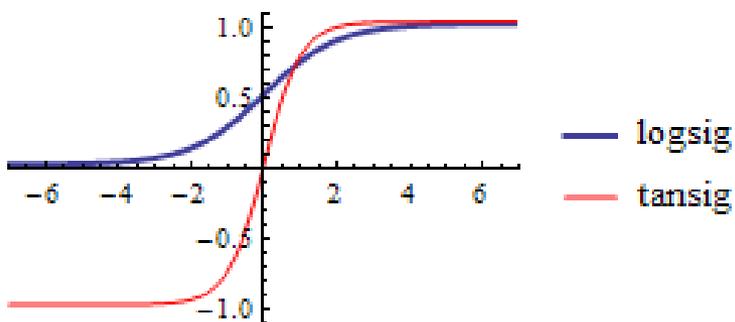
切线斜率 = 0.46



切线每次旋转的幅度叫做学习率(Learning Rate)，加大学习率会加快拟合速度，但是如果调得太大会导致切线旋转过度而无法收敛。[学习率其实是个预先设置好的参数，不会每次变化，不过可以影响每次变化的幅度。]

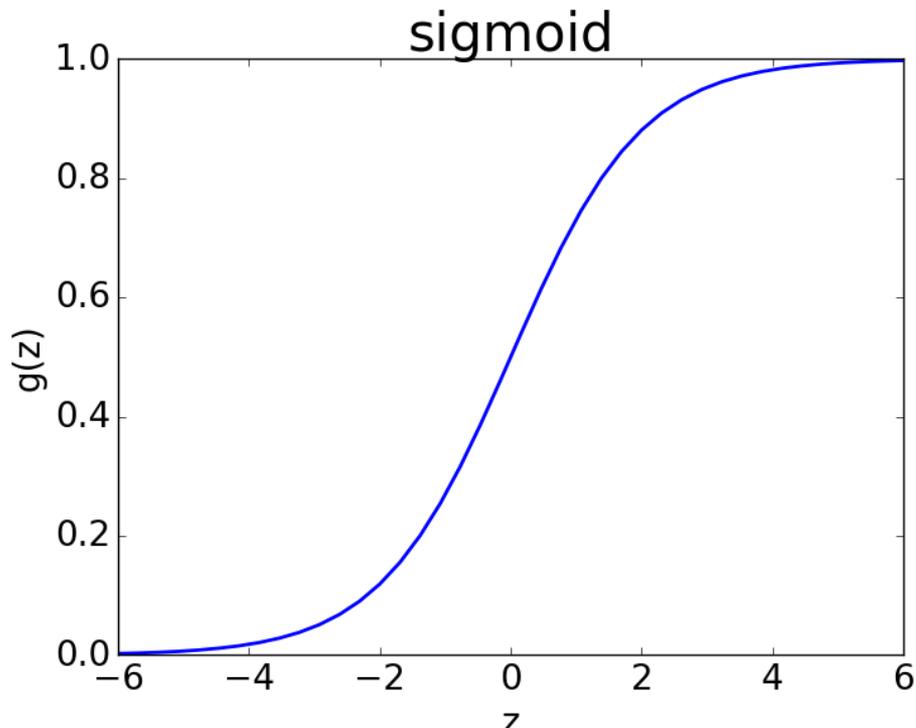
Neural network

逻辑回归(Logistic Regression, 逻辑斯谛函数)



$$y = \text{logsig}(x) = \frac{1}{1 + e^{-x}}$$

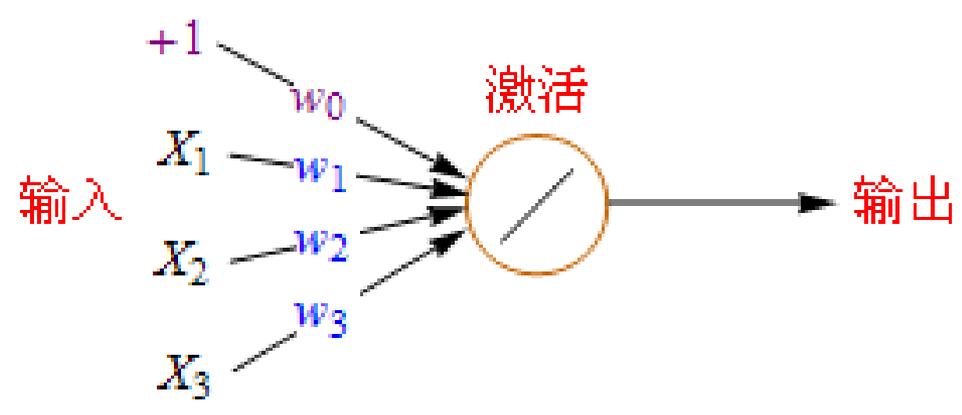
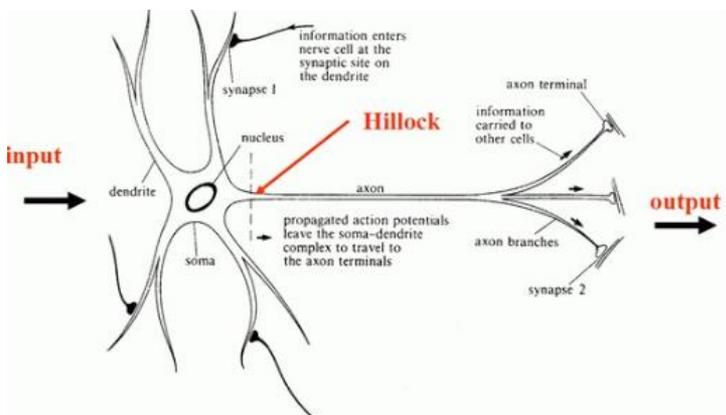
$$y = \text{tansig}(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}}$$



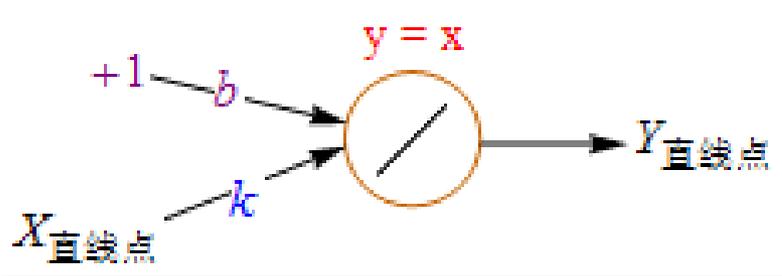
y 的阈值处于 $(-\infty, +\infty)$, 此时不能很好的给出属于某一类的概率, 因为概率的范围是 $[0,1]$, 我们需要一个更好的映射函数, 能够将分类的结果很好的映射成为 $[0,1]$ 之间的概率, 并且这个函数能够具有很好的可微分性。

Neural network

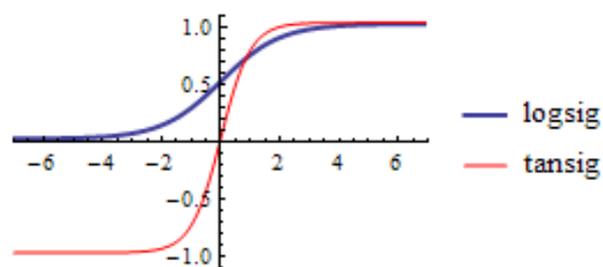
逻辑回归



Purelin 激活函数



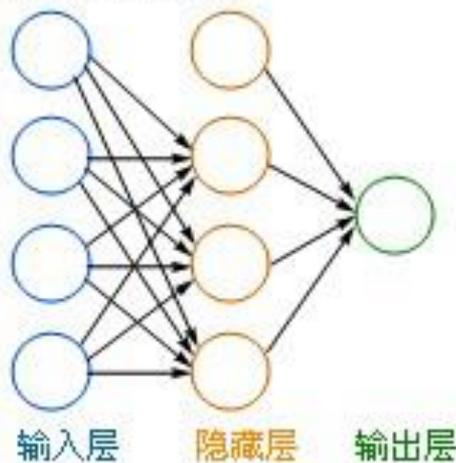
Sigmoid 激活函数



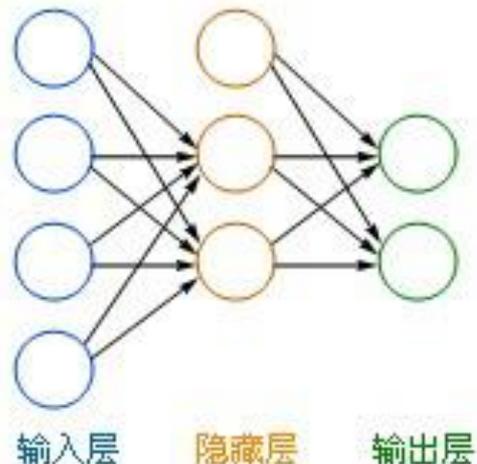
Neural network

神经网络

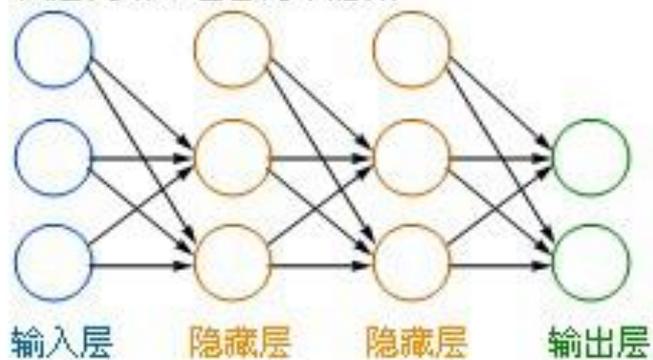
典型的三层网络



三层网络，输出层有两个节点

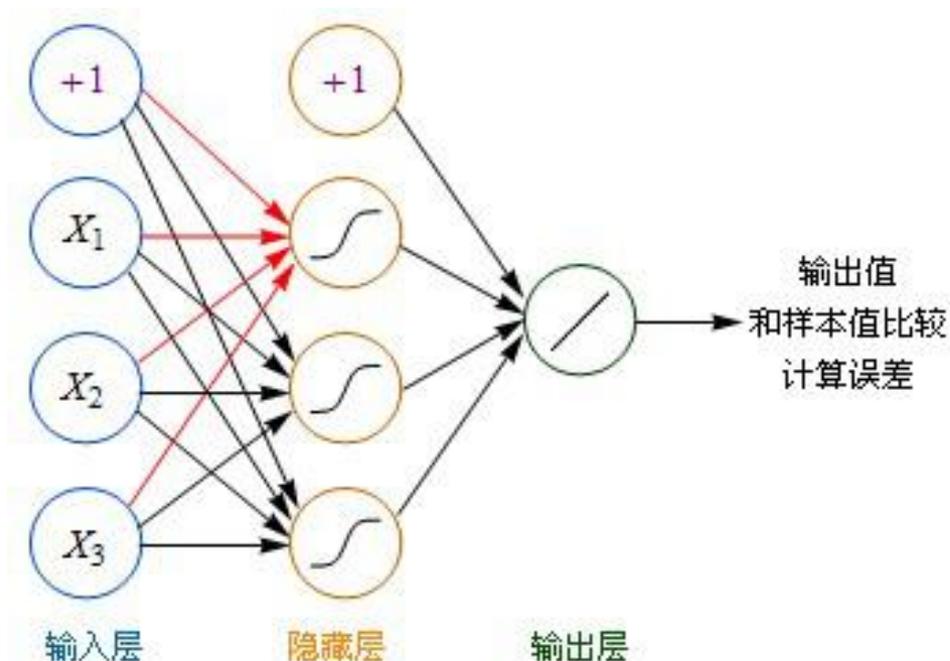


四层网络，包含两个隐藏



Neural network

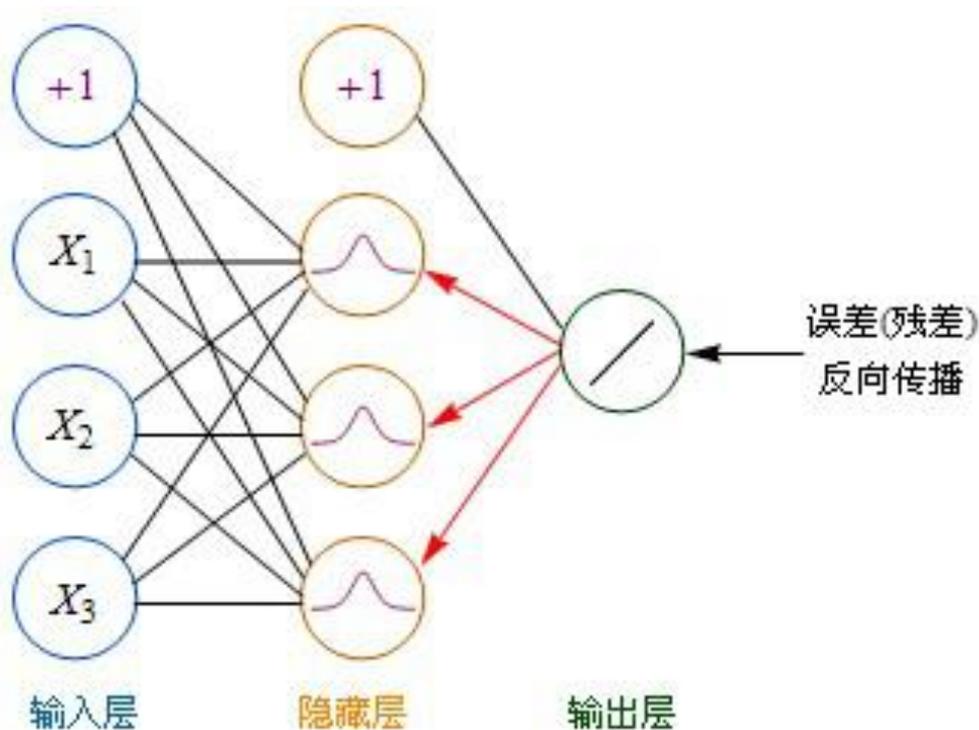
神经网络



- 隐藏层用都是用Sigmoid作激活函数，而输出层用的是Purelin。这是因为Purelin可以保持之前任意范围的数值缩放，便于和样本值作比较，而Sigmoid的数值范围只能在0~1之间。
- 起初输入层的数值通过网络计算分别传播到隐藏层，再以相同的方式传播到输出层，最终的输出值和样本值作比较，计算出误差，这个过程叫前向传播(Forward Propagation)。

Neural network

反向传播算法(BP算法)

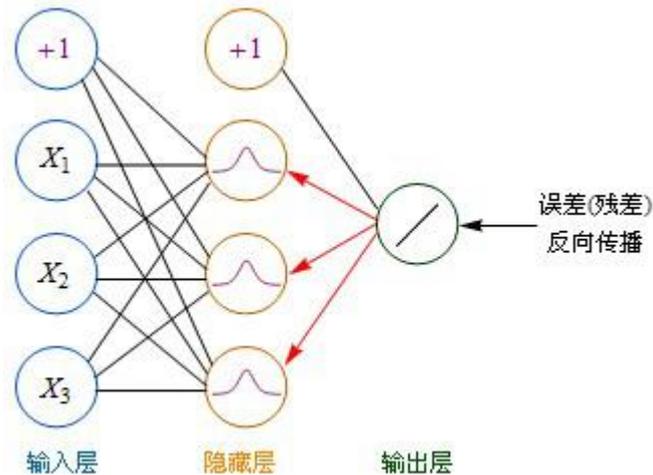


利用前向传播最后输出的结果来计算误差的偏导数，再用这个偏导数和前面的隐藏层进行加权求和，如此一层一层的向后传下去，直到输入层(不计算输入层)，最后利用每个节点求出的偏导数来更新权重

Neural network

反向传播算法(BP算法)

前馈神经网络(FeedForward Neural Network), 也叫BP神经网络(Back Propagation Neural Network)。



如果输出层用Purelin作激活函数, Purelin的导数是1, 输出层 \rightarrow 隐藏层: 残差 = -(输出值-样本值)

如果用Sigmoid(logsig)作激活函数, 那么: Sigmoid导数 = Sigmoid*(1-Sigmoid)

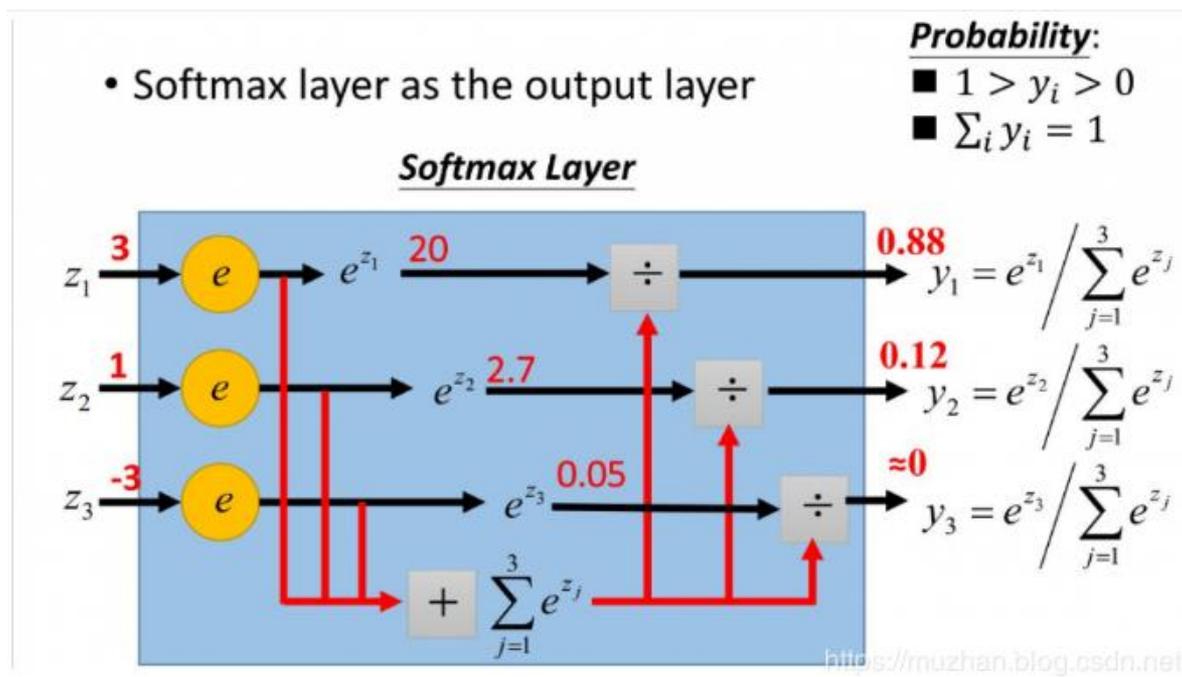
输出层 \rightarrow 隐藏层: 残差 = -(Sigmoid输出值-样本值) * Sigmoid*(1-Sigmoid) = -(输出值-样本值) * 输出值*(1-输出值)

隐藏层 \rightarrow 隐藏层: 残差 = (右层每个节点的残差加权求和) * 当前节点的Sigmoid*(1-当前节点的Sigmoid)

Neural network

softmax分类器

$$\text{softmax}(x_0) = \frac{e^{x_0}}{e^{x_0} + e^{x_1} + e^{x_2}}$$



softmax对神经元的输出信号进行加工，输出为分类的概率值。

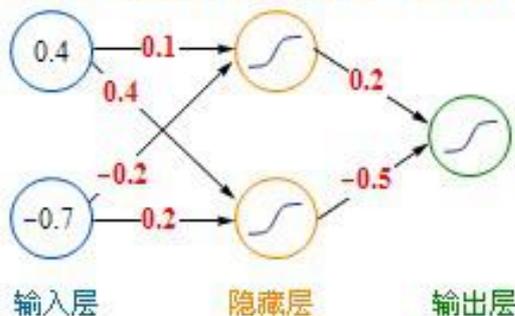
Neural network

示例

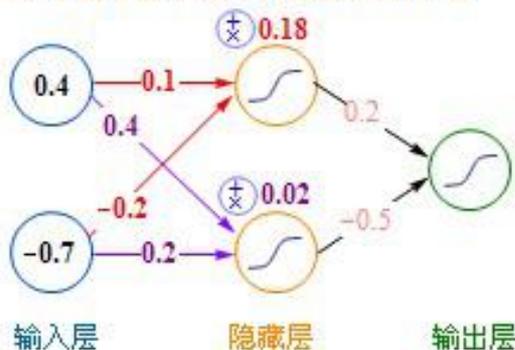
图例：

-  logsig激活函数
-  残差(误差偏导数)
-  加权求和

1. 每个节点之间的初始权重一般是随机生成的



2. 对输入层节点进行加权求和计算



$$0.4 \times 0.1 + (-0.7) \times (-0.2) = 0.18$$

$$0.4 \times 0.4 + (-0.7) \times (0.2) = 0.02$$

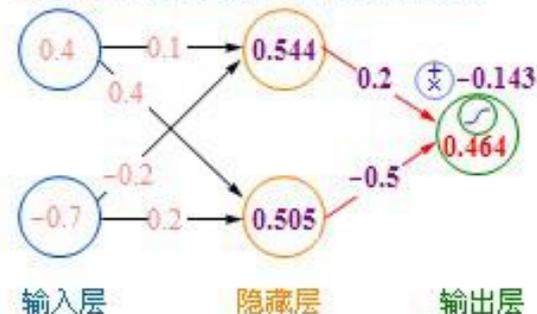
3. 执行Sigmoid激活



$$\text{logsig}(0.18) = \frac{1}{1 + e^{-0.18}} = 0.544$$

$$\text{logsig}(0.02) = \frac{1}{1 + e^{-0.02}} = 0.505$$

4. 用相同的方法计算出输出层的值



$$0.544 \times 0.2 + 0.505 \times (-0.5) = -0.143$$

$$\text{logsig}(-0.143) = \frac{1}{1 + e^{0.143}} = 0.464$$

训练集的数据，首先对第一行进行处理

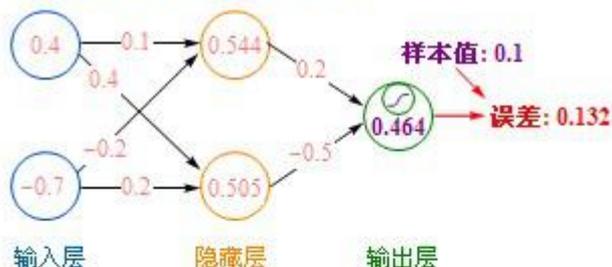
X1	X2	样本值
0.4	-0.7	0.1
0.3	-0.5	0.05
0.6	0.1	0.3
0.2	0.4	0.25
0.1	-0.2	0.12

为了适应输出层的logsig变换，X1,X2的数值范围规定在至0~1之间

Neural network

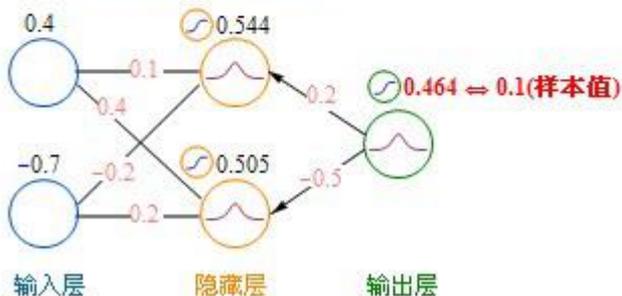
示例

5. 计算误差，误差接近0时收敛

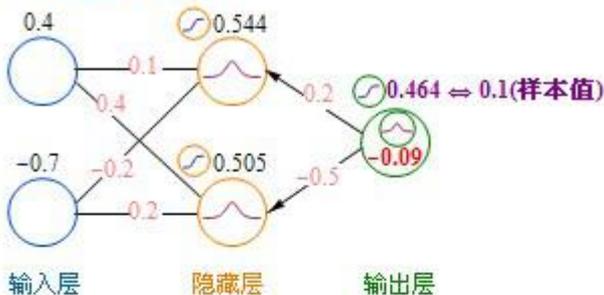


误差 = $(0.464 - 0.1)^2 = 0.132$ (误差 < 0.0001, 可以收敛)

6. 开始准备误差的反向传播



7. 计算输出层的残差



残差 = $-(0.464 - 0.1) \times 0.464 \times (1 - 0.464) = -0.09$

8. 输出层节点的残差加权求和



$-0.09 \times 0.2 = -0.018$ (输出层只有一个节点, 不用求和)

$-0.09 \times 0.5 = -0.045$

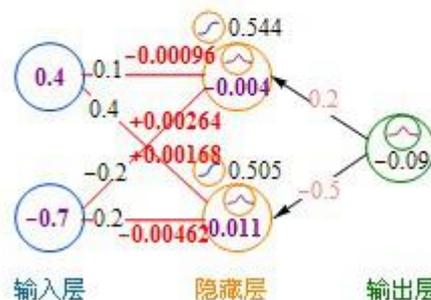
9. 继续求隐藏层的残差



残差1 = $(-0.018) \times (0.544) \times (1 - 0.544) = -0.004$

残差2 = $0.045 \times 0.505 \times (1 - 0.505) = 0.011$

10. 开始准备更新第一层权重，设学习率为0.6



$0.4 \times (-0.004) \times 0.6 = -0.0016 \times 0.6 = -0.00096$

$0.4 \times 0.011 \times 0.6 = 0.0044 \times 0.6 = 0.00264$

$-0.7 \times (-0.004) \times 0.6 = 0.0028 \times 0.6 = 0.00168$

$-0.7 \times 0.011 \times 0.6 = -0.0077 \times 0.6 = -0.00462$

Neural network

示例

11. 更新第一层权重



$$\begin{aligned}0.1 - 0.00096 &= 0.09904 \\0.4 + 0.00264 &= 0.40264 \\-0.2 + 0.00168 &= -0.19832 \\0.2 - 0.00462 &= 0.19538\end{aligned}$$

12. 使用前面两步的方法，计算出后两层的权重



$$\begin{aligned}0.544 \times (-0.09) \times 0.6 &= -0.04896 \times 0.6 = -0.029376 \\0.505 \times (-0.09) \times 0.6 &= -0.04545 \times 0.6 = -0.02727 \\0.2 - 0.029376 &= 0.170624 \\-0.5 - 0.02727 &= -0.52727\end{aligned}$$

利用更新之后的权重，对训练集的每一条数据反复进行前面1~12步的计算，直到最后收敛

■ Step 15 :

With the updated weights $[V]$ and $[W]$, error is calculated again and next training set is taken and the error will then get adjusted.

■ Step 16 :

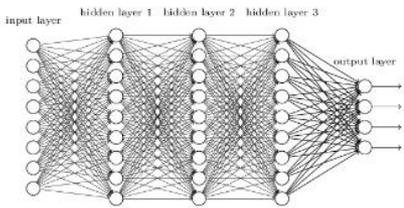
Iterations are carried out till we get the error less than the tolerance.

■ Step 17 :

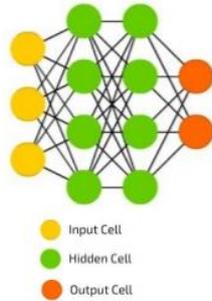
Once the weights are adjusted the network is ready for inferencing new objects .

Neural network

Deep Neural Network (DNN)

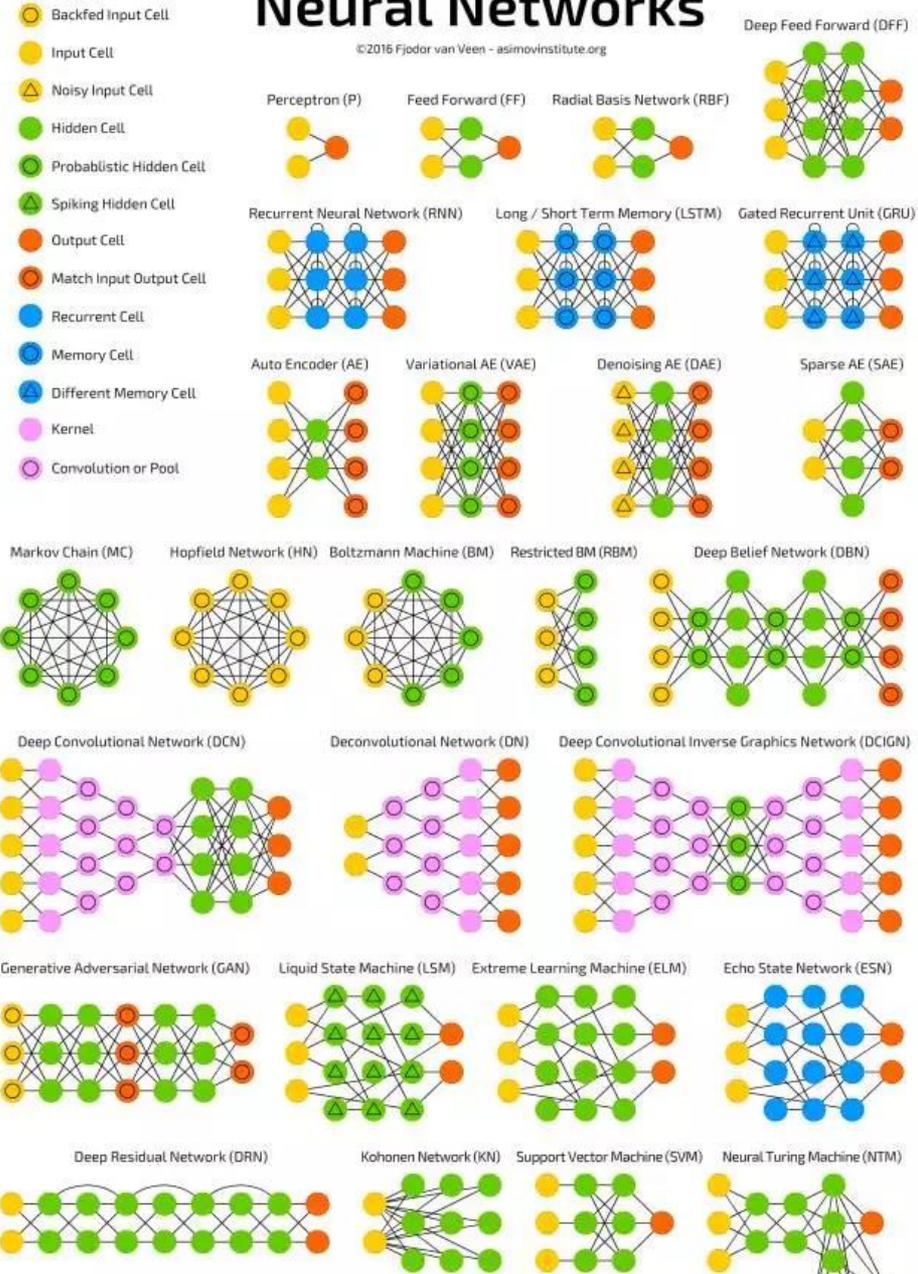


Deep Feed Forward (DFF)



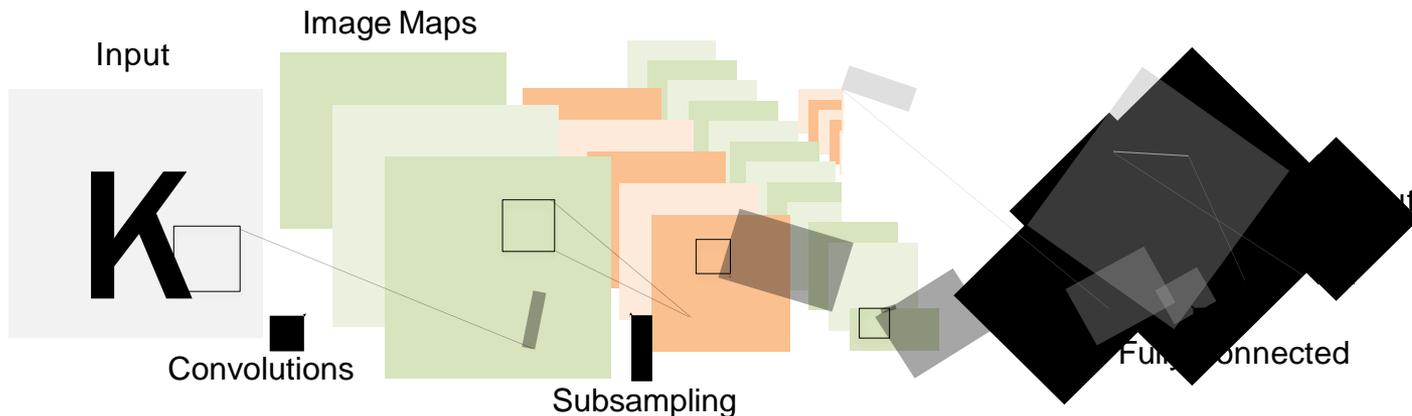
A mostly complete chart of Neural Networks

©2016 Fjodor van Veen - asimovinstitute.org



1998

LeCun et al.



of transistors



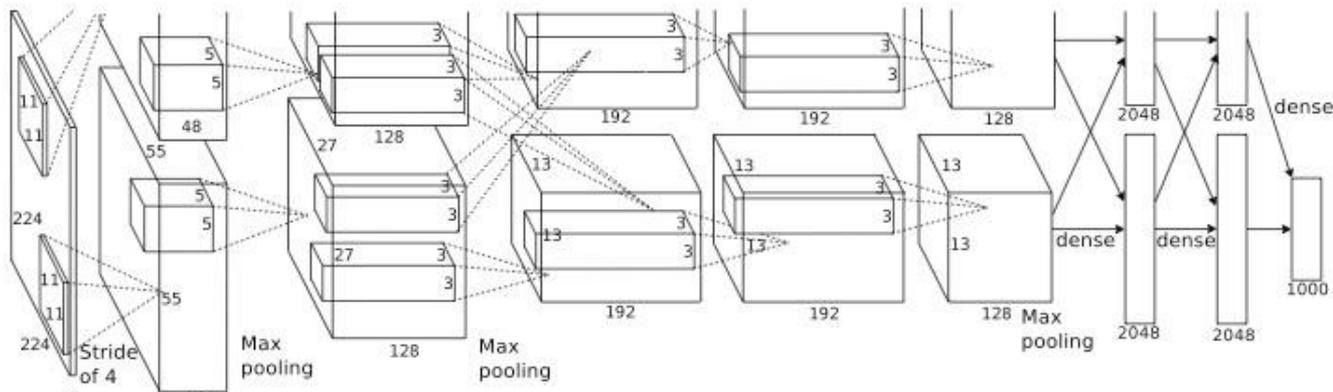
10^6

of pixels used in training

10^7 **NIST**

2012

Krizhevsky et al.



of transistors



10^9

GPUs



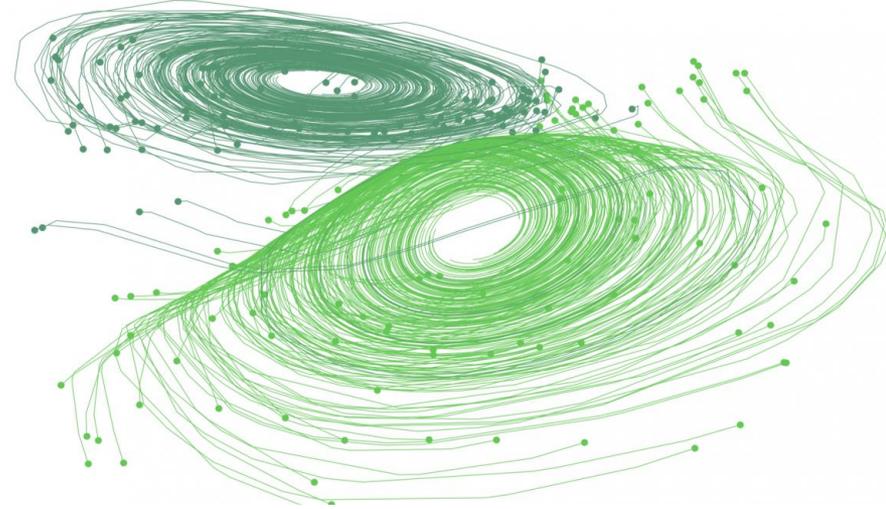
of pixels used in training

10^{14} **IMAGENET**

Figure copyright Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton, 2012. Reproduced with permission.

Convolutional Neural Networks (CNN) were NOT invented overnight!

Neural network



Intelligent Machines

A radical new neural network design could overcome big challenges in AI

Researchers borrowed equations from calculus to redesign the core machinery of deep learning so it can model continuous processes like changes in health.

by Karen Hao December 12, 2018

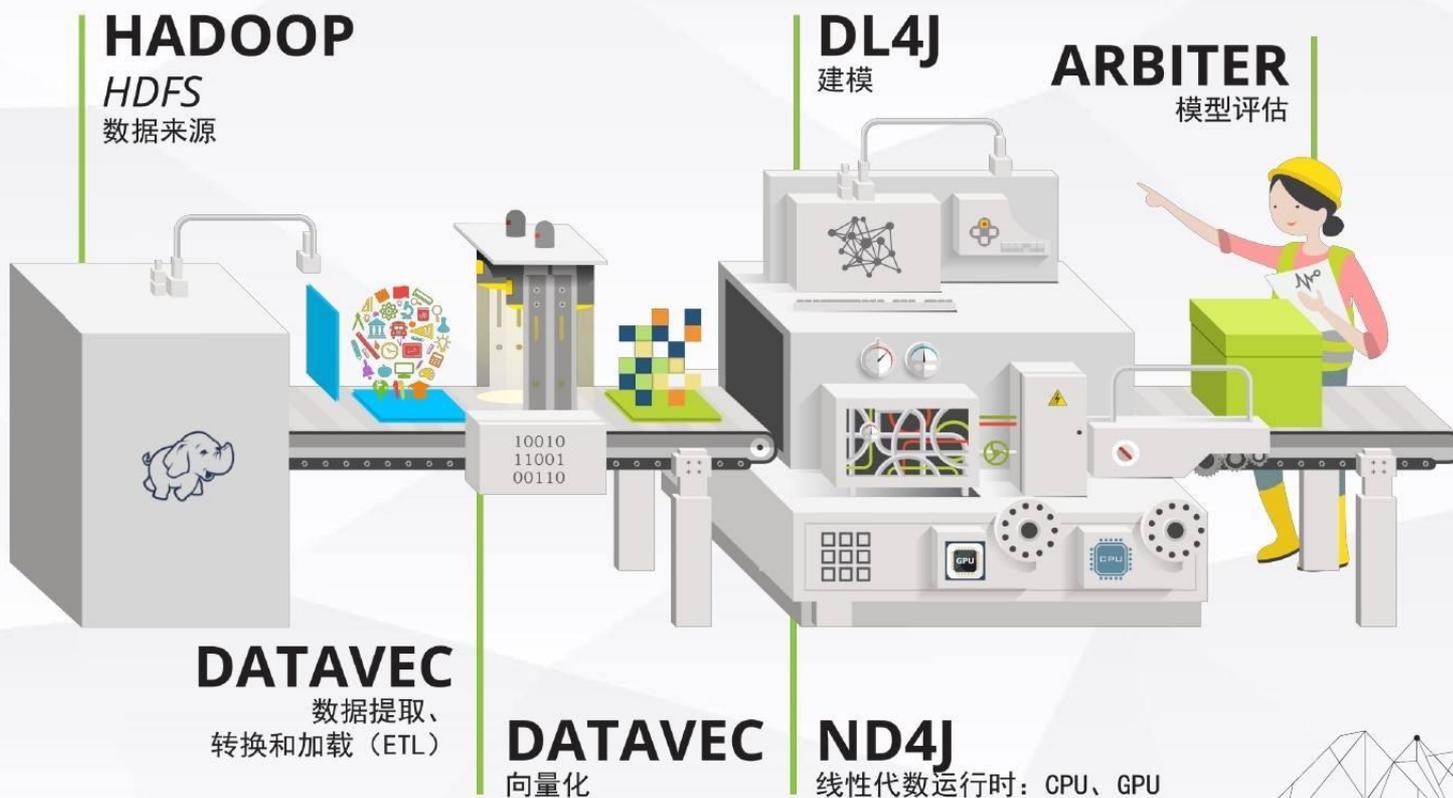
Advertisement

An advertisement for EmTech Asia 2019. The ad features the MIT Technology Review logo and the event title 'EmTech Asia' in large white and red text. Below the title, it says '22 - 23 January 2019 | Singapore'. Two speakers are listed: Jonah Myerberg from Desktop Metal and Samantha Payne from Open Bionics. At the bottom, there is a red button that says 'BOOK NOW!' with a hand cursor icon. A small 'X' icon is in the top right corner of the ad.

深度学习的基本过程



深度学习建模（模型训练）流程



Data-Driven Approach

1. Collect a dataset of images and labels
2. Use Machine Learning to train a classifier
3. Evaluate the classifier on new images

Example training set

```
def train(images, labels):  
    # Machine learning!  
    return model
```

```
def predict(model, test_images):  
    # Use model to predict labels  
    return test_labels
```

airplane



automobile



bird



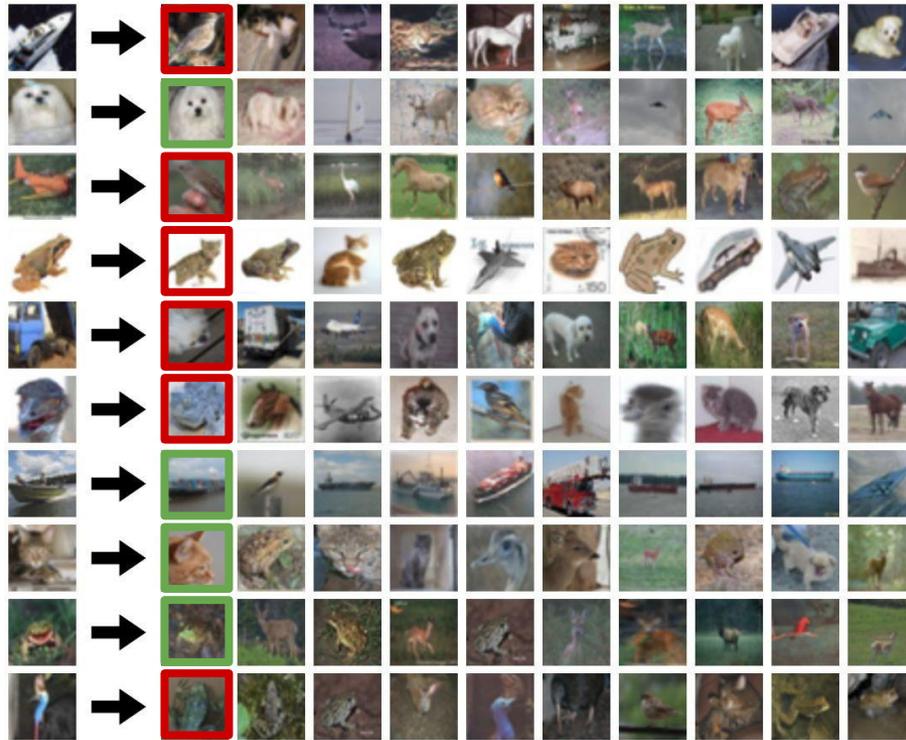
cat



deer



What does this look like?



Parametric Approach: Linear Classifier

3072x1

Image



Array of **32x32x3** numbers
(3072 numbers total)

$$f(x, W) = Wx$$

10x1

10x3072



10 numbers giving
class scores

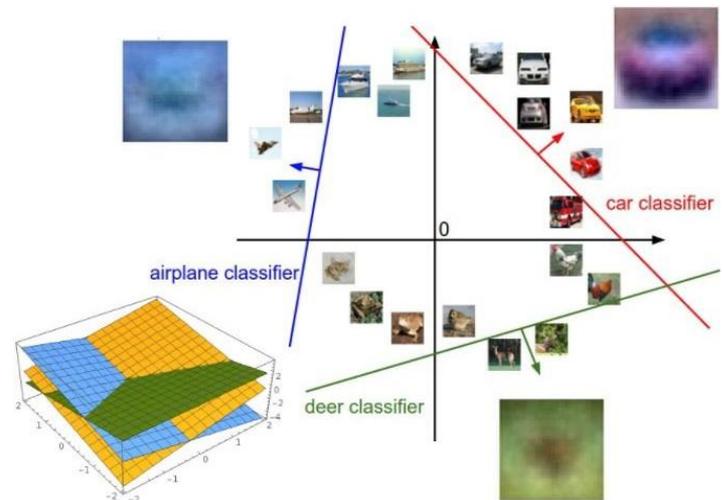
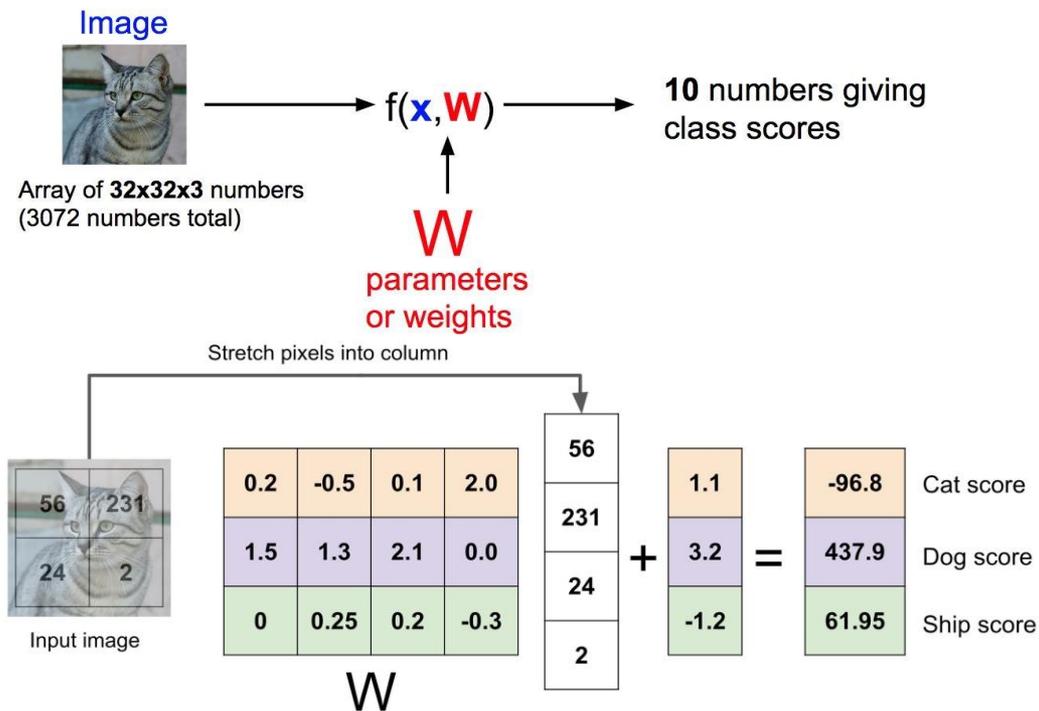


W

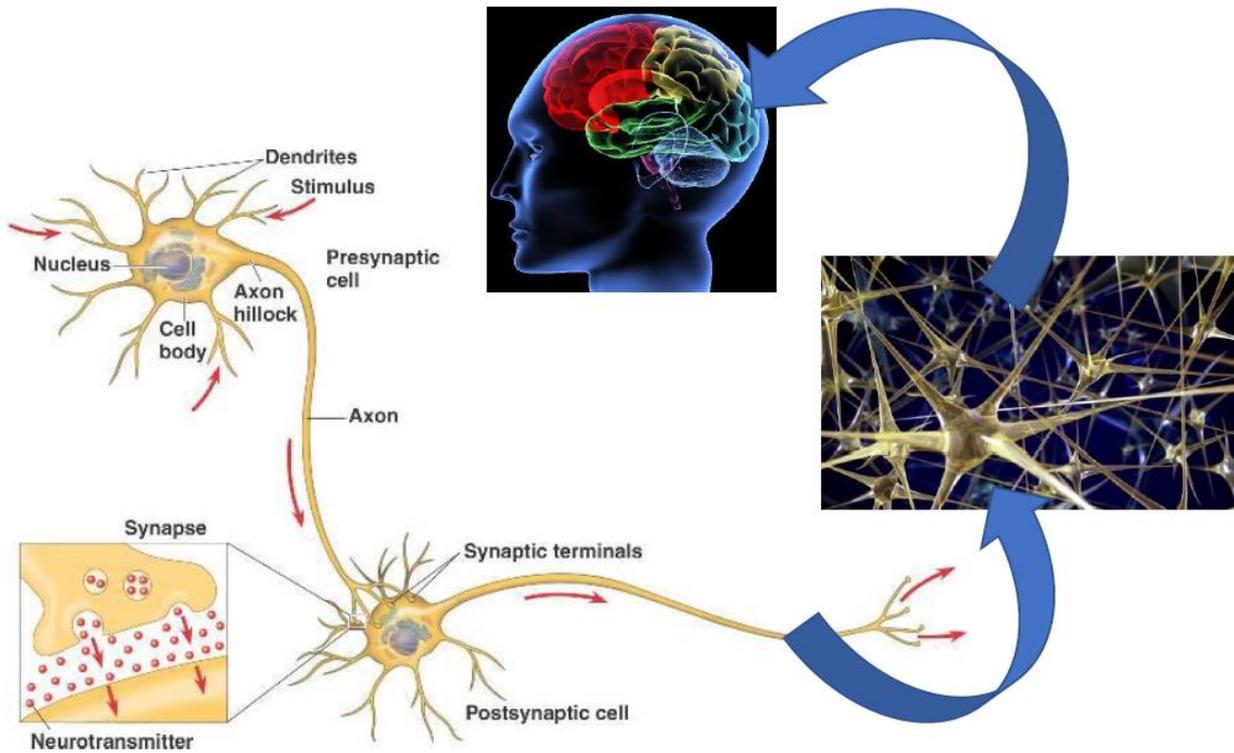
parameters
or weights

Parametric Approach: Linear Classifier

$$f(x, W) = Wx + b$$



Neural Network introduction

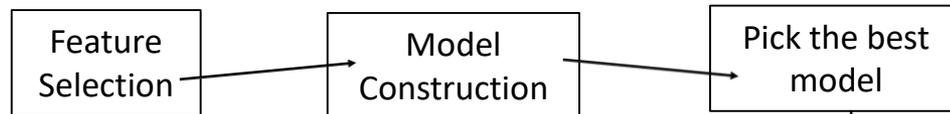
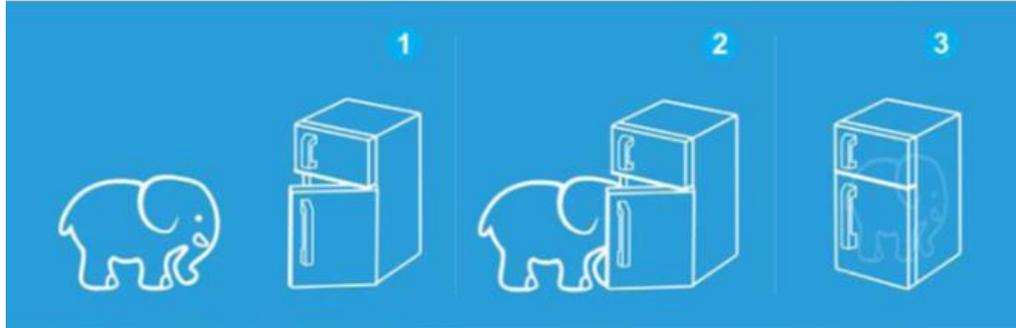


Neural network model is inspired by human neuron

Each neuron is a judgement unit

Neural Network workflow

Neural network is so simple.....



The key point

Deep learning should be a problem-oriented tool.

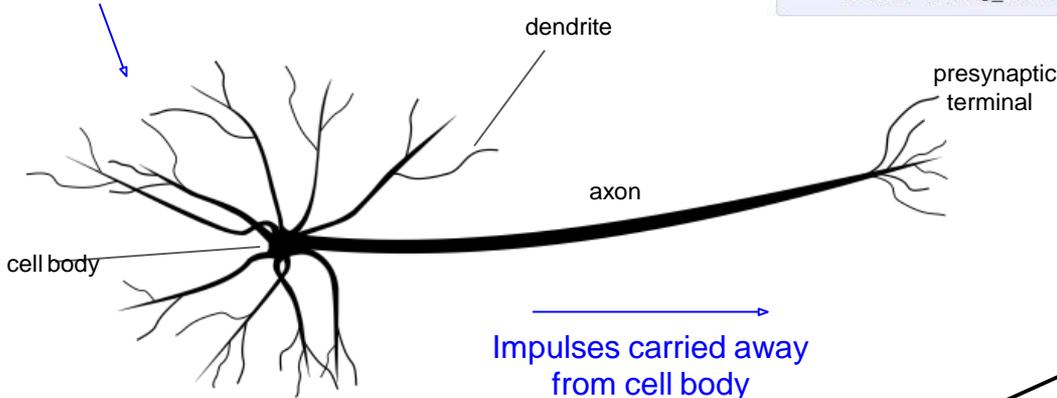
Number of layers
Functions Selection

Parameters Adjustment

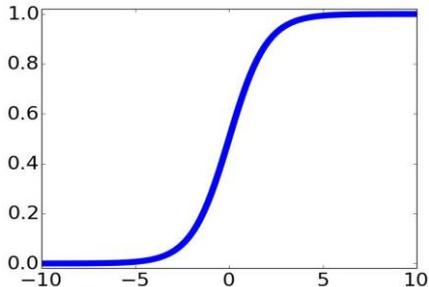
Neuron Network

```
class Neuron:  
    # ...  
    def neuron_tick(inputs):  
        """ assume inputs and weights are 1-D numpy arrays and bias is a number """  
        cell_body_sum = np.sum(inputs * self.weights) + self.bias  
        firing_rate = 1.0 / (1.0 + math.exp(-cell_body_sum)) # sigmoid activation function  
        return firing_rate
```

Impulses carried toward cell body

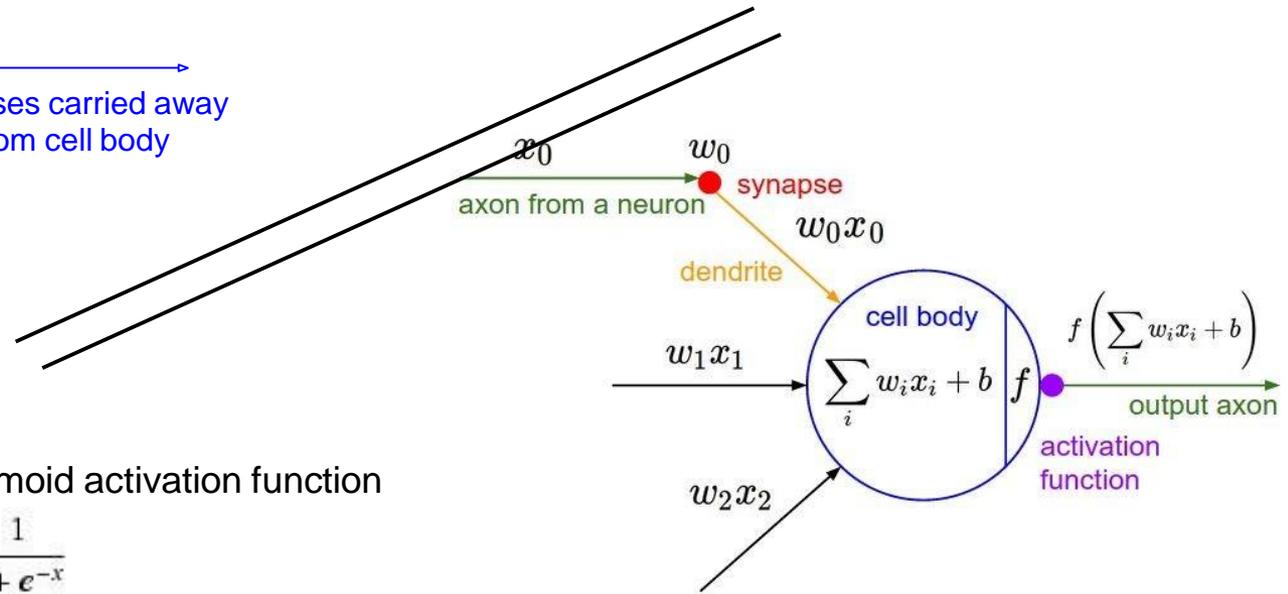


[This image](#) by Felipe Perucho is licensed under [CC-BY3.0](#)



sigmoid activation function

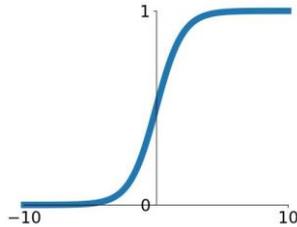
$$\frac{1}{1 + e^{-x}}$$



Activation Functions

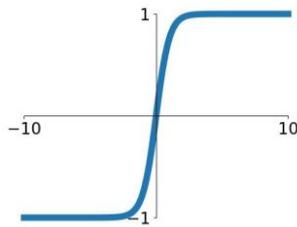
Sigmoid

$$\sigma(x) = \frac{1}{1+e^{-x}}$$



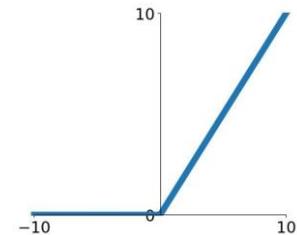
tanh

$$\tanh(x)$$



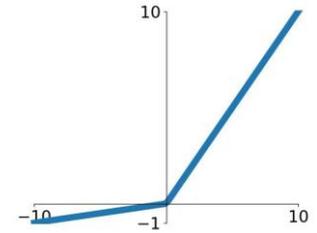
ReLU

$$\max(0, x)$$



Leaky ReLU

$$\max(0.1x, x)$$

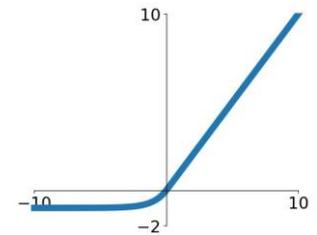


Maxout

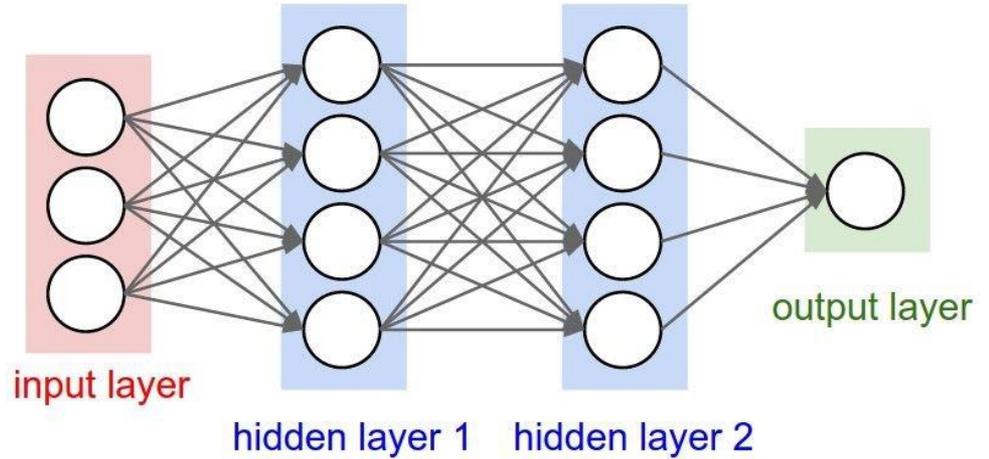
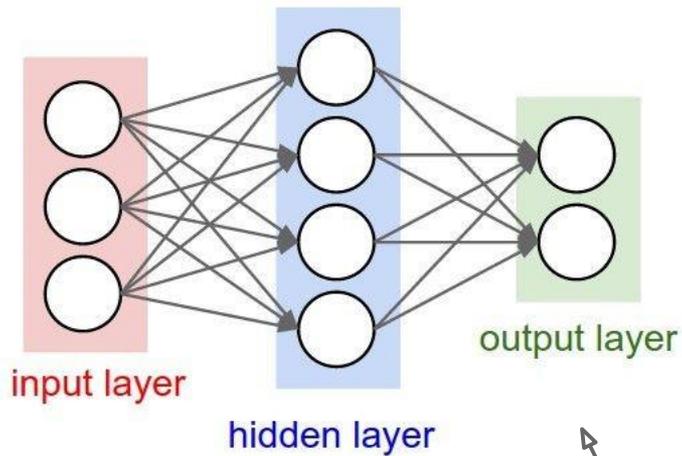
$$\max(w_1^T x + b_1, w_2^T x + b_2)$$

ELU

$$\begin{cases} x & x \geq 0 \\ \alpha(e^x - 1) & x < 0 \end{cases}$$



Neural networks: Architectures



“2-layer Neural Net”, or
“1-hidden-layer Neural Net”

“3-layer Neural Net”, or
“2-hidden-layer Neural Net”

“Fully-connected” layers

Convolutional Neural Networks

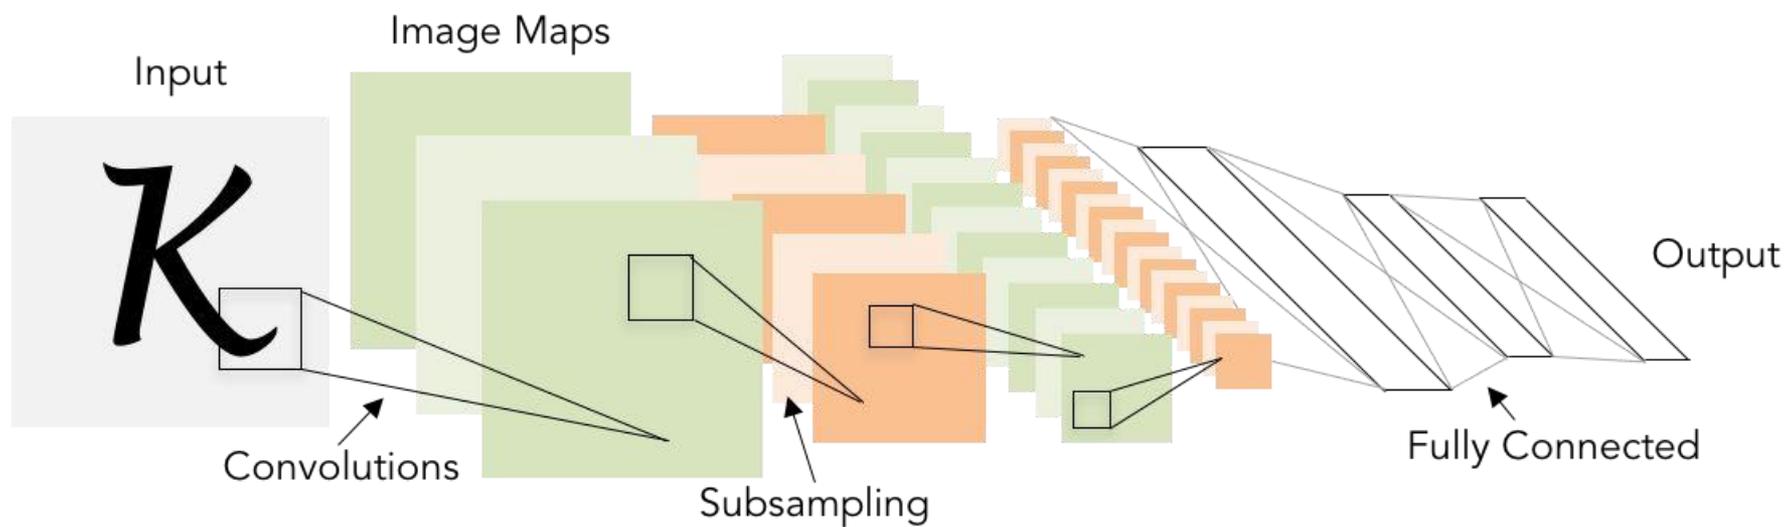
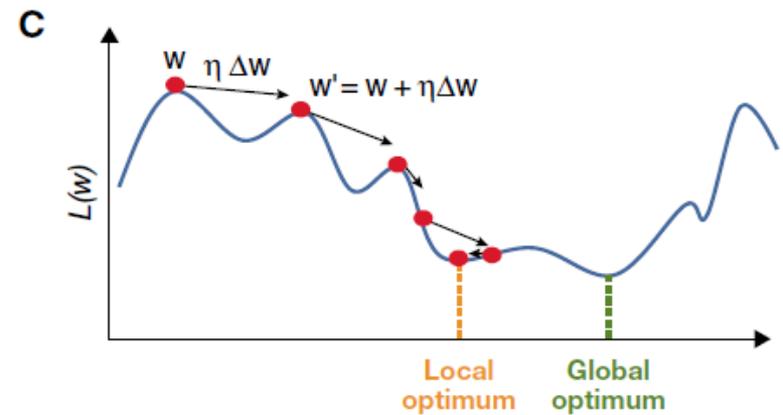
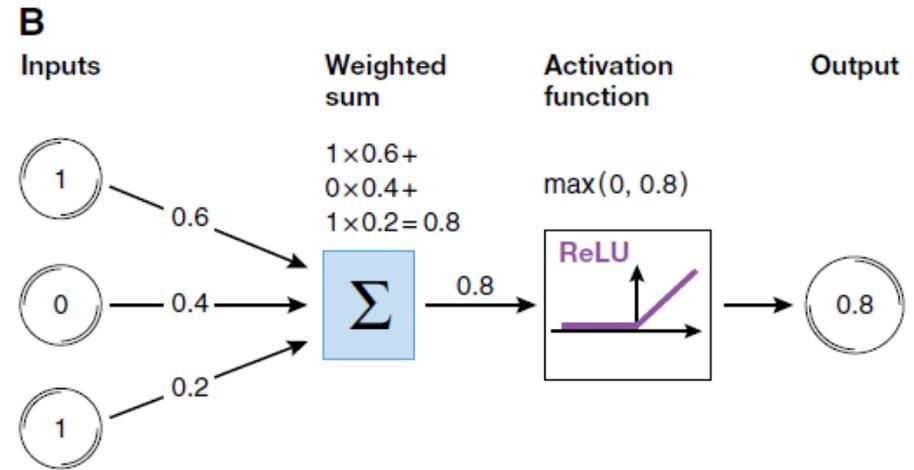
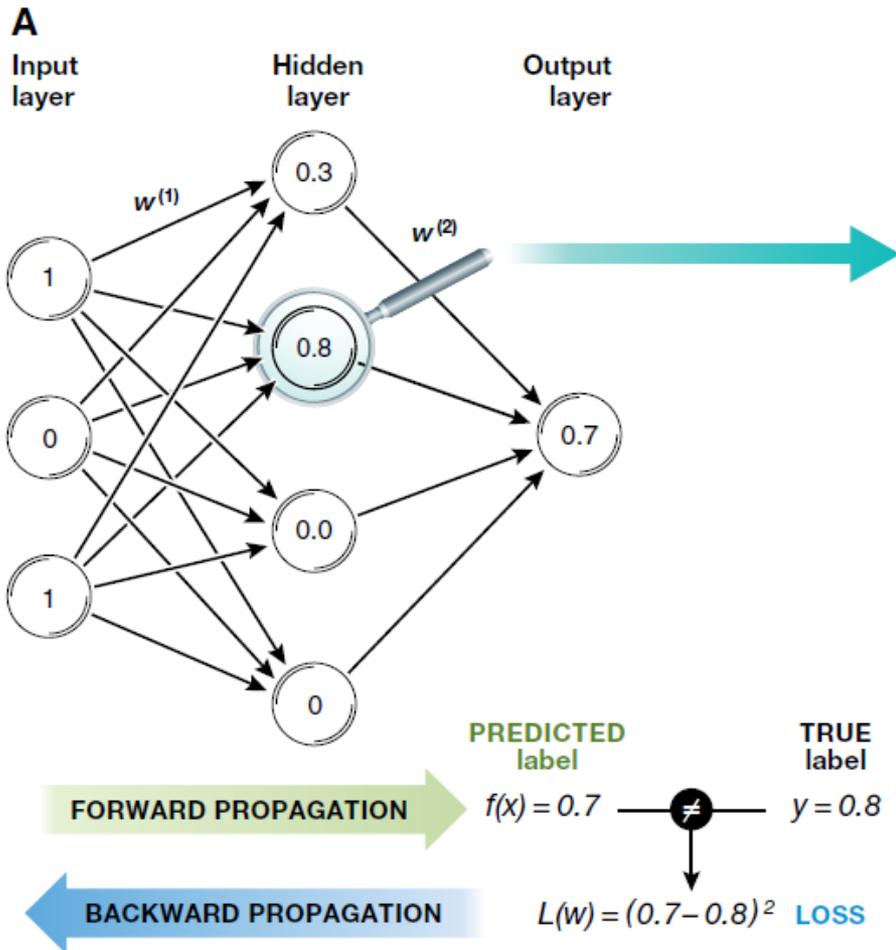


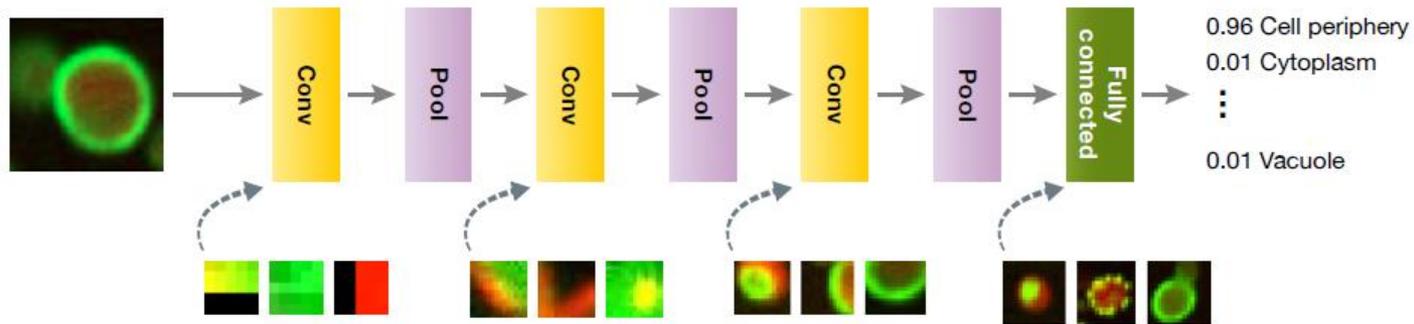
Illustration of LeCun et al. 1998 from CS231n 2017 Lecture 1

Artificial Neural Network

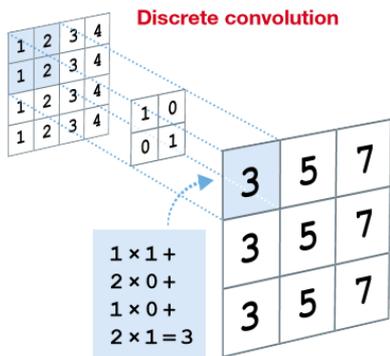


Convolution Neural Network

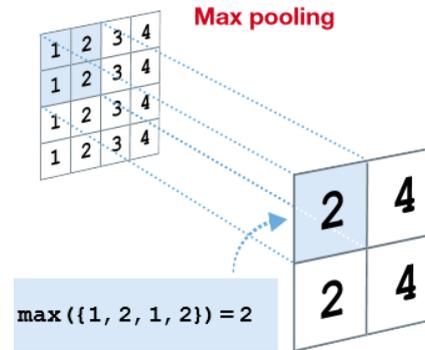
Convolution neural network is similar to neural network, it is a good approach to process figure



Convolution and pooling process : edge features extraction



Convolution



Pooling

Convolution Neural Network

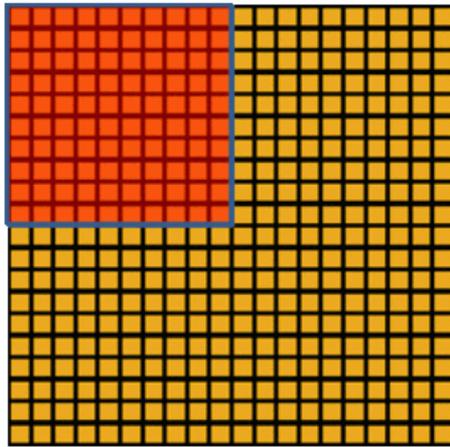
1 _{x1}	1 _{x0}	1 _{x1}	0	0
0 _{x0}	1 _{x1}	1 _{x0}	1	0
0 _{x1}	0 _{x0}	1 _{x1}	1	1
0	0	1	1	0
0	1	1	0	0

Image

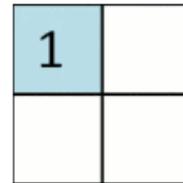
4		

Convolved
Feature

Convolution Neural Network



Convolved
feature



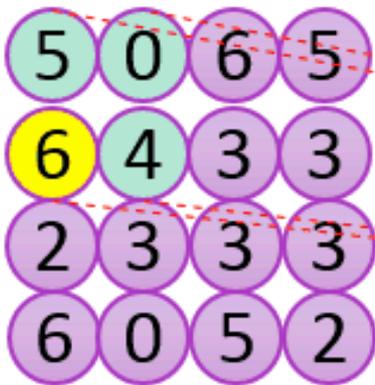
Pooled
feature

Convolution Neural Network

Input
(4*4)

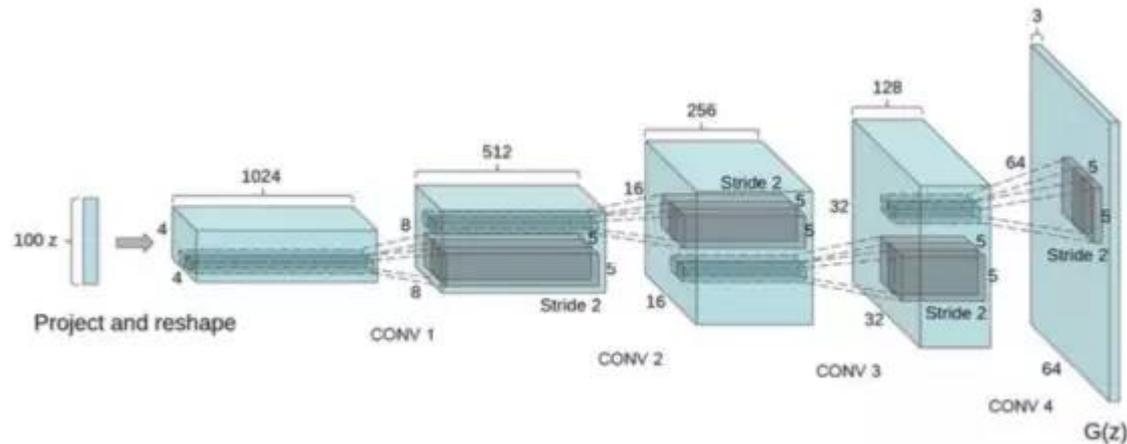
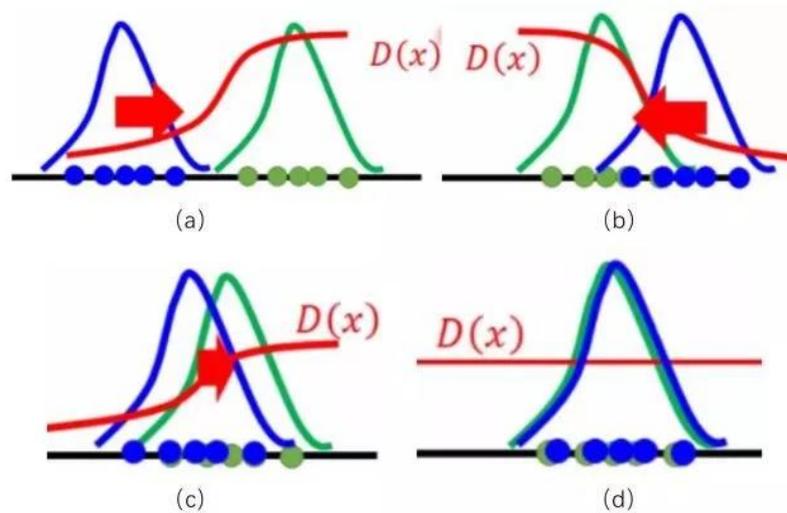
Kernal
(2*2)

Output

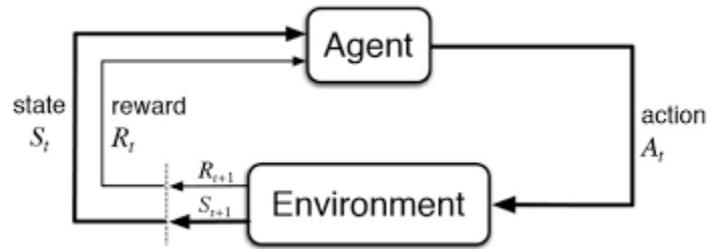


Pooling...

Generative adversarial network, GAN

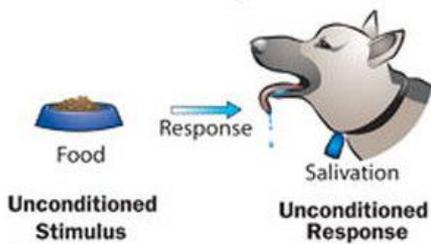


Reinforcement learning

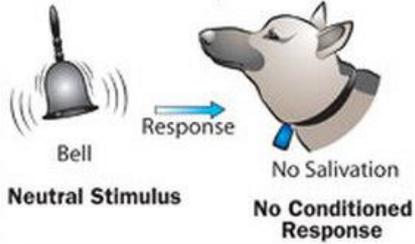


How Dog Training Works

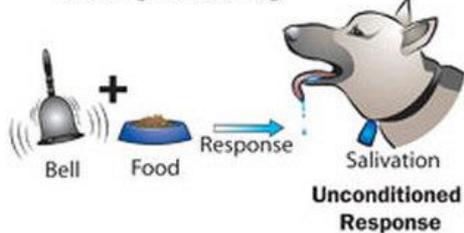
1. Before Conditioning



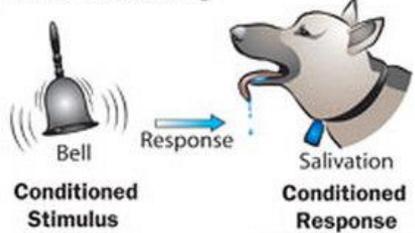
2. Before Conditioning



3. During Conditioning

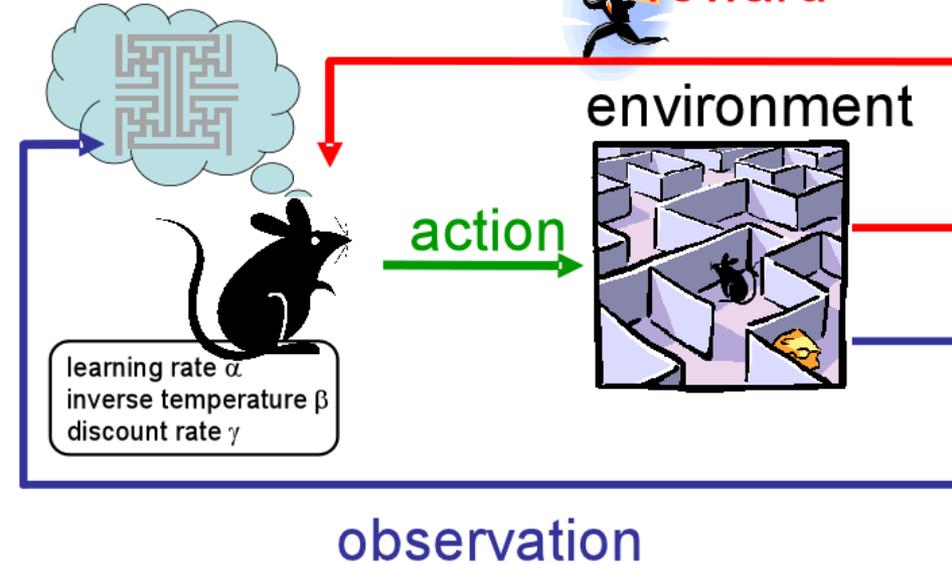


4. After Conditioning

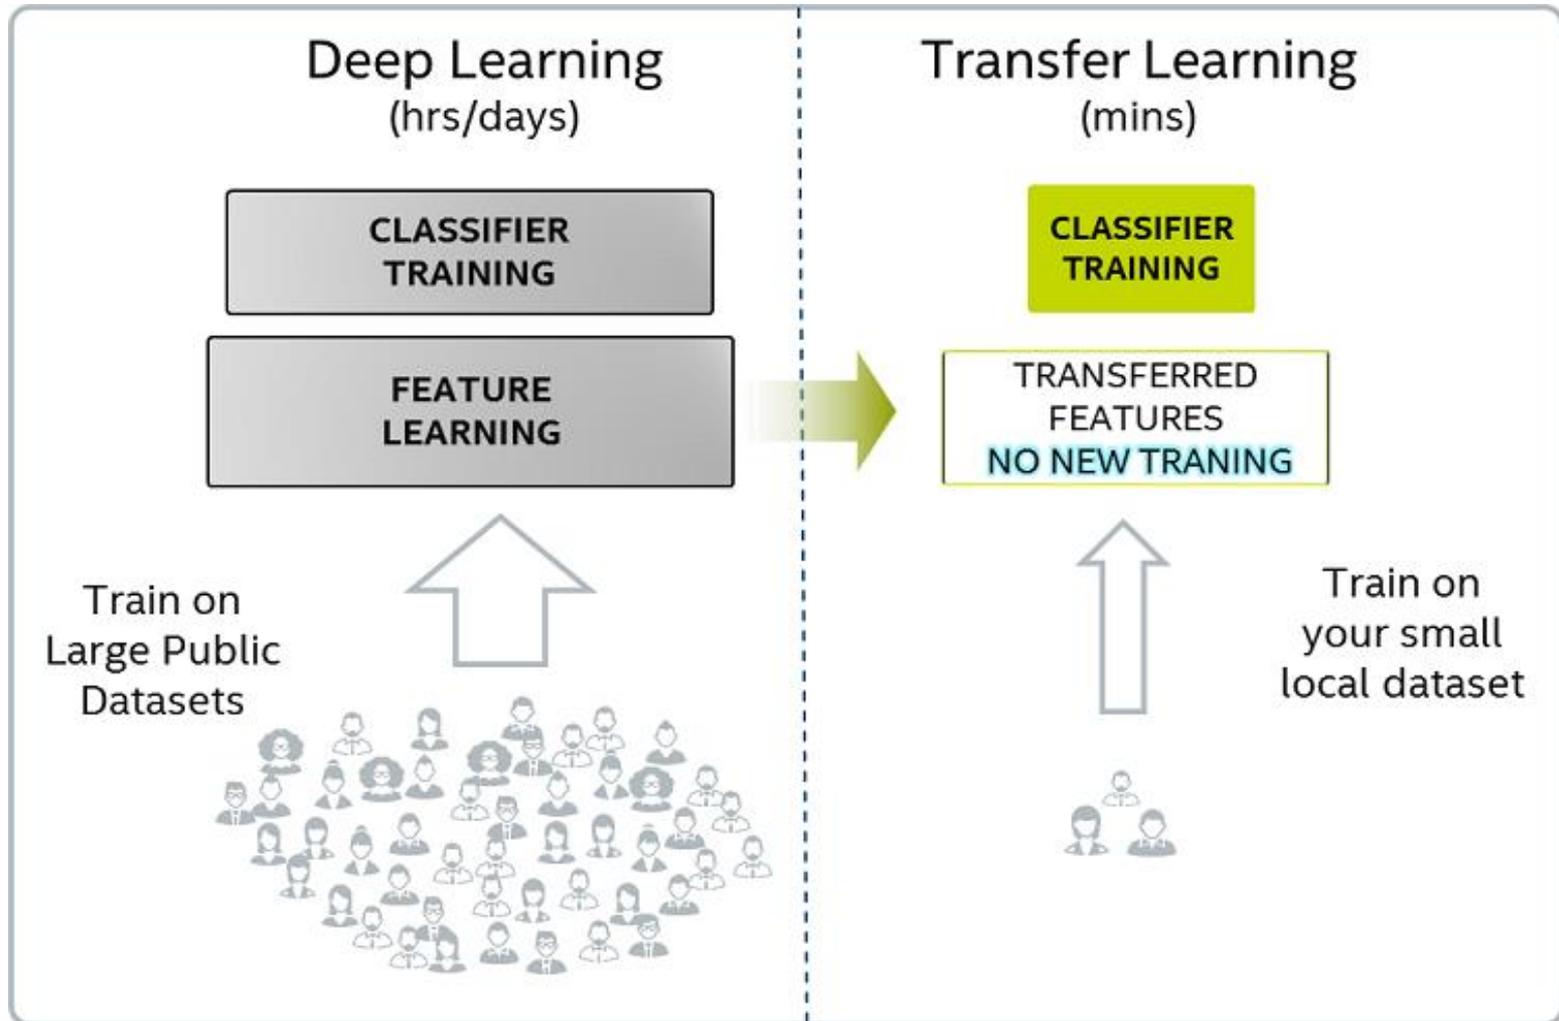


©2006 HowStuffWorks

internal state



Transfer learning



深度学习的特点

深度学习常用算法介绍

深度学习常用框架介绍

深度学习

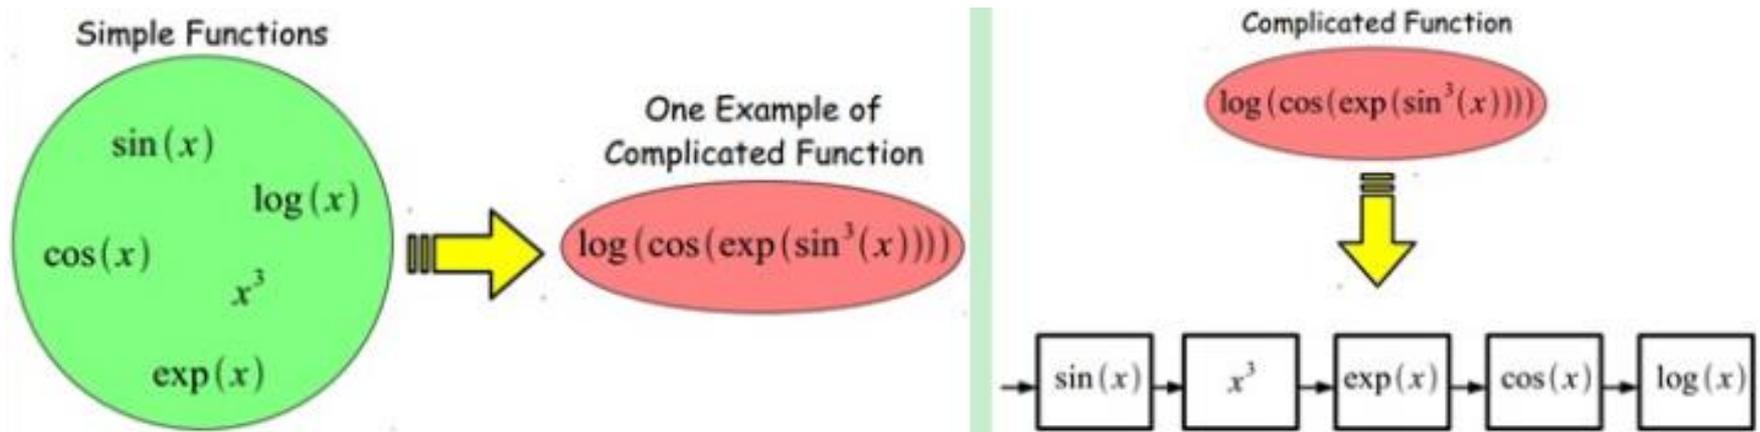
- 2006年，加拿大多伦多大学教授、机器学习领域的泰斗Geoffrey Hinton在《科学》上发表论文提出深度学习主要观点：
 - 1) 多隐层的人工神经网络具有优异的特征学习能力，学习得到的特征对数据有更本质的刻画，从而有利于可视化或分类；
 - 2) 深度神经网络在训练上的难度，可以通过“逐层初始化”（layer-wise pre-training）来有效克服，逐层初始化可通过无监督学习实现的。

深度学习

- **本质：**通过构建多隐层的模型和海量训练数据（可为无标签数据），来学习更有用的特征，从而最终提升分类或预测的准确性。“深度模型”是手段，“特征学习”是目的。
- **与浅层学习区别：**
 - 1) 强调了模型结构的深度，通常有5-10多层的隐层节点；
 - 2) 明确突出了特征学习的重要性，通过逐层特征变换，将样本在原空间的特征表示变换到一个新特征空间，从而使分类或预测更加容易。与人工规则构造特征的方法相比，利用大数据来学习特征，更能够刻画数据的丰富内在信息。

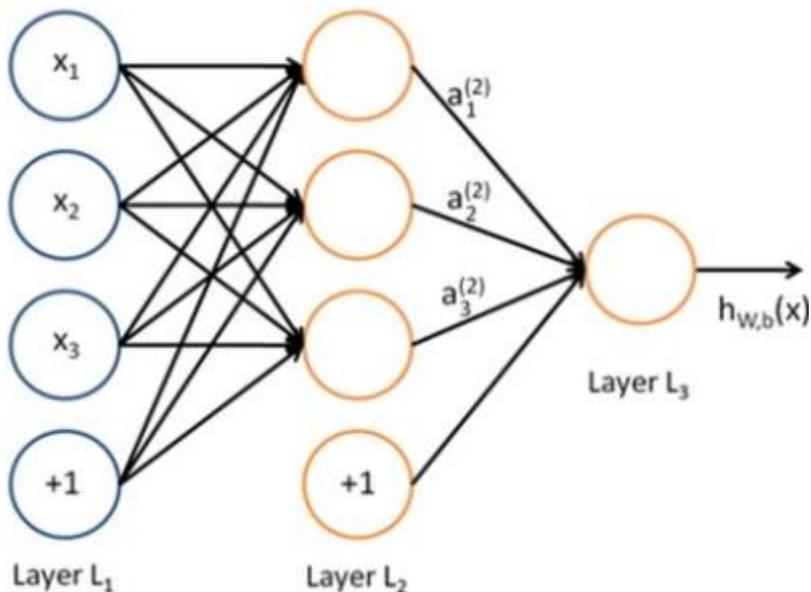
深度学习

- 好处：可通过学习一种深层非线性网络结构，实现复杂函数逼近，表征输入数据分布式表示。

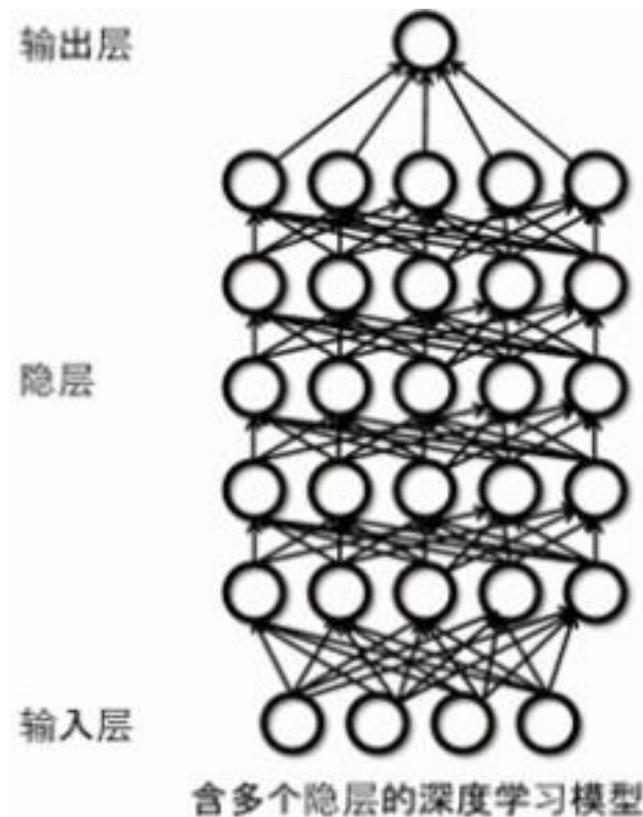


深度学习 vs. 神经网络

神经网络 :



深度学习 :



深度学习 vs. 神经网络

相同点：二者均采用分层结构，系统包括输入层、隐层（多层）、输出层组成的多层网络，只有相邻层节点之间有连接，同一层以及跨层节点之间相互无连接，每一层可以看作是一个logistic 回归模型。

不同点：

神经网络：采用BP算法调整参数，即采用迭代算法来训练整个网络。随机设定初值，计算当前网络的输出，然后根据当前输出和样本真实标签之间的差去改变前面各层的参数，直到收敛；

深度学习：采用逐层训练机制。采用该机制的原因在于如果采用BP机制，对于一个deep network（7层以上），残差传播到最前面的层将变得很小，出现所谓的gradient diffusion（梯度扩散）。

深度学习 vs. 神经网络

- 神经网络的局限性:
 - 1) 比较容易过拟合，参数比较难调整，而且需要不少技巧；
 - 2) 训练速度比较慢，在层次比较少（小于等于3）的情况下效果并不比其它方法更优；

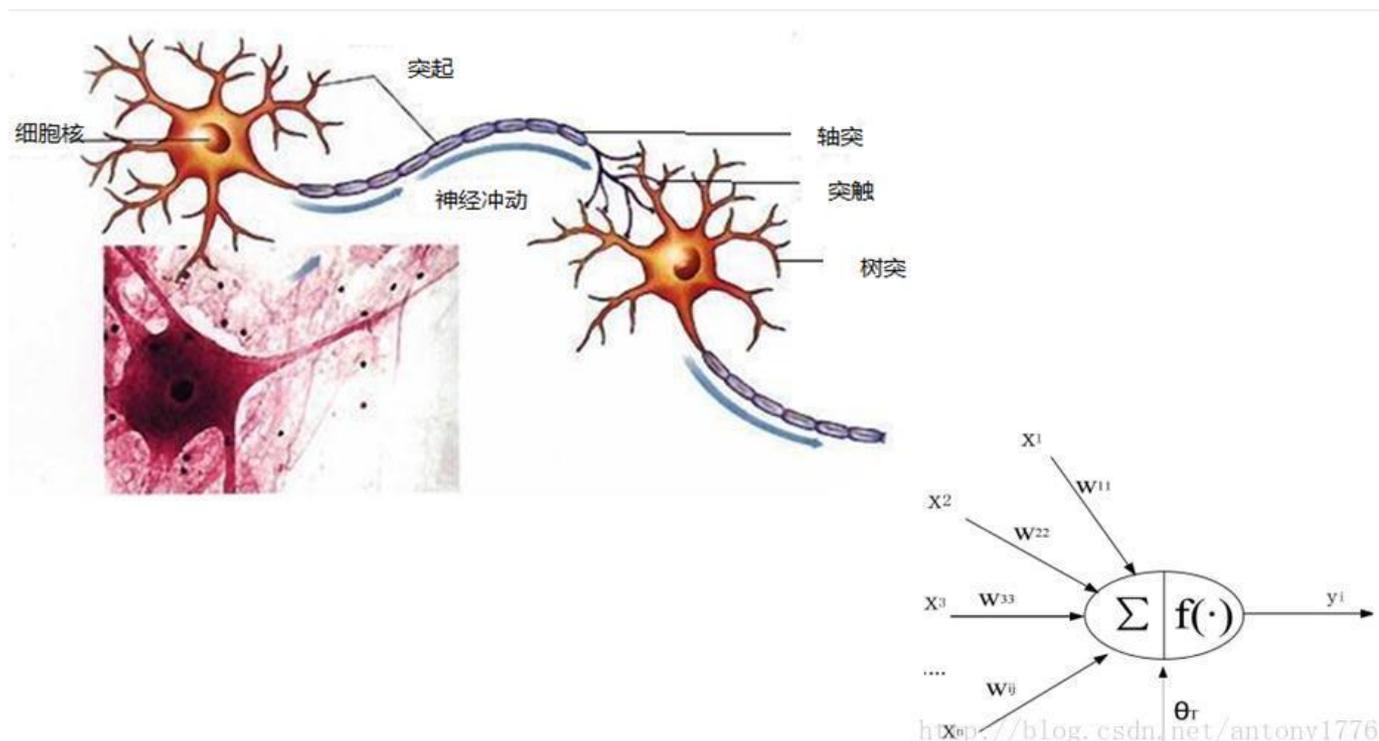
深度学习的特点

深度学习常用算法介绍

深度学习常用框架介绍

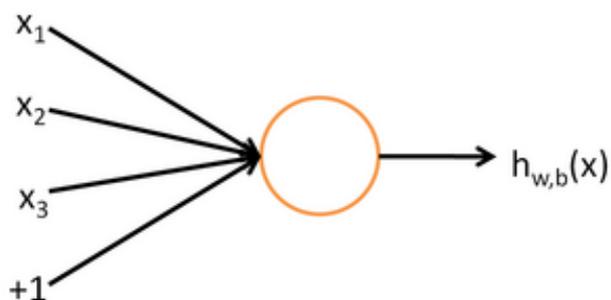
人工神经网络 (ANN)

- 人工神经网络 (Artificial Neural Networks) 是一种模仿生物神经网络行为特征，进行分布式并行信息处理的算法数学模型。这种网络依靠系统的复杂程度，通过调整内部大量节点（神经元）之间相互连接的权重，从而达到处理信息的目的。

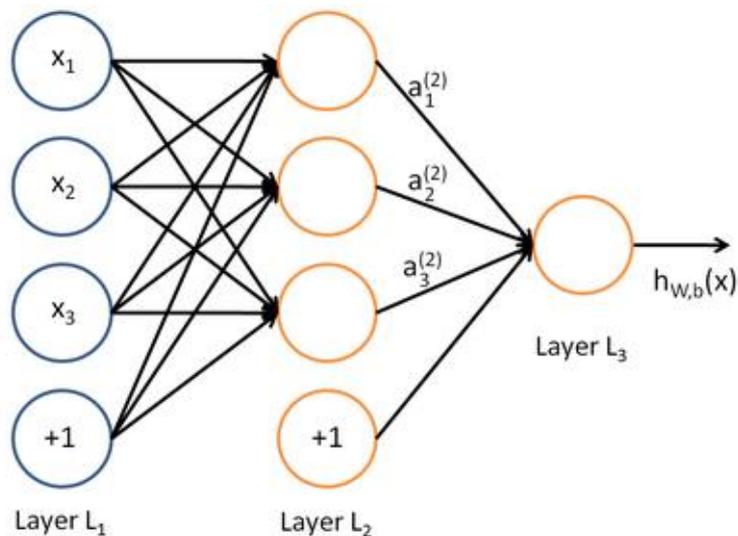


人工神经网络 (ANN)

- 神经网络



$$h_{W,b}(x) = f(W^T x) = f(\sum_{i=1}^3 W_i x_i + b)$$



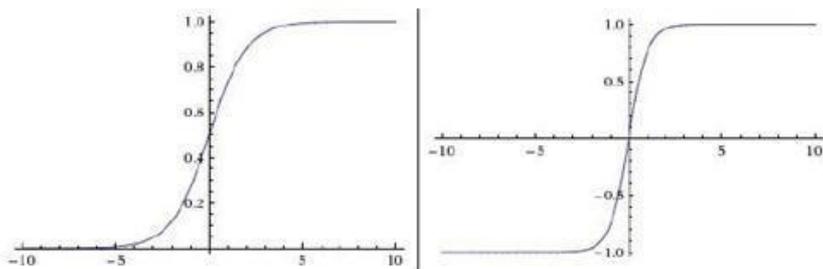
$$\begin{aligned} a_1^{(2)} &= f(W_{11}^{(1)} x_1 + W_{12}^{(1)} x_2 + W_{13}^{(1)} x_3 + b_1^{(1)}) \\ a_2^{(2)} &= f(W_{21}^{(1)} x_1 + W_{22}^{(1)} x_2 + W_{23}^{(1)} x_3 + b_2^{(1)}) \\ a_3^{(2)} &= f(W_{31}^{(1)} x_1 + W_{32}^{(1)} x_2 + W_{33}^{(1)} x_3 + b_3^{(1)}) \\ h_{W,b}(x) &= a_1^{(3)} = f(W_{11}^{(2)} a_1^{(2)} + W_{12}^{(2)} a_2^{(2)} + W_{13}^{(2)} a_3^{(2)} + b_1^{(2)}) \end{aligned}$$

人工神经网络 (ANN)

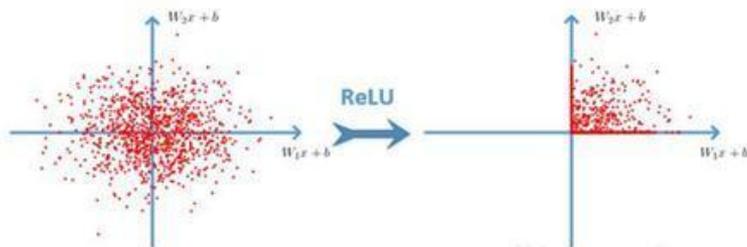
- 人工神经网络的重要概念:

1 权值矩阵: 相当于神经网络的记忆! 在训练的过程中, 动态调整和适应。

2 激励函数:

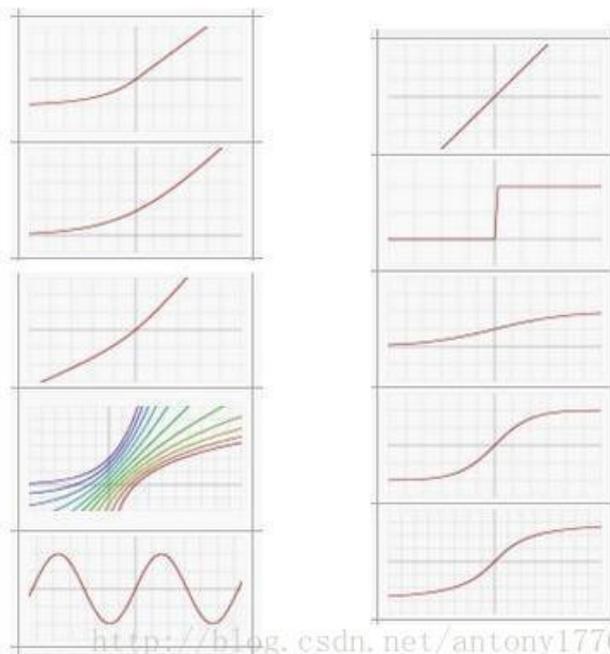


Sigmoid



Relu

Others



<http://blog.csdn.net/antony1776>

人工神经网络 (ANN)

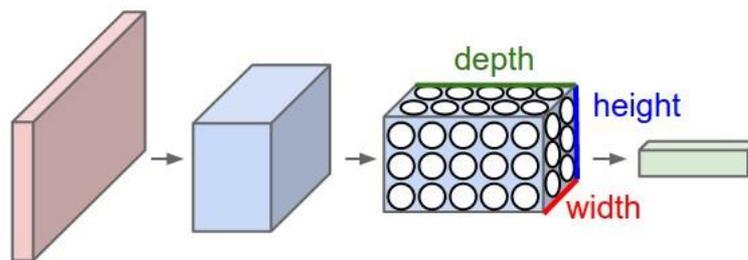
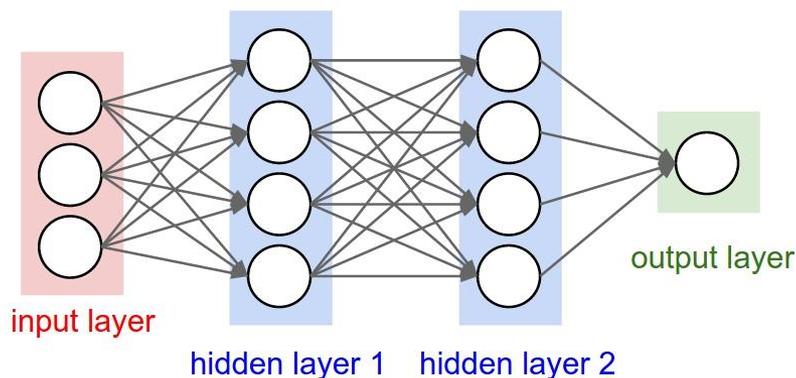
- 人工神经网络的重要概念:

激励函数很重要，无论是对建立神经网络的模型，还是理解神经网络。首先要了解，它有以下几个影响：

- 1 如何能更好的求解目标函数的极值！——高等数学中求解函数极值的知识！
可微，单调！
- 2 如何提升训练效率，让梯度的优化方法更稳定；
- 3 权值的初始值，不影响训练结果！

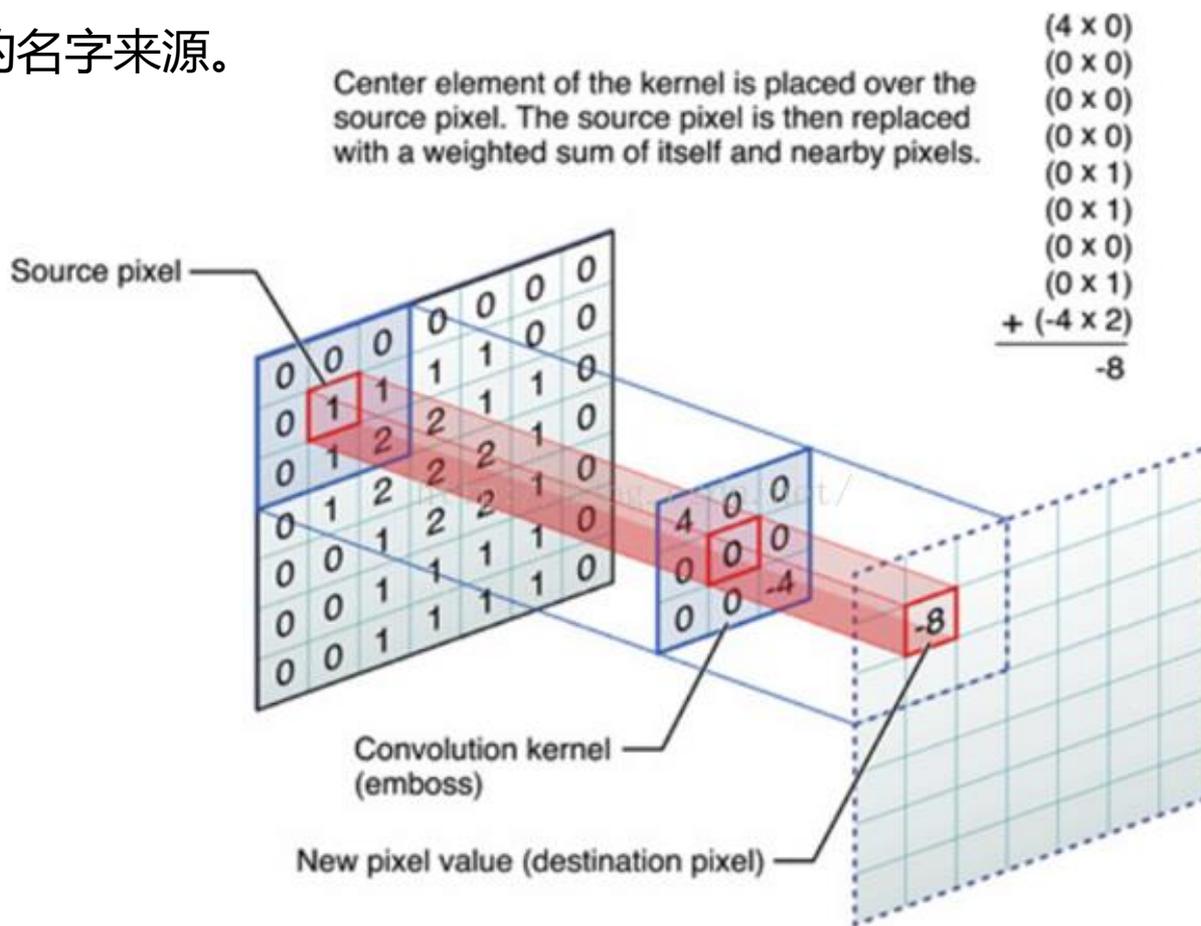
卷积神经网络 (CNN)

- 卷积神经网络 (Convolutional Neural Networks / CNNs / ConvNets) 与普通神经网络非常相似，它们都由具有可学习的权重和偏置常量(biases)的神经元组成。每个神经元都接收一些输入，并做一些点积计算，输出是每个分类的分数，普通神经网络里的一些计算技巧到这里依旧适用。
- 与普通神经网络不同之处：卷积神经网络默认输入是图像，可以让我们把特定的性质编码入网络结构，使是我们的前馈函数更加有效率，并减少了大量参数。

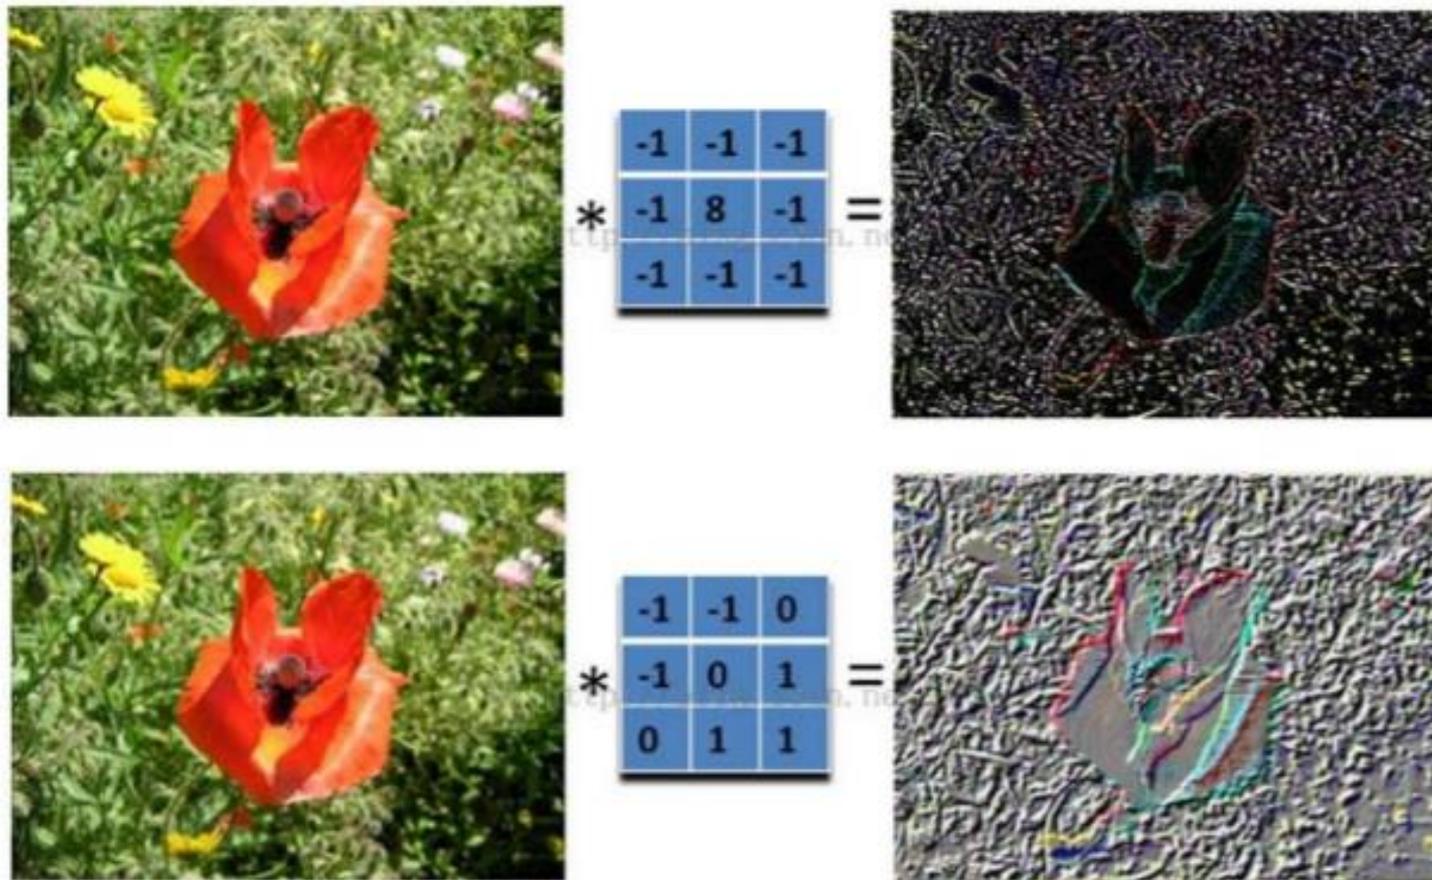


卷积神经网络 (CNN)

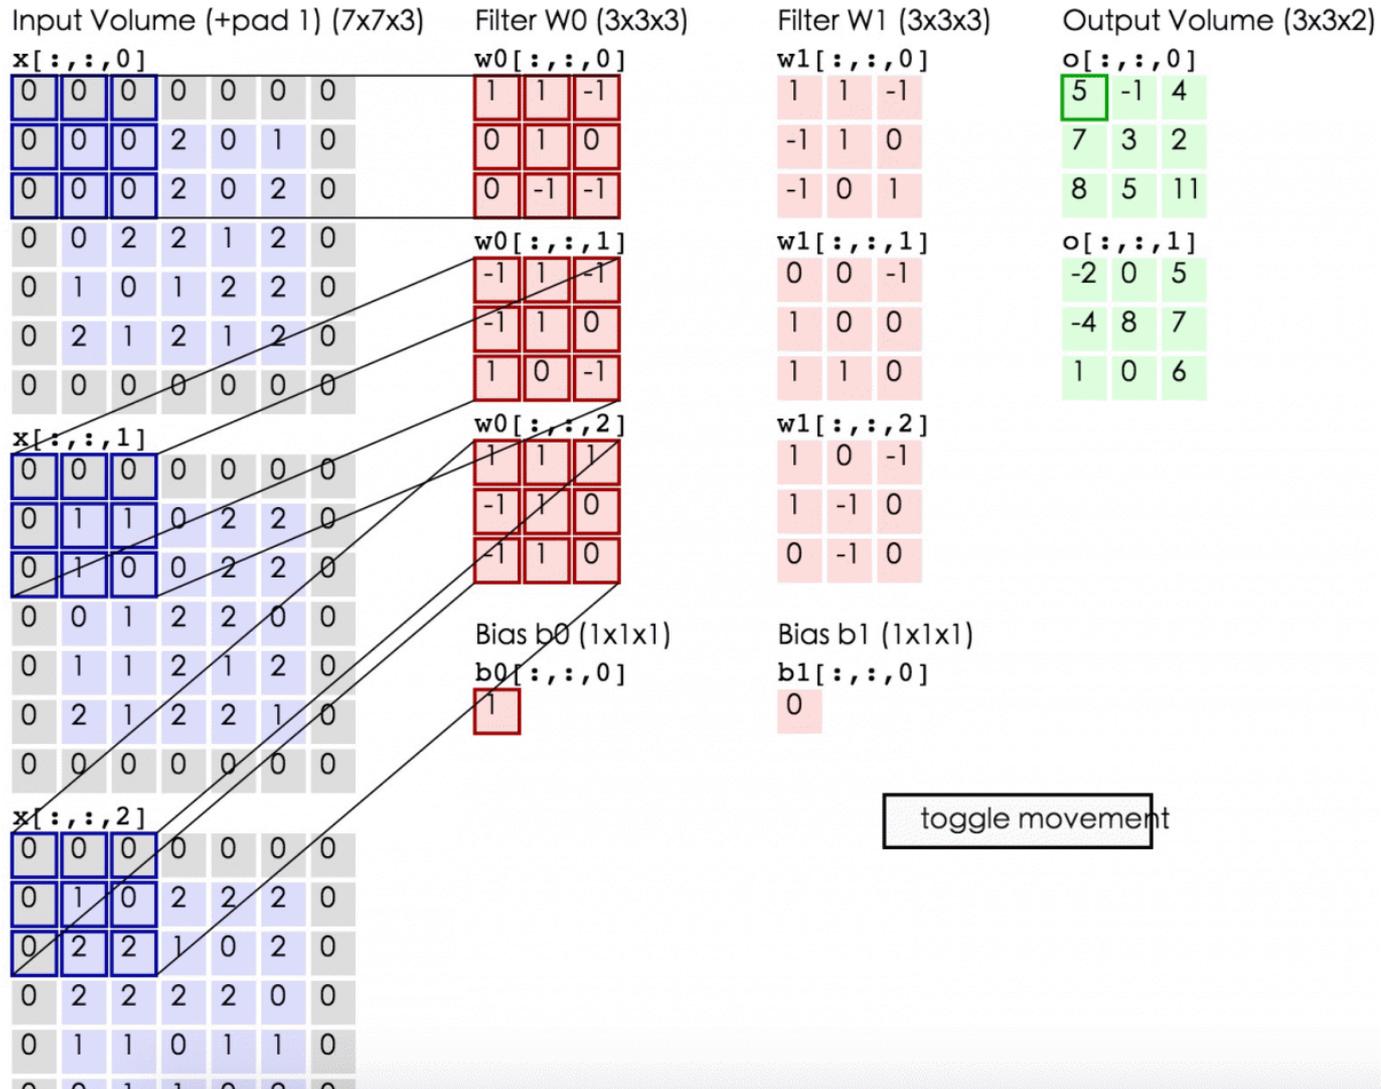
- 卷积操作：对图像（不同的数据窗口数据）和滤波矩阵（一组固定的权重：因为每个神经元的权重固定，所以又可以看做一个恒定的滤波器filter）做内积（逐个元素相乘再求和）的操作就是所谓的『卷积』操作，也是卷积神经网络的名字来源。



卷积神经网络 (CNN)



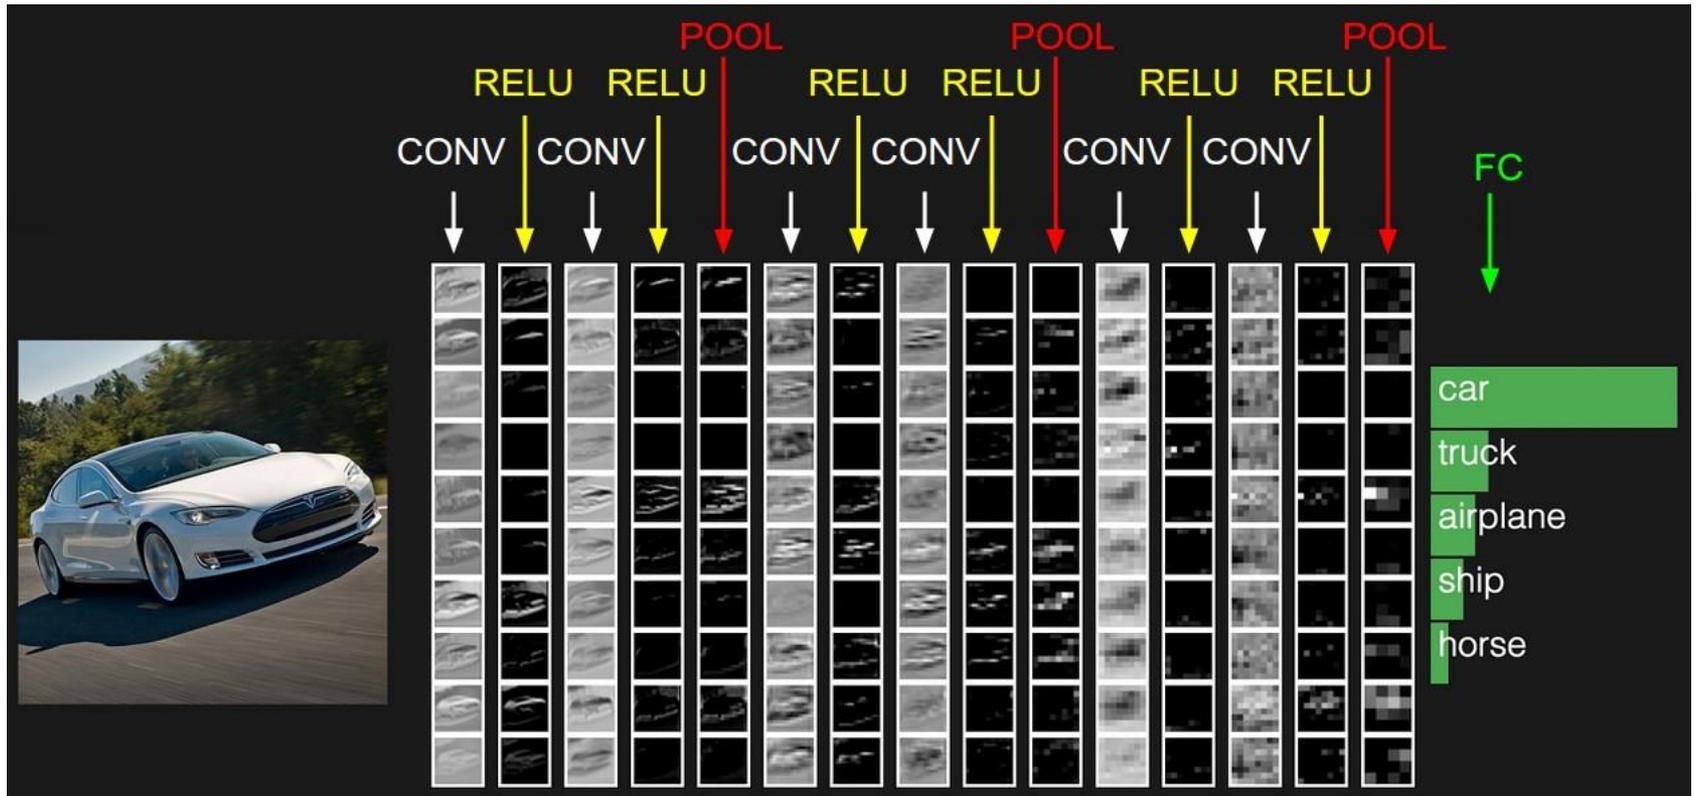
卷积神经网络 (CNN)



卷积神经网络 (CNN)

- 卷积层 (Convolutional layer) , 卷积神经网络中每层卷积层由若干卷积单元组成, 每个卷积单元的参数都是通过反向传播算法优化得到的。卷积运算的目的是提取输入的不同特征, 第一层卷积层可能只能提取一些低级的特征如边缘、线条和角等层级, 更多层的网络能从低级特征中迭代提取更复杂的特征。
- 线性整流层 (Rectified Linear Units layer, ReLU layer) , 这一层神经的激励函数 (Activation function) 使用线性整流 (Rectified Linear Units, ReLU) $f(x)=\max(0,x)$ 。
- 池化层 (Pooling layer) , 通常在卷积层之后会得到维度很大的特征, 将特征切成几个区域, 取其最大值或平均值, 得到新的、维度较小的特征。
- 全连接层 (Fully-Connected layer) , 把所有局部特征结合变成全局特征, 用来计算最后每一类的得分。

卷积神经网络 (CNN)

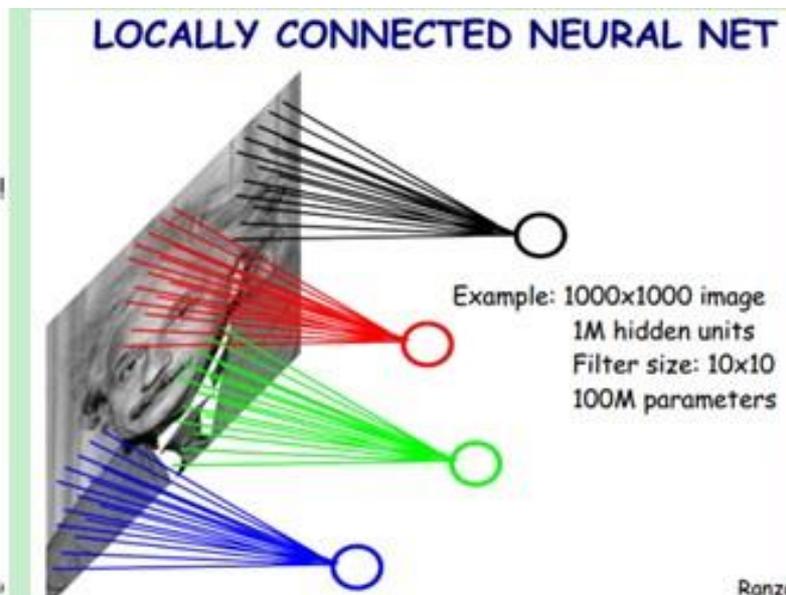
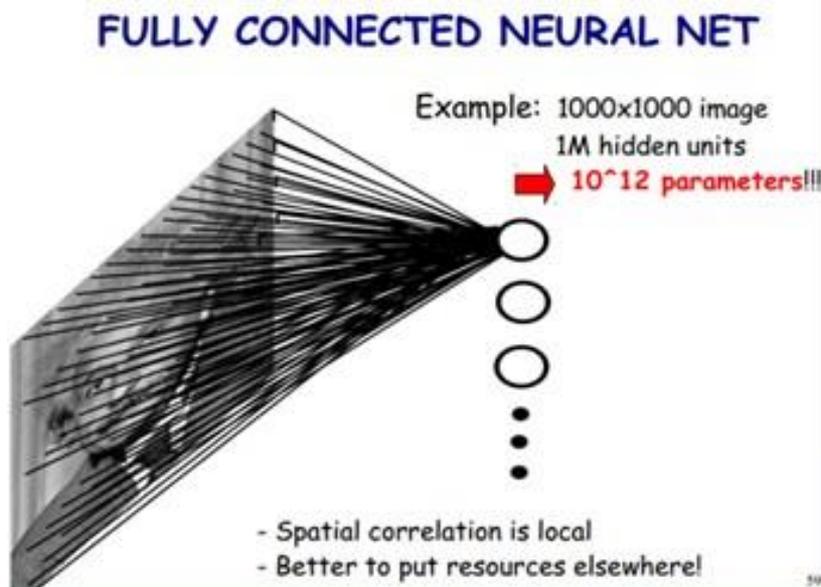


卷积神经网络 (CNN)

- 在图像处理中，往往把图像表示为像素的向量，比如一个 1000×1000 的图像，可以表示为一个 1000000 的向量。在上一节中提到的神经网络中，如果隐含层数目与输入层一样，即也是 1000000 时，那么输入层到隐含层的参数数据为 $1000000 \times 1000000 = 10^{12}$ ，这样就太多了，基本没法训练。所以图像处理要想练成神经网络大法，必先减少参数加快速度。就跟辟邪剑谱似的，普通人练得很挫，一旦自宫后内力变强剑法变快，就变的很牛了。

卷积神经网络 (CNN)

- 卷积神经网络有两种神器可以降低参数数目，第一种神器叫做局部感知。
- 在下方右图中，假如每个神经元只和 10×10 个像素值相连，那么权值数据为 1000000×100 个参数，减少为原来的万分之一。而那 10×10 个像素值对应的 10×10 个参数，其实就相当于卷积操作。

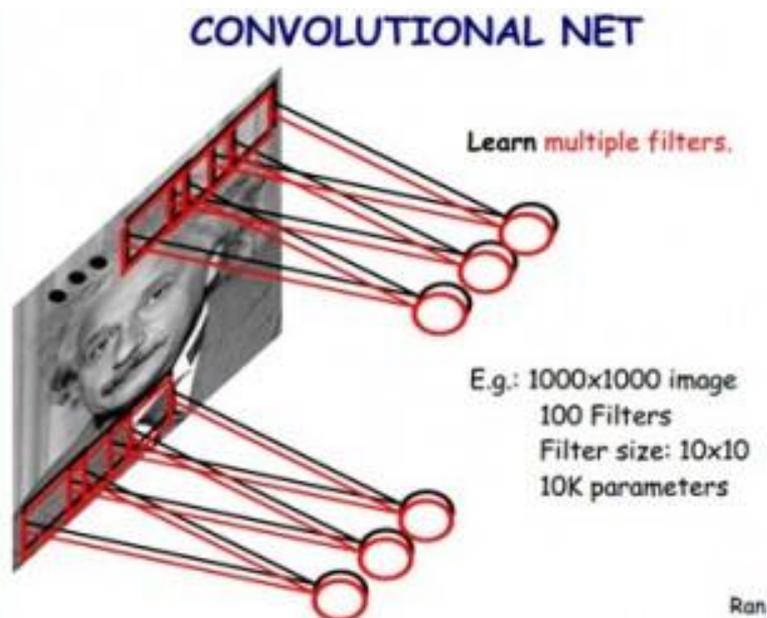
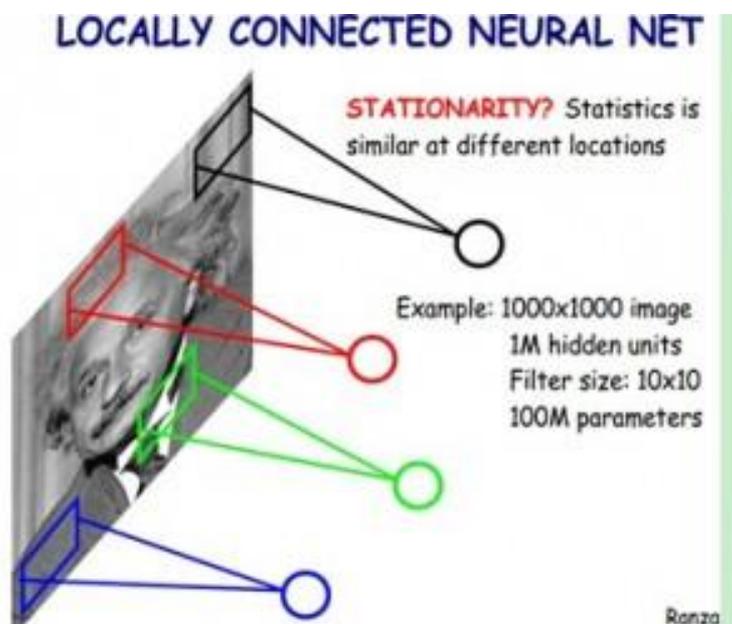


卷积神经网络 (CNN)

- 这样的话参数仍然过多，那么就启动第二级神器，即权值共享。在上面的局部连接中，每个神经元都对应100个参数，一共1000000个神经元，如果这1000000个神经元的100个参数都是相等的，那么参数数目就变为100了。
- 怎么理解权值共享呢？我们可以这100个参数（也就是卷积操作）看成是提取特征的方式，该方式与位置无关。这其中隐含的原理则是：图像的一部分的统计特性与其他部分是一样的。这也意味着我们在这一部分学习的特征也能用在另一部分上，所以对于这个图像上的所有位置，我们都能使用同样的学习特征。

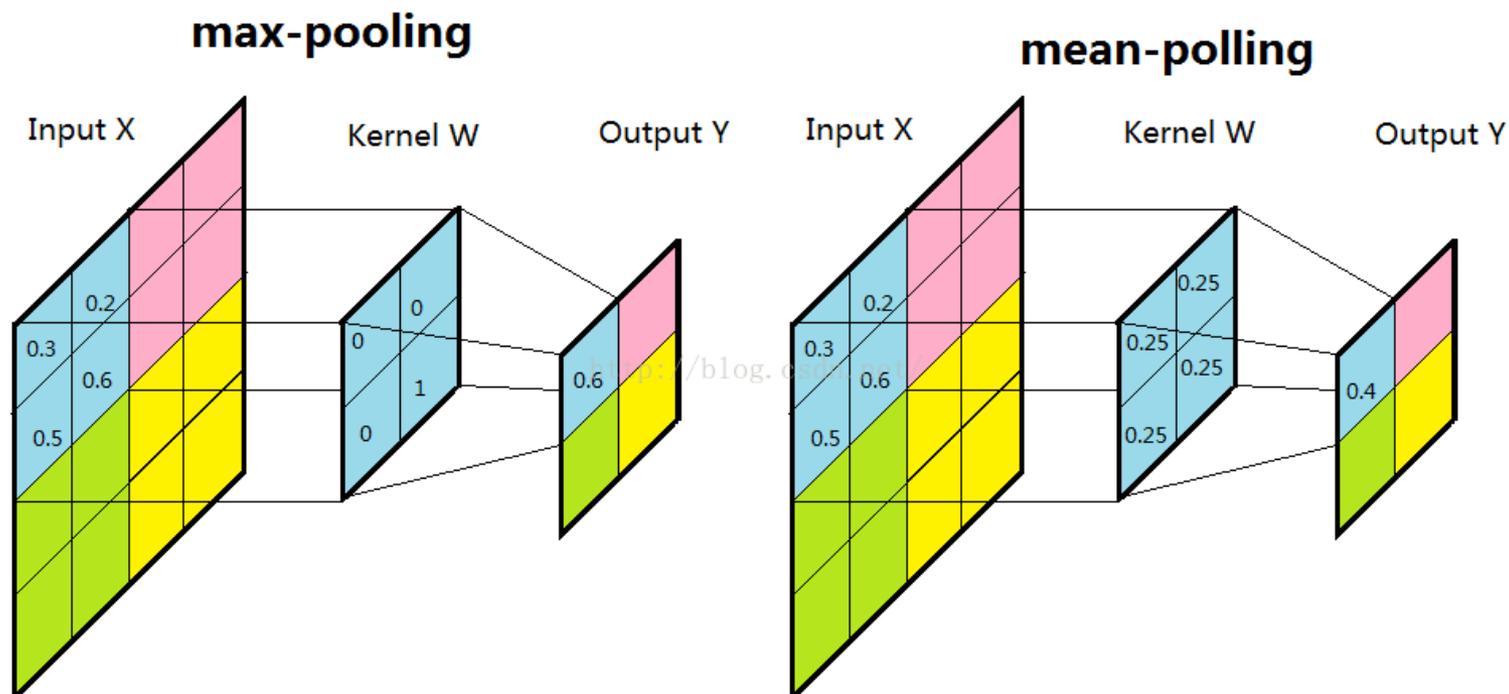
卷积神经网络 (CNN)

- 上面所述只有100个参数时，表明只有1个100*100的卷积核，显然，特征提取是不充分的，我们可以添加多个卷积核，比如32个卷积核，可以学习32种特征。在有多个卷积核时，如下图所示：



卷积神经网络 (CNN)

- 池化，也称作下采样，可以实现降维。常用有最大值池化和均值池化。

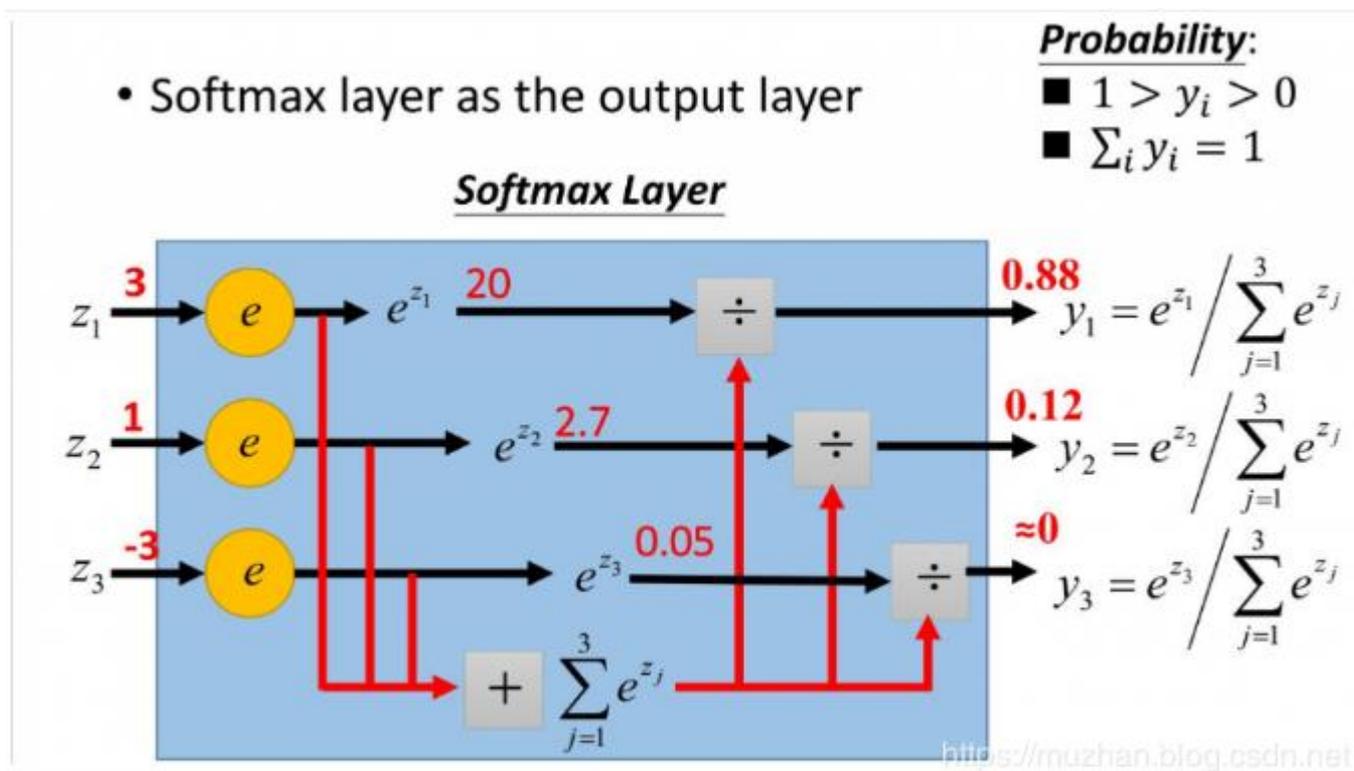


卷积神经网络 (CNN)

- 全连接层：连接所有的特征，将输出值送给分类器（如softmax分类器），最终得出识别结果。

softmax分类器

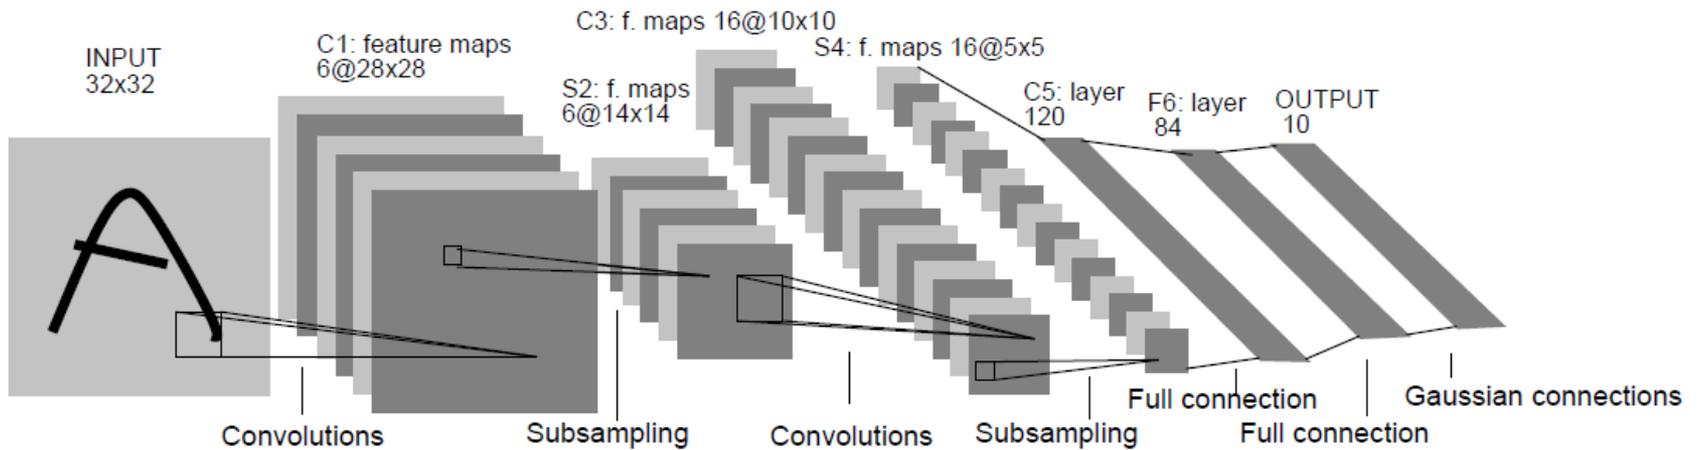
$$\text{softmax}(x_0) = \frac{e^{x_0}}{e^{x_0} + e^{x_1} + e^{x_2}}$$



softmax对神经元的输出信号进行加工，输出为分类的概率值。

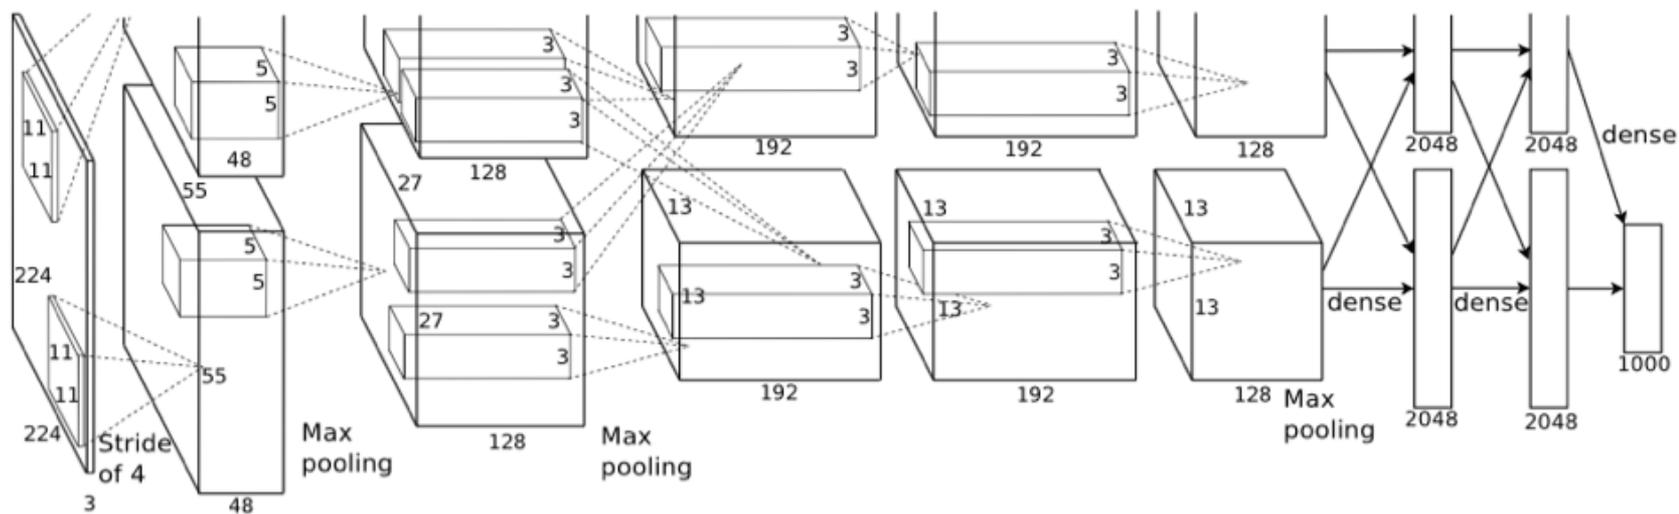
常见网络模型

- LeNet



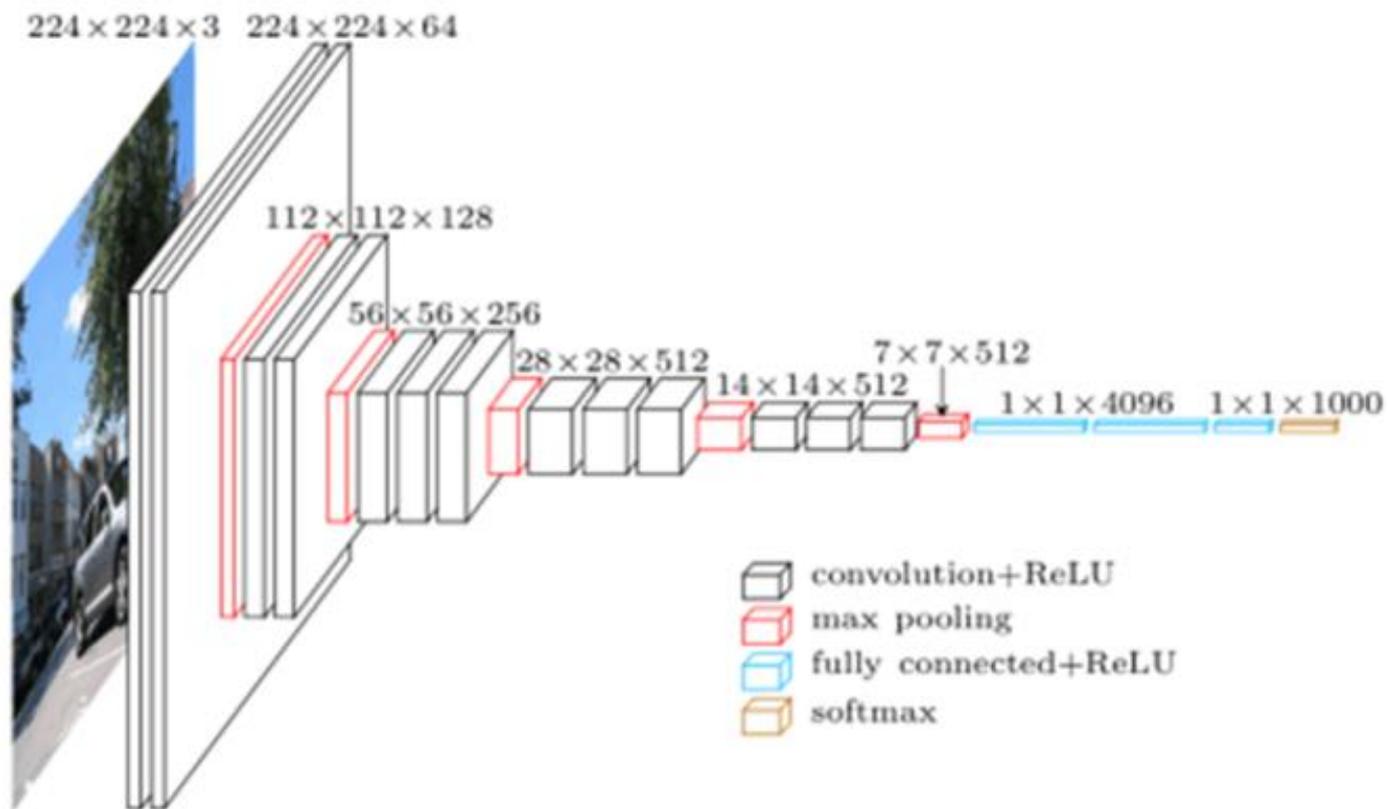
常见网络模型

- AlexNet



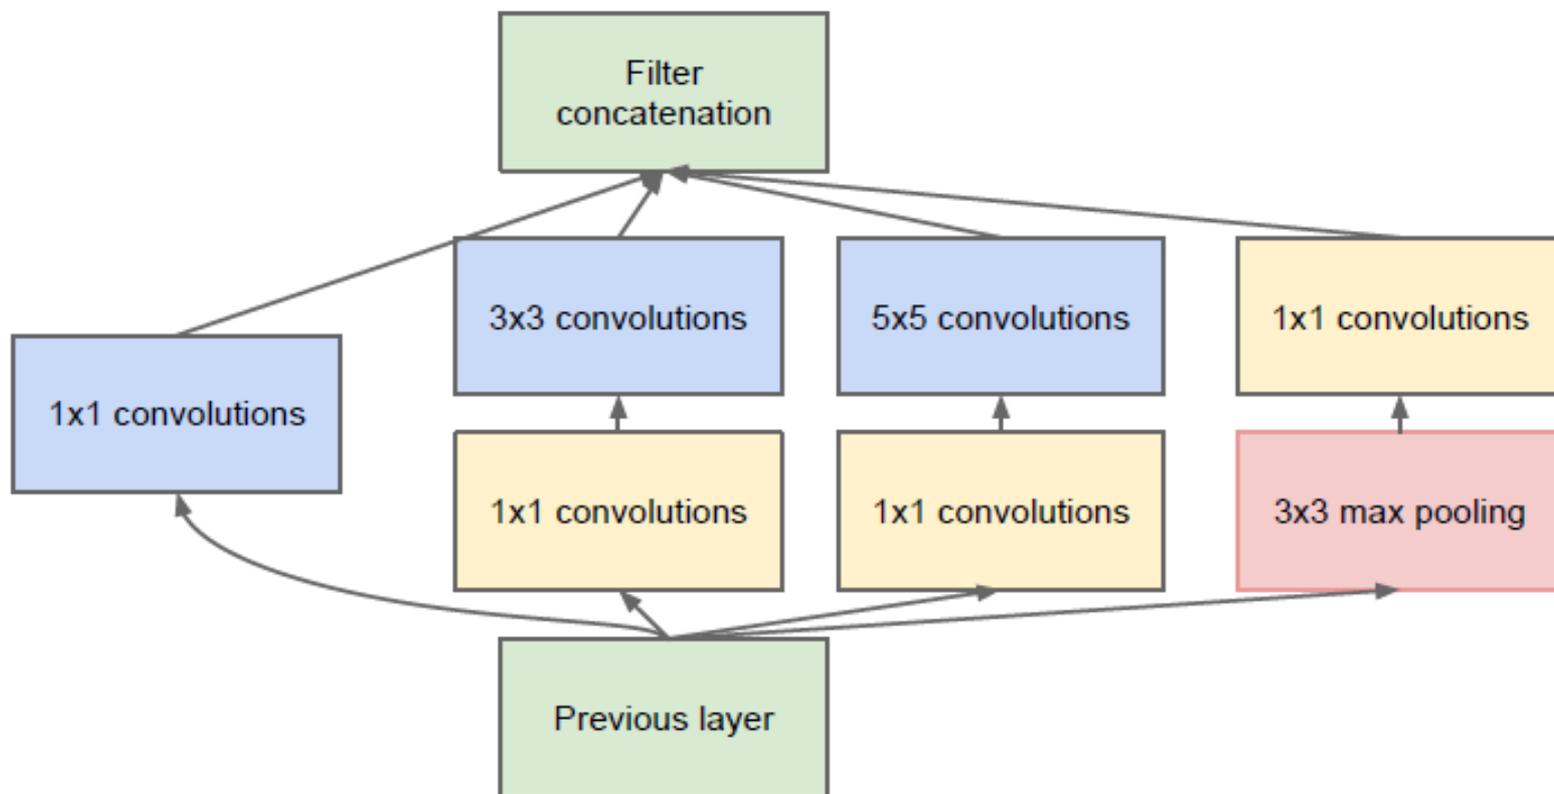
常见网络模型

- VGG16



常见网络模型

- GoogleNet (InceptionV4)



常见网络模型



CNN模型比较.doc

- 比较

模型名	AlexNet	VGG	GoogLeNet	ResNet
初入江湖	2012	2014	2014	2015
层数	8	19	22	152
Top-5错误	16.4%	7.3%	6.7%	3.57%
Data Augmentation	+	+	+	+
Inception(NIN)	-	-	+	-
卷积层数	5	16	21	151
卷积核大小	11,5,3	3	7,1,3,5	7,1,3,5
全连接层数	3	3	1	1
全连接层大小	4096,4096,1000	4096,4096,1000	1000	1000
Dropout	+	+	+	+
Local Response Normalization	+	-	+	-
Batch Normalization	-	-	-	+

其他深度学习算法

- 自动编码器 (AutoEncoder)
- 稀疏编码 (Sparse Coding)
- 限制玻尔兹曼机 (RBM)

深度学习的特点

深度学习常用算法介绍

深度学习常用框架介绍

开源框架概述

- 深度学习研究的热潮持续高涨，各种开源深度学习框架也层出不穷，其中包括TensorFlow、Caffe、Keras、CNTK、Torch7、MXNet、Leaf、Theano、DeepLearning4J、Lasagne、Neon等等。下图是各个开源框架在GitHub上的数据统计（2017年初）。

框 架	机 构	支持语言	Stars	Forks	Contributors
TensorFlow	Google	Python/C++/Go/...	41628	19339	568
Caffe	BVLC	C++/Python	14956	9282	221
Keras	fchollet	Python	10727	3575	322
CNTK	Microsoft	C++	9063	2144	100
MXNet	DMLC	Python/C++/R/...	7393	2745	241
Torch7	Facebook	Lua	6111	1784	113
Theano	U. Montreal	Python	5352	1868	271
Deeplearning4J	DeepLearning4J	Java/Scala	5053	1927	101
Leaf	AutumnAI	Rust	4562	216	14
Lasagne	Lasagne	Python	2749	761	55
Neon	NervanaSystems	Python	2633	573	52

开源框架概述

- Google、Microsoft、Facebook等巨头都参与了这场深度学习框架大战，此外，还有毕业于伯克利大学的贾扬清主导开发的Caffe，蒙特利尔大学Lisa Lab团队开发的Theano，以及其他个人或商业组织贡献的框架。下表是主流深度学习框架在各个维度的评分。

	模型设计	接口	部署	性能	架构设计	总体评分
TensorFlow	80	80	90	90	100	88
Caffe	60	60	90	80	70	72

续表

	模型设计	接口	部署	性能	架构设计	总体评分
CNTK	50	50	70	100	60	66
Theano	80	70	40	50	50	58
Torch	90	70	60	70	90	76
MXNet	70	100	80	80	90	84
DeepLearning4J	60	70	80	80	70	72

TensorFlow

- TensorFlow最初是由研究人员和Google Brain团队针对机器学习和深度神经网络进行研究所开发的，目前开源之后可以在几乎各种领域适用。
- TensorFlow灵活的架构可以部署在一个或多个CPU、GPU的台式以及服务器中，或者使用单一的API应用在移动设备中。



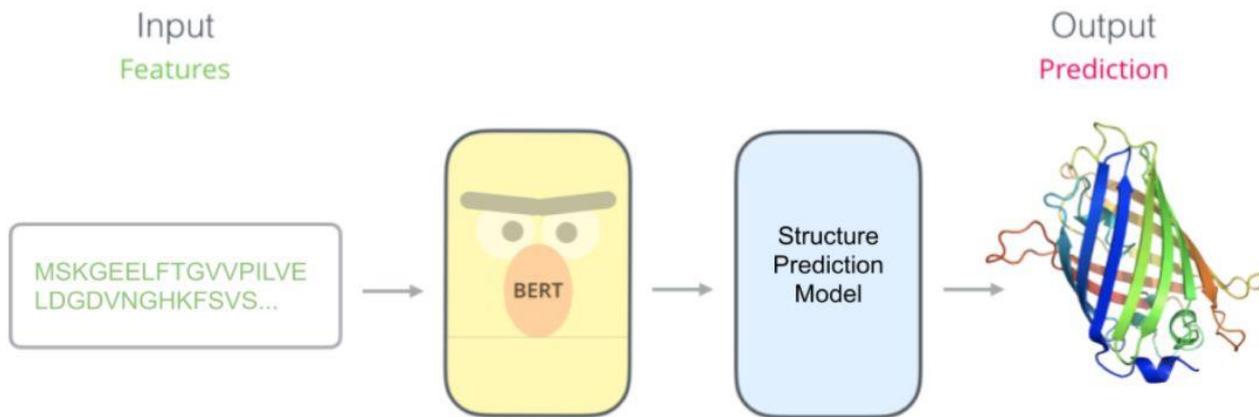
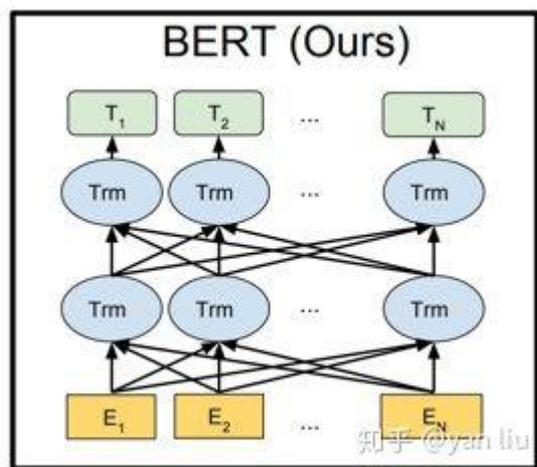
Torch

- Torch是一个有大量机器学习算法支持的科学计算框架，其诞生已经有十年之久，但是真正起势得益于Facebook开源了大量Torch的深度学习模块和扩展。Torch另外一个特殊之处是采用了编程语言Lua(该语言曾被用来开发视频游戏)。
- PyTorch是基于Torch的衍生，支持Python语言，实现了机器学习框架 Torch 在 Python 语言环境的执行。



Bert

- 使用Transformer的结构将已经走向瓶颈期的Word2Vec带向了一个新的方向，并再一次炒火了《Attention is All you Need》这篇论文；
- 11个NLP任务的精度大幅提升足以震惊整个深度学习领域；
- 无私的开源了多种语言的源码和模型，具有非常高的商业价值；
- 迁移学习又一次胜利，而且这次是在NLP领域的大胜，狂胜。



Caffe

- Caffe由加州大学伯克利的PHD贾扬清开发，全称Convolutional Architecture for Fast Feature Embedding，是一个清晰而高效的开源深度学习框架，目前由伯克利视觉学中心（Berkeley Vision and Learning Center, BVLC）进行维护。（贾扬清曾就职于MSRA、NEC、Google Brain，他也是TensorFlow的作者之一，目前任职于Facebook FAIR实验室。）
- Caffe2脸书 (Facebook) 出品，为生产环境设计，提供在各种平台（包括移动设备）的运行。



Theano

- 2008年诞生于蒙特利尔理工学院,Theano派生出了大量深度学习Python软件包,最著名的包括Blocks和Keras。Theano的核心是一个数学表达式的编译器,它知道如何获取你的结构。并使之成为一个使用numpy、高效本地库的高效代码,如BLAS和本地代码(C++)在CPU或GPU上尽可能快地运行。它是为深度学习中处理大型神经网络算法所需的计算而专门设计的,是这类库的首创之一(发展始于2007年),被认为是深度学习研究和开发的行业标准。

Theano logo, featuring the word "theano" in white lowercase letters on a blue rectangular background.

theano

Deeplearning4j

- Deeplearning4j是“for Java”的深度学习框架，也是首个商用级别的深度学习开源库。Deeplearning4j由创业公司Skymind于2014年6月发布，使用Deeplearning4j的不乏埃森哲、雪弗兰、博斯咨询和IBM等明星企业。DeepLearning4j是一个面向生产环境和商业应用的高成熟度深度学习开源库，可与Hadoop和Spark集成，即插即用，方便开发者在APP中快速集成深度学习功能。

MXNet

- 出自CXXNet、Minerva、Purine 等项目的开发者之手，主要用C++ 编写。MXNet 强调提高内存使用的效率，甚至能在智能手机上运行诸如图像识别等任务。

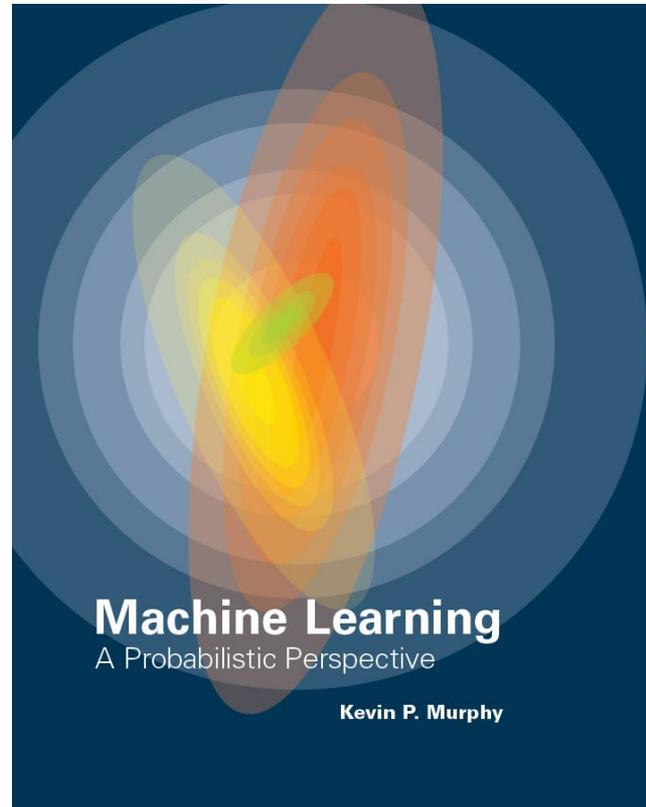
mxnet

CNTK

- CNTK (Computational Network Toolkit) 是微软研究院 (MSR) 开源的深度学习框架。它最早由start the deep learning craze的演讲人创建，目前已经发展成一个通用的、跨平台的深度学习系统，在语音识别领域的使用尤其广泛。

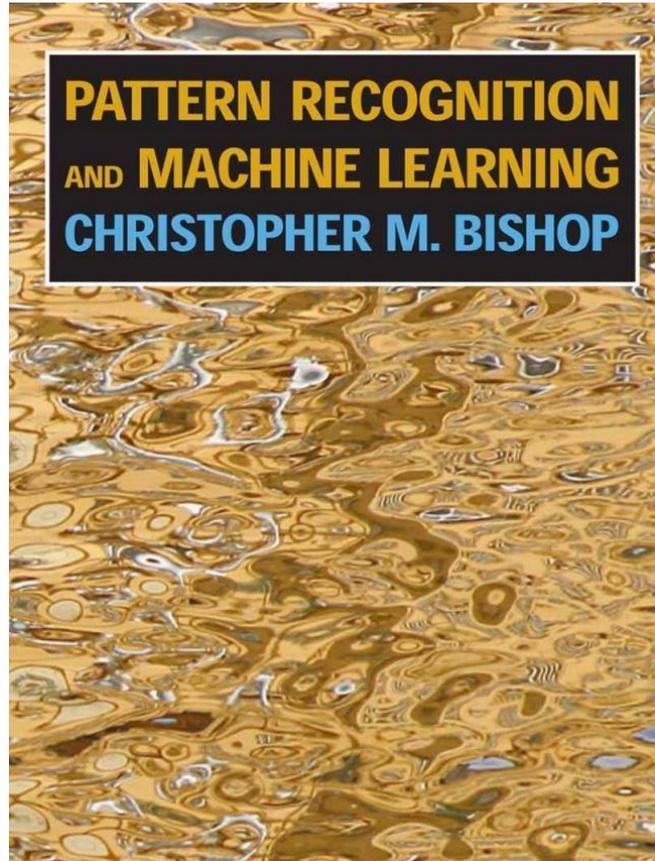


References

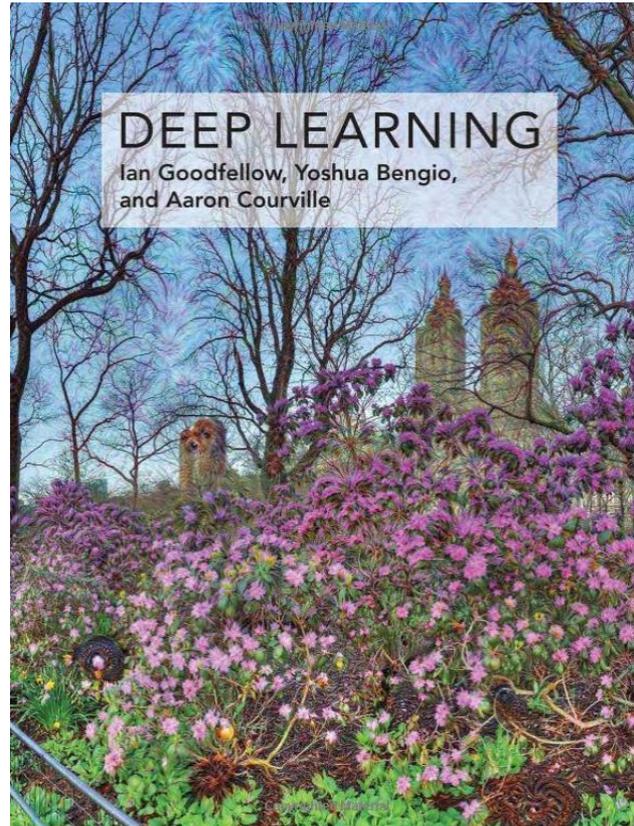


<https://github.com/probml/pyprobml>

References



References



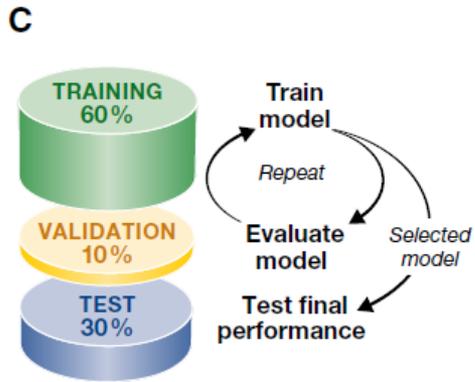
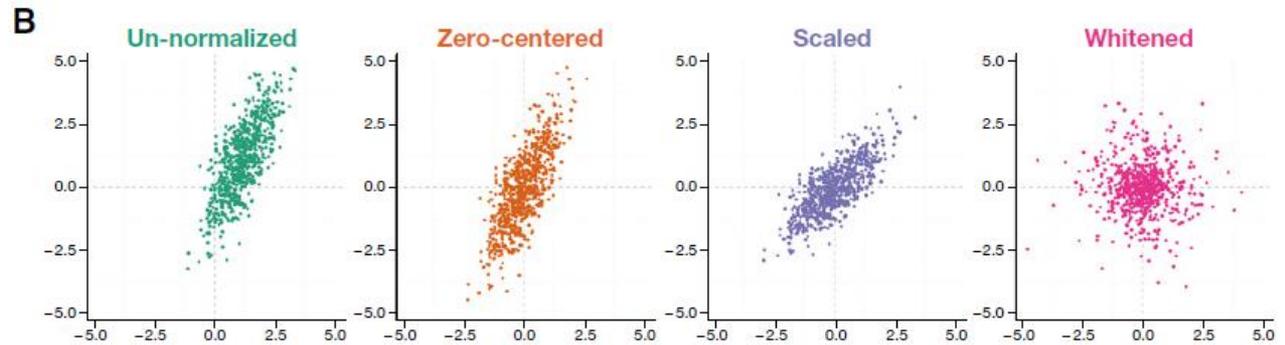
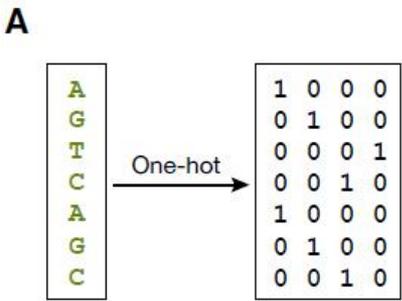
<https://github.com/exacity/deeplearningbook-chinese>

Part III

生物大数据的深度学习的方法

Feature Selection

Collection→Partition→Normalization
Data label

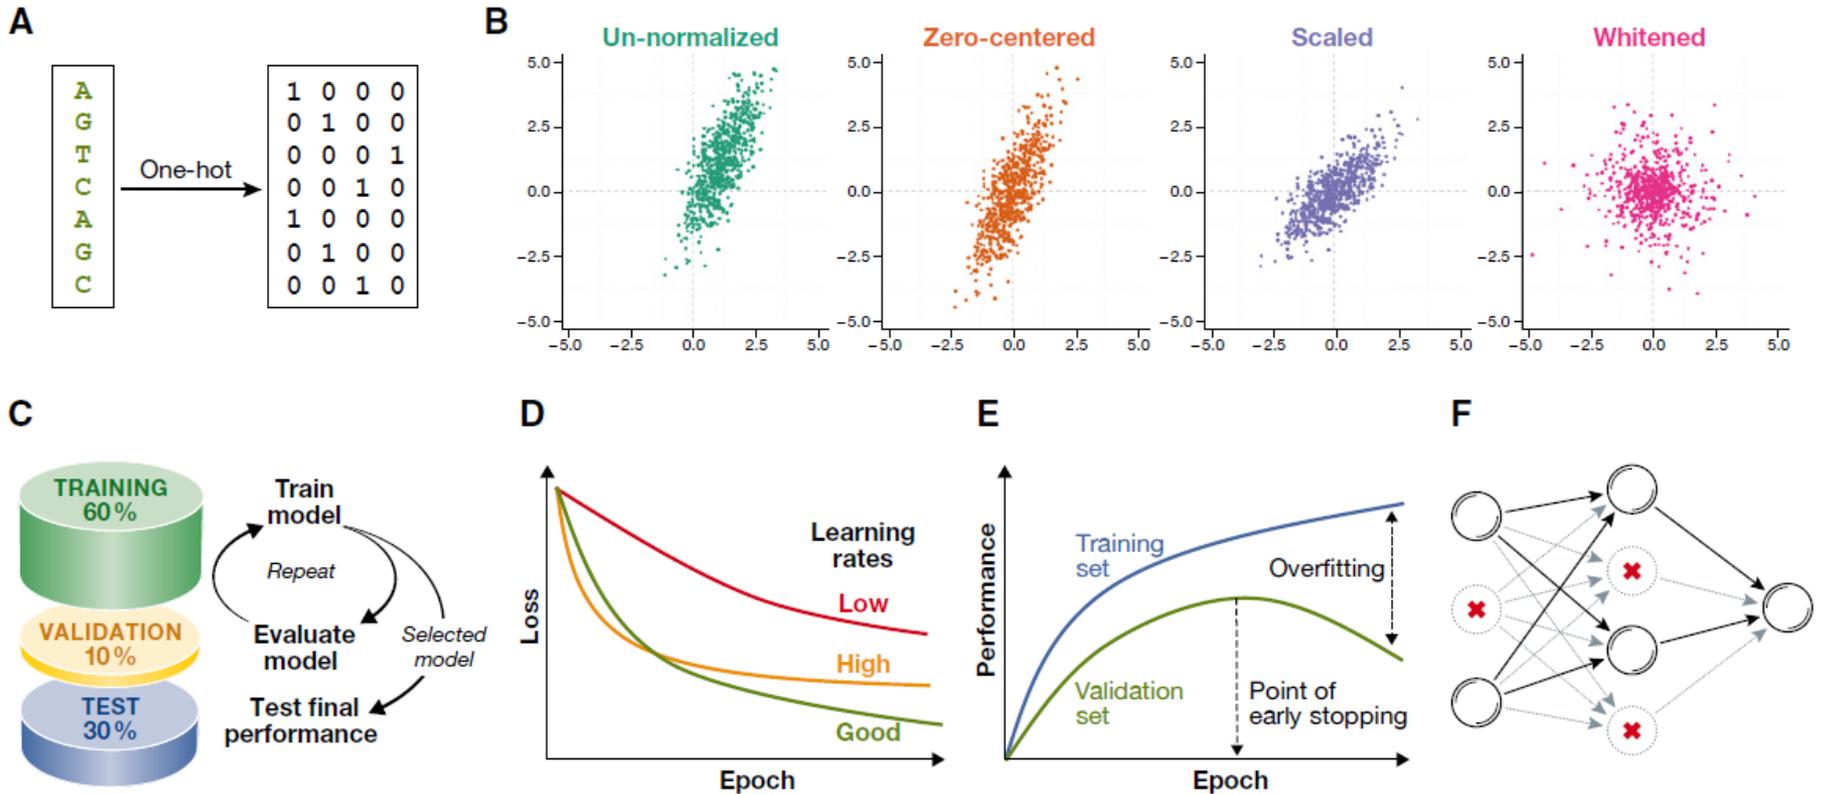


(A) Data One-hot Coding

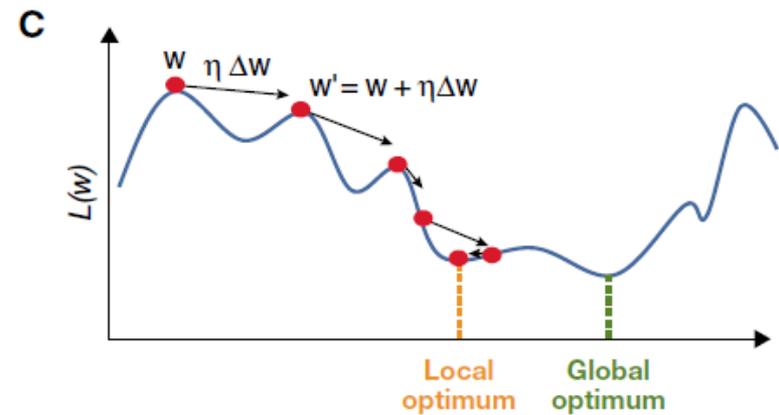
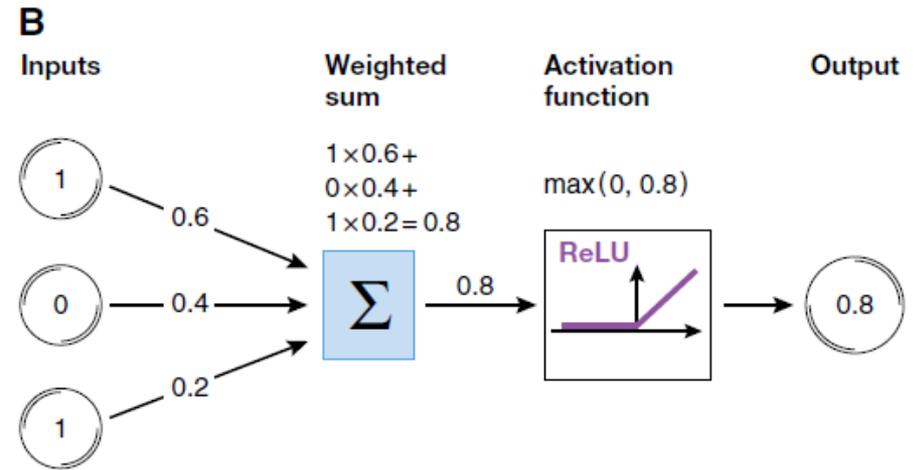
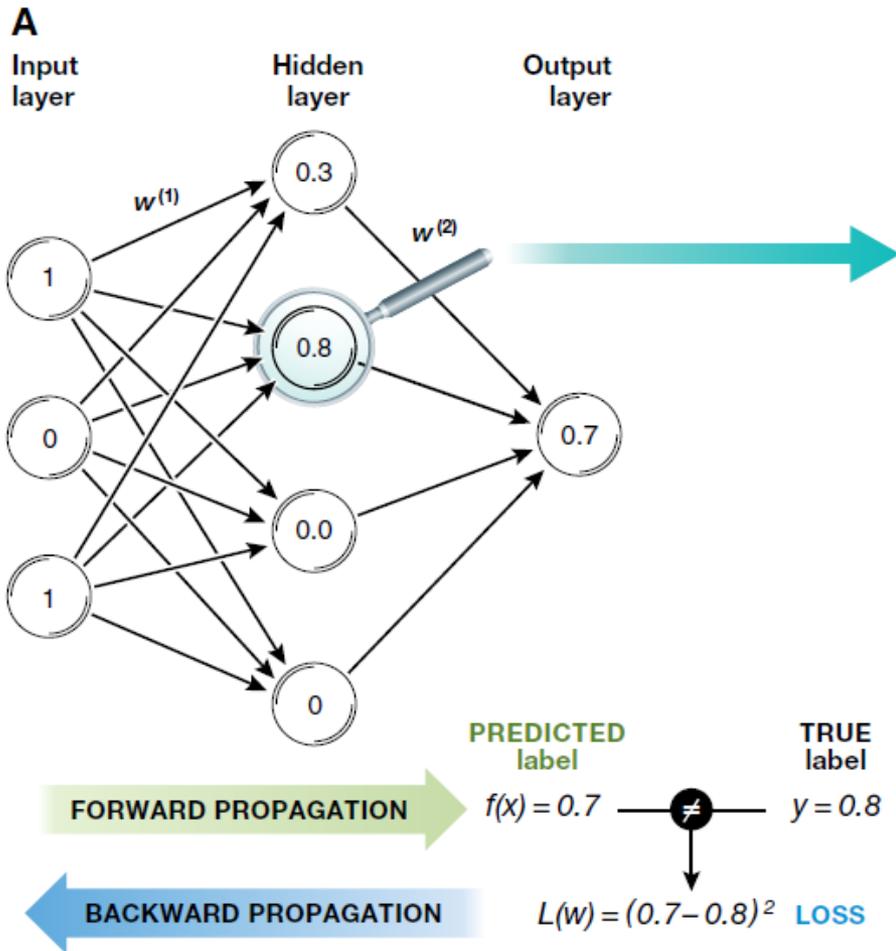
(B) Different data normalization methods

(C) Training data, validation data(adjust model structure), Test data

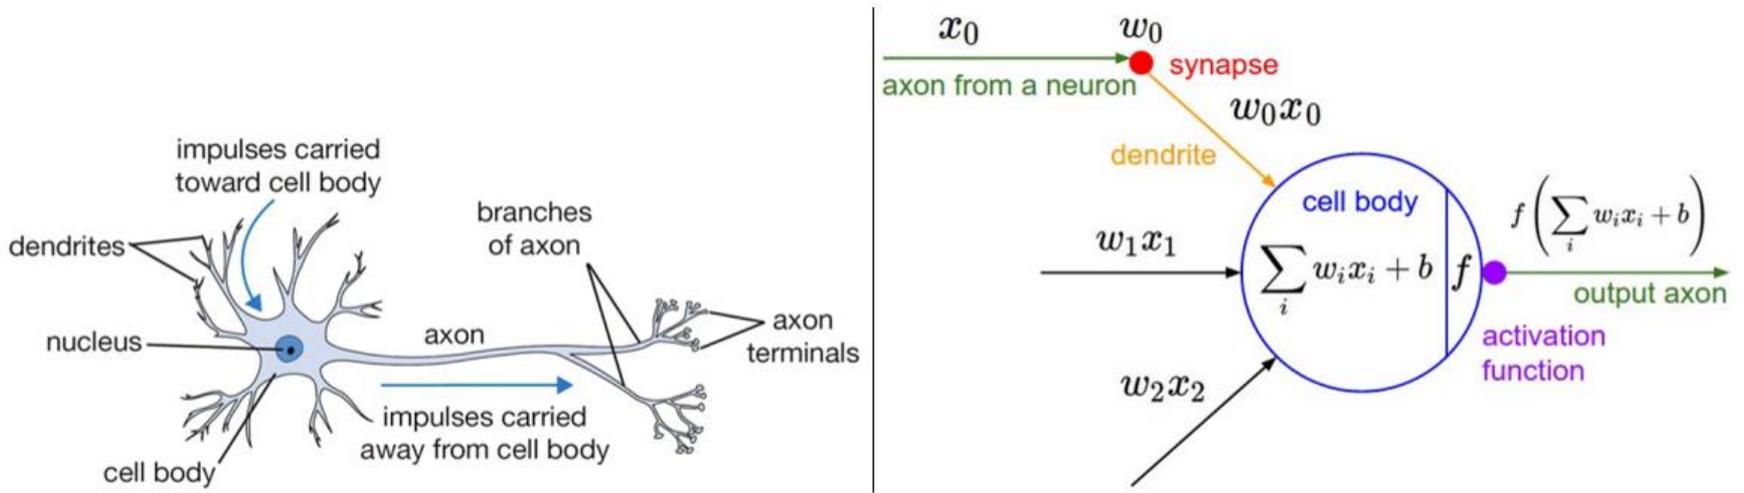
Data normalization



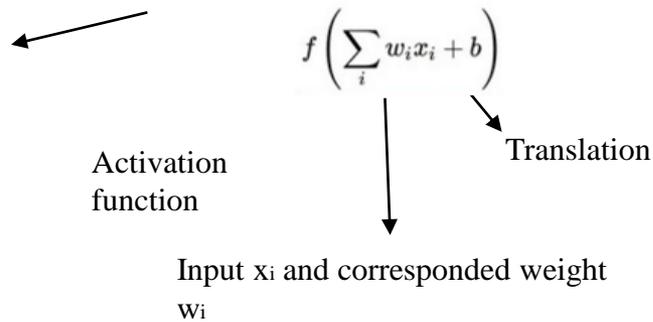
Artificial Neural Network



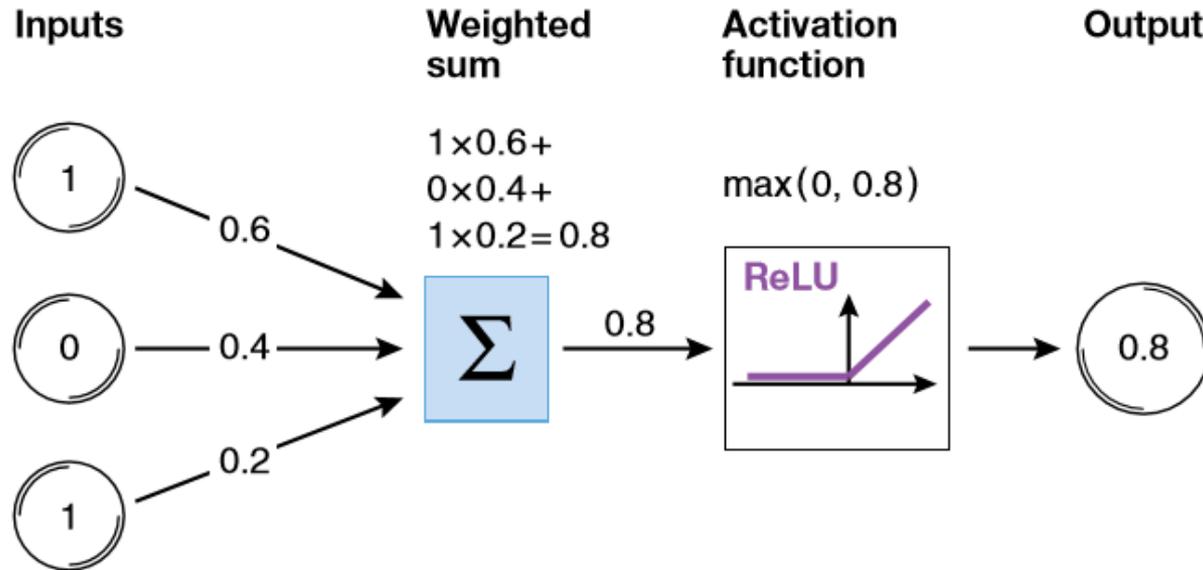
Model Construction



A Cartoon drawing of a biological neuron(left) and its mathematical model(right)

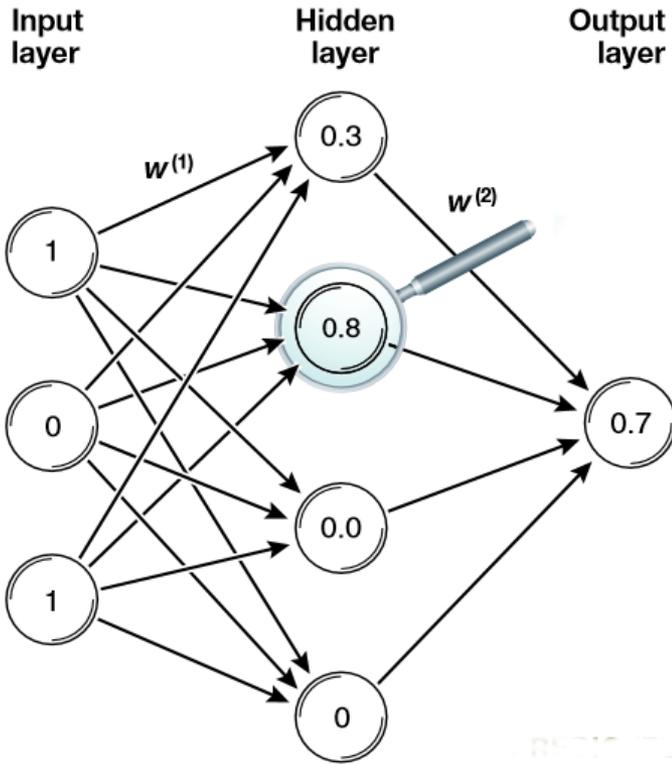


Model Construction



An example for a “transmission of neural signal”

Model Construction



Model construction:

Input label: 1, 0, 1

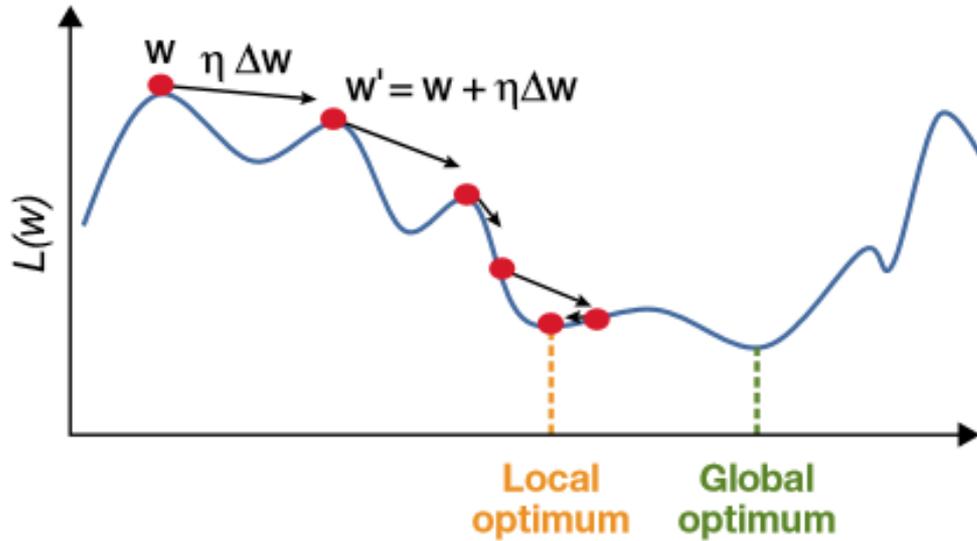
Predicted label: 0.7

True data label: 0.8

Neuron network construction
process

Pick the best model

loss function: $L(w) = (\text{predicted label} - \text{true label})^2$



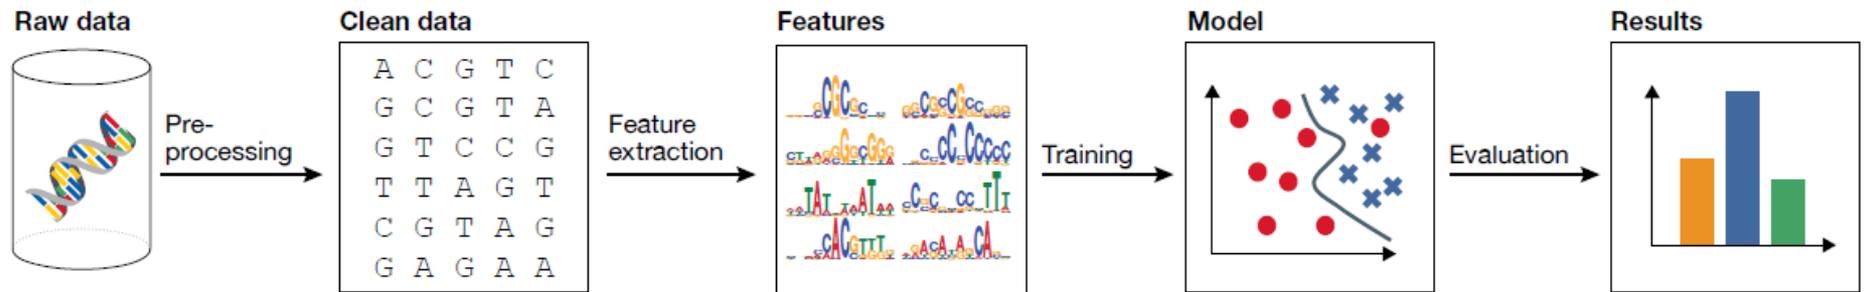
Many parameters are adjusted to find the global minimums
then update the neuron weight



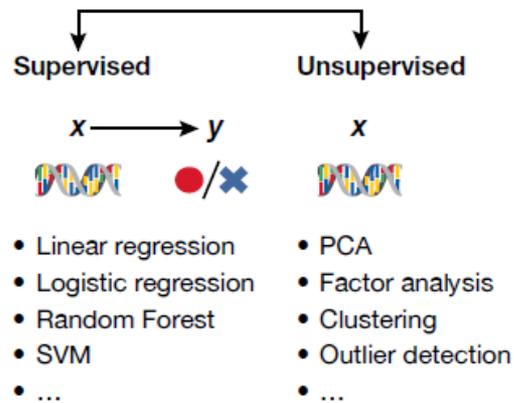
Predicted label:0.8 → True label:0.8

Machine learning and representation learning

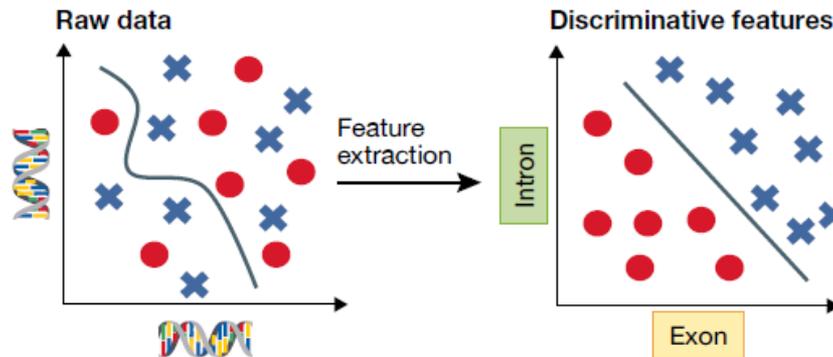
A



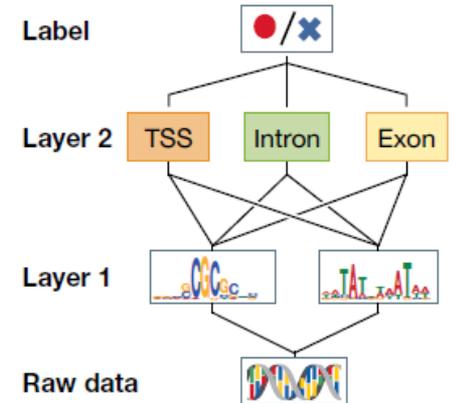
B



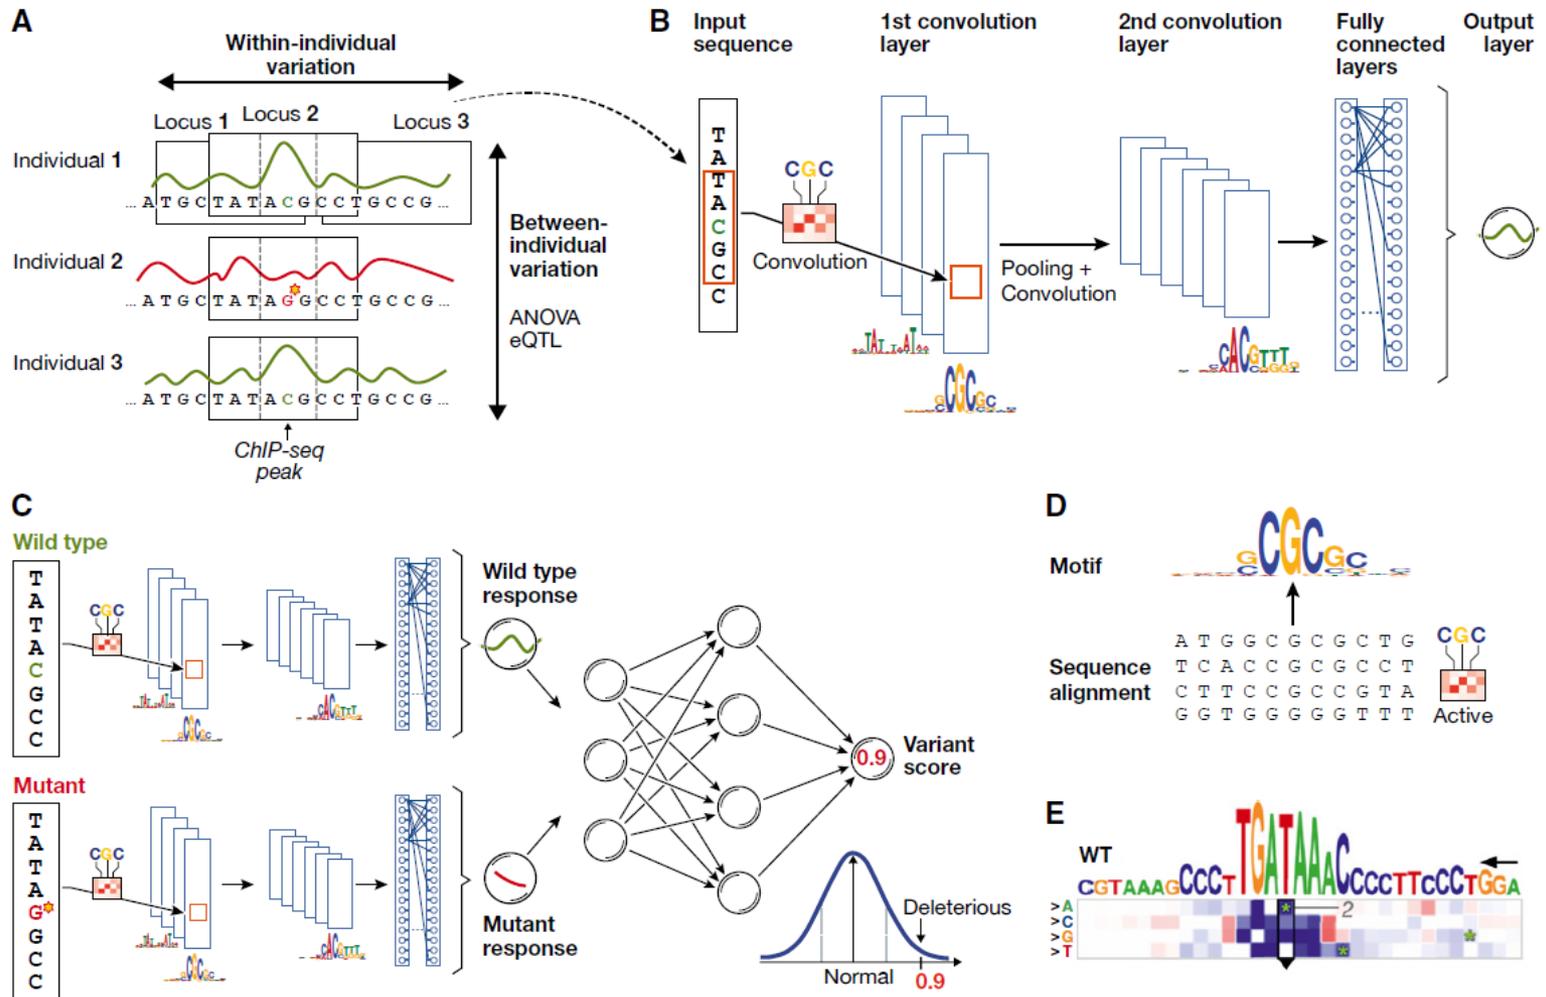
C



D



Neural Network for DNA sequences



Part IV

生物大数据深度学习的应用案例

INPUT

PROCESSING

APPLICATION

DNA sequences

RNA

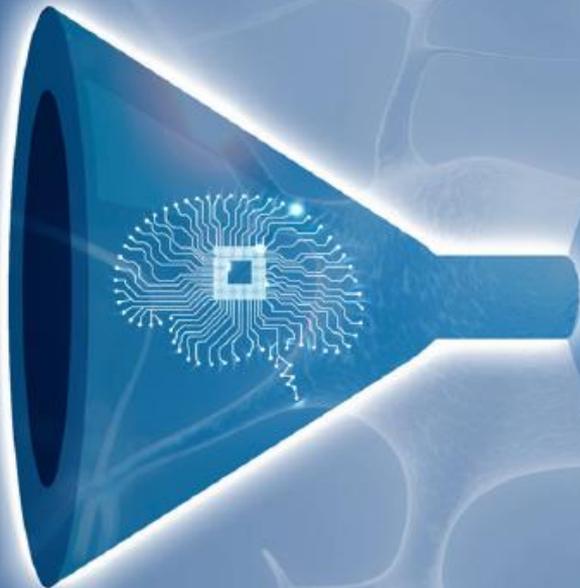
MicroRNA

Gene expressions

Gene alleles

Molecule compounds

Protein structures



Disease stratification



Cancer diagnostic



Gene variations



Drug design

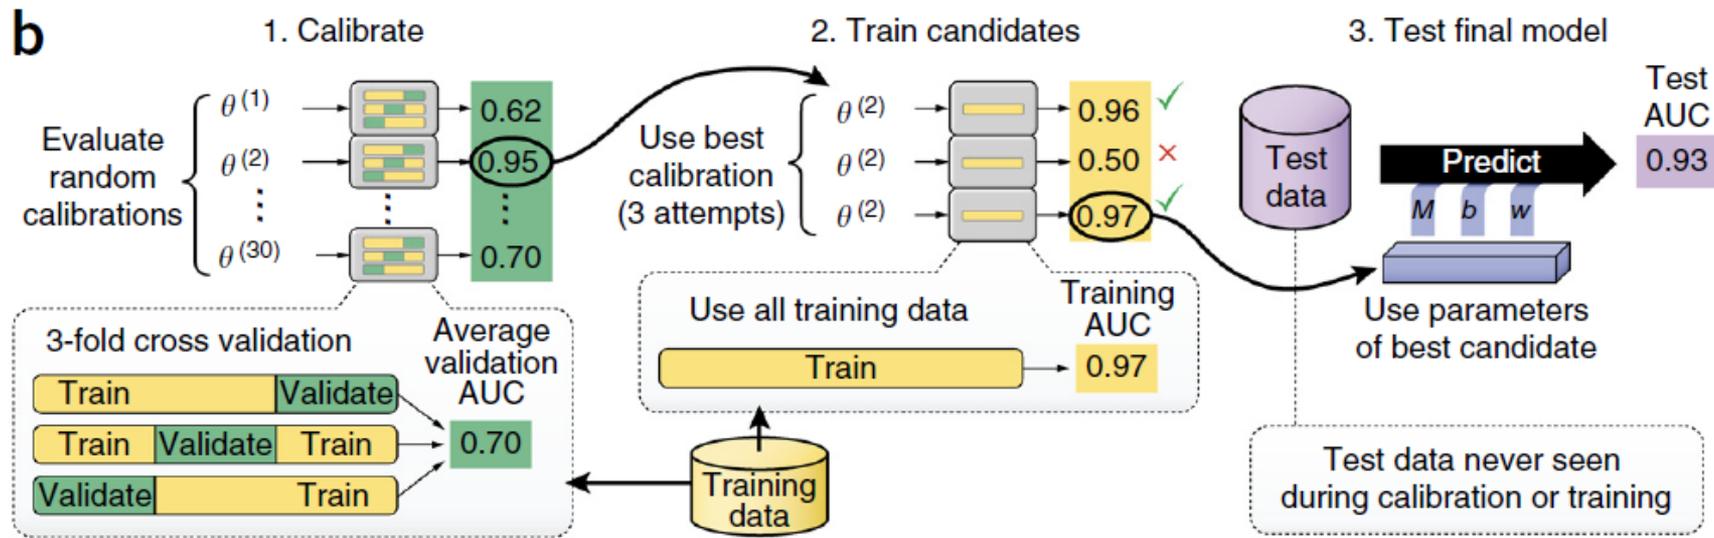
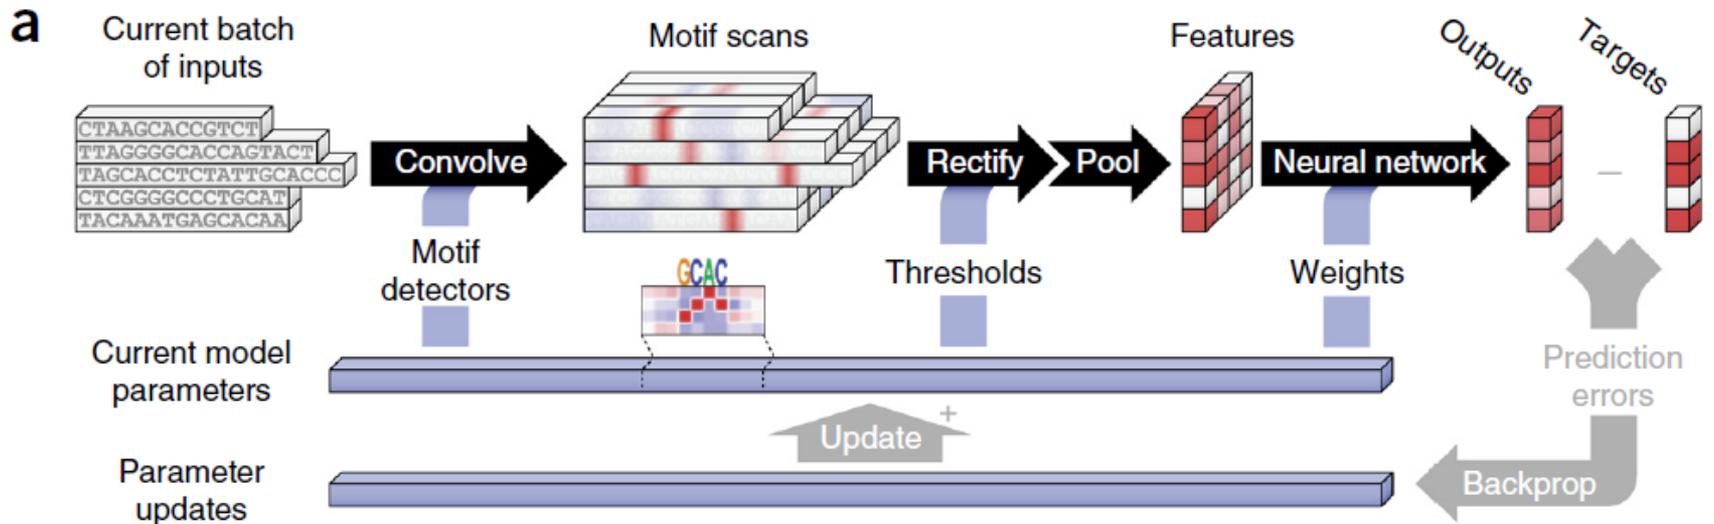


RNA Splicing



Protein structure and interaction

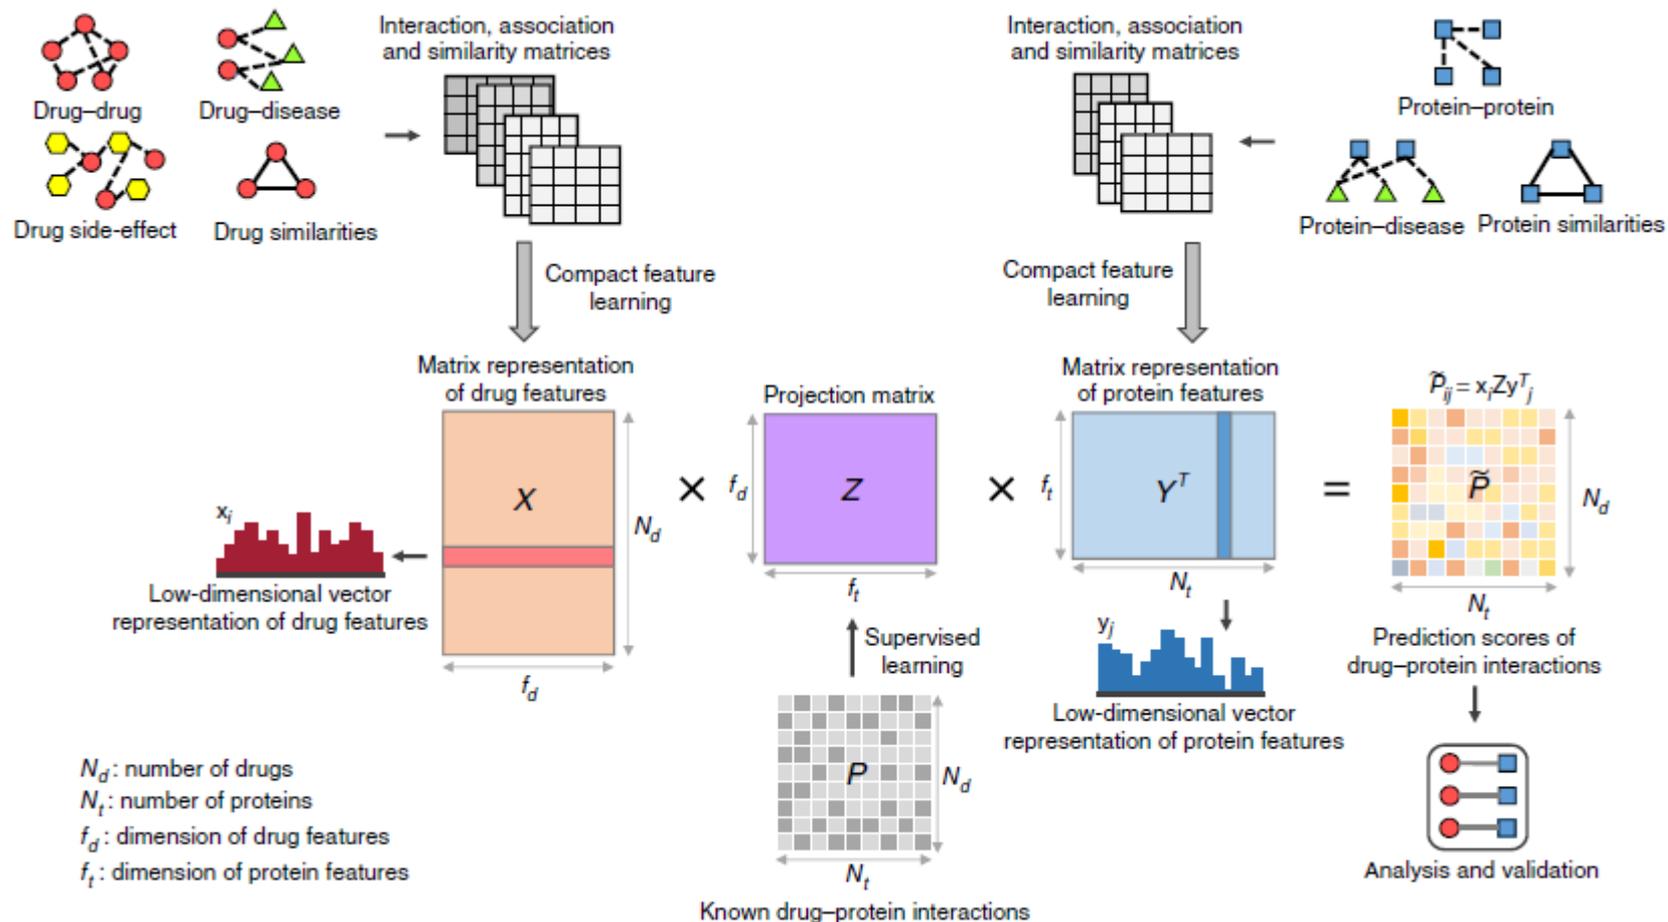
DNA- and RNA-binding protein prediction



DeepBind总结

- CNN算法自动学习了模版M的参数，这个模版就是能够决定蛋白质是否Binding的特征！
- 而CNN网络就是在前端自动总结出可能的模版，在后端根据这些模版在序列中出现与否而进行判断；
- CNN算法巧妙地将统计归纳和MLP的学习功能相结合，减少了MLP中要学习的参数；

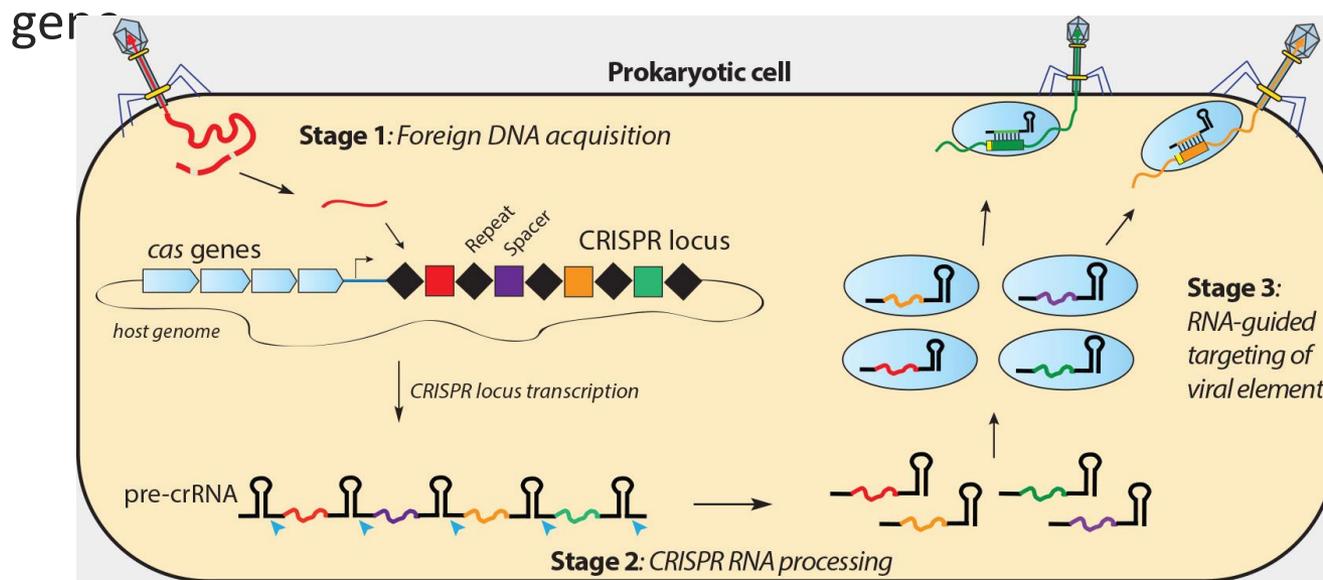
Drug-target interaction prediction



What is Cas?

- Cas (CRISPR-associated system) genes

bacteria immune system: cleavage the foreign



CRISPR-Cas system workflow: foreign DNA acquisition → CRISPR RNA processing → RNA-guided targeting of viral element

The existing method for Cas genes prediction

- The existing method for Cas genes prediction

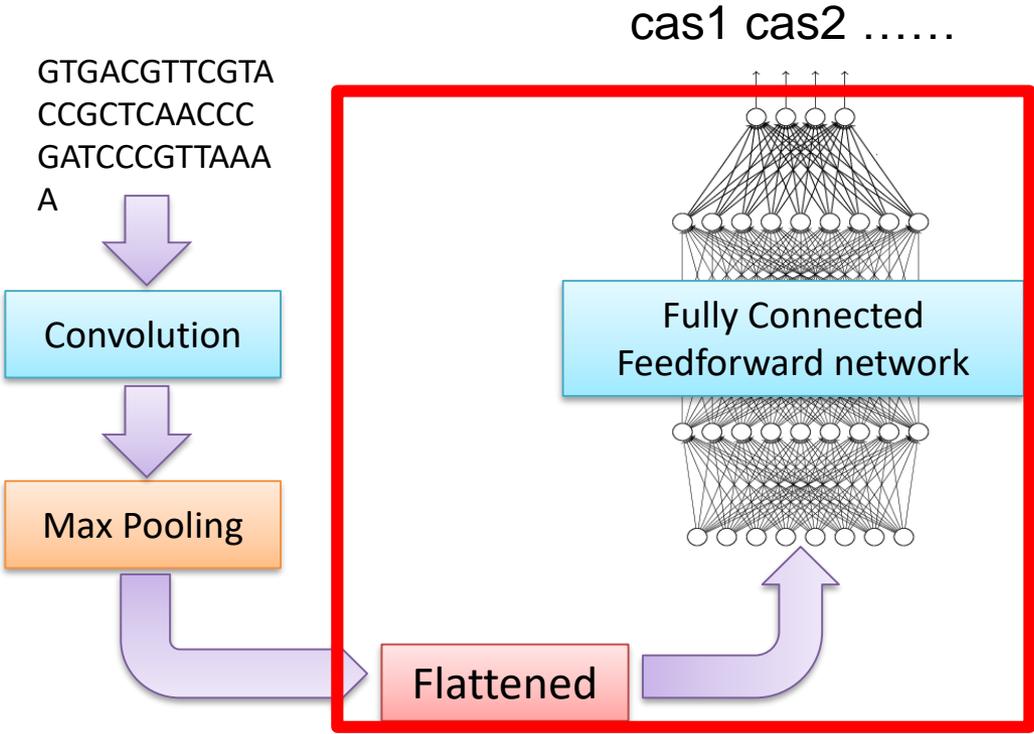
Existing tools	Description
HMMCAS	A web tool for Cas protein identification and domain annotation
CRISPRone	A website tool to predict CRISPR arrays of repeat-spacer units, and cas genes
BLAST	A website tool to find regions of local similarity between sequences
CD-search	A tool allows the conserved domain annotation for large sets of protein queries
Custom HMMs	HMMs deposited in the TIGRFAMs and Pfam protein family databases for known protein families

- Limitations

- Scan for known Cas proteins.
- Identify or predict based on the “best hits” of homology searches against existing databases.
- Work well for detecting known and highly conserved types of Cas genes but may fail to detect novel Cas genes.

The whole neural network to predict Cas genes

Mainly based on Python Tensorflow package

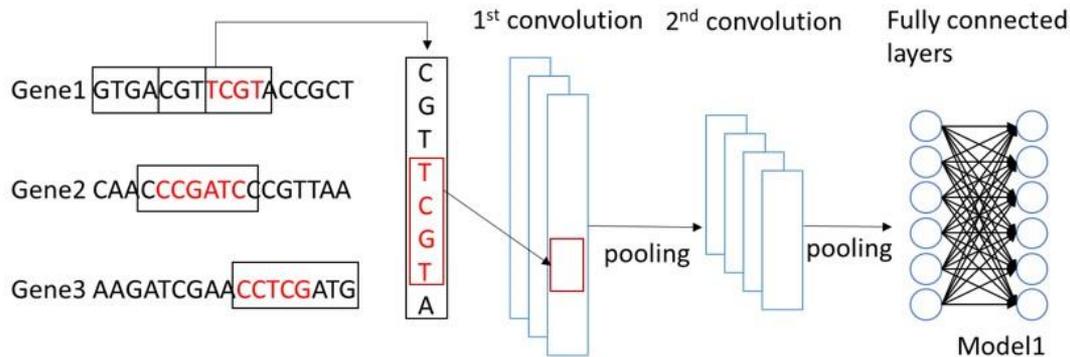


Data preparation

10 core family, 75 family, 7500 Cas genes in total

Core family	Structural features
Cas1	N-terminal β stranded domain and catalytic C-terminal α -helical domain
Cas2	RuvC-like nuclease domains
Cas3	Helicase and HD domain
Cas4	RecB-like nuclease homolog with three-cysteine C-terminal cluster
Cas5	RRM (ferredoxin) fold, RAMP superfamily
Cas6	Double RRM (ferredoxin) fold, RAMP superfamily
Cas7	RRM (ferredoxin) fold with subdomains, RAMP superfamily
Cas8	Subunit of Cascade complex, involved in PAM recognition
Cas9	RuvC-like (RNase H fold) and HNH nuclease domains
Cas10	Two domains homologous to Palm domain polymerases and cyclases

Convolution Neural Networks model



- Design a neural network that will classify cas proteins.
- Given a cas gene sequence, we want our neural net to let us know which cas gene that it belong to.
- The sequence is be made up of A, T, C,G, and fasta format:
- the neural net will have **xxx inputs**, each one representing a particular sequence and a hidden layer consisting of a number of neurons (more on this later) all feeding their output into just **one neuron in the output** layer.

A convolutional layer

A CNN is a neural network with some convolutional layers (and some other layers). A convolutional layer has a number of filters that does convolutional operation.

1	0	0	0
0	1	0	0
0	0	0	1
1	0	0	0
⋮	⋮	⋮	⋮
⋮	⋮	⋮	⋮
0	0	1	0

80 x 4

These are the network parameters to be learned

1	-1	-1	-1
-1	1	-1	-1
-1	-1	-1	1
1	-1	-1	-1

Filter 1

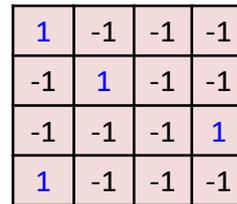
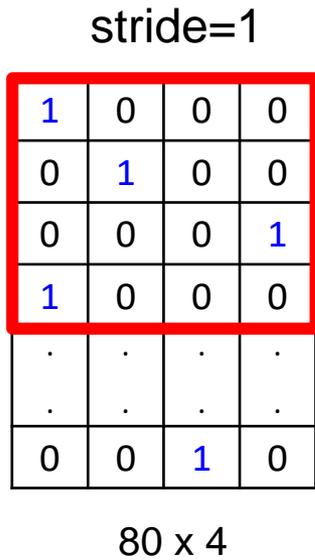
-1	1	-1	-1
-1	-1	1	-1
-1	-1	-1	1
1	-1	-1	-1

Filter 2

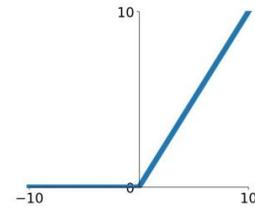
.....

Each filter detects a small pattern (4 x 4).

A convolutional layer

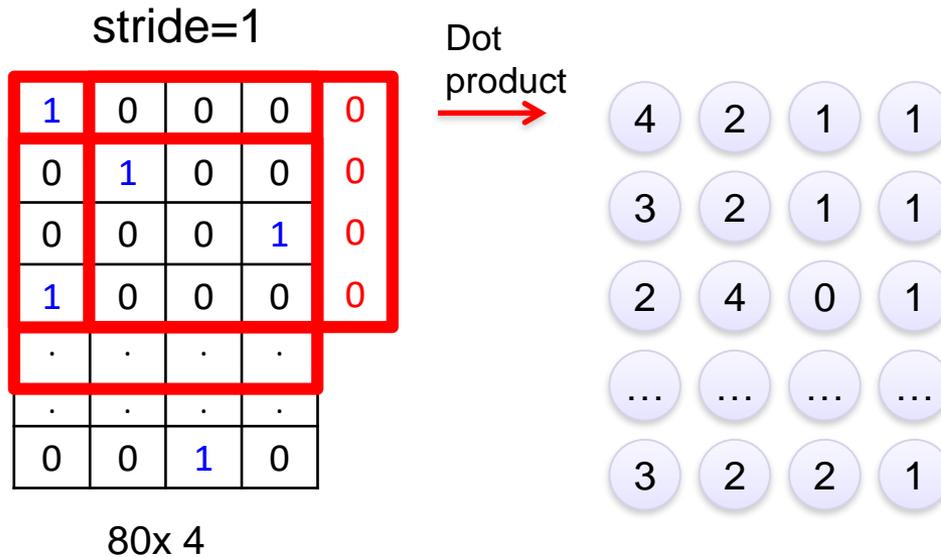


Filter 1



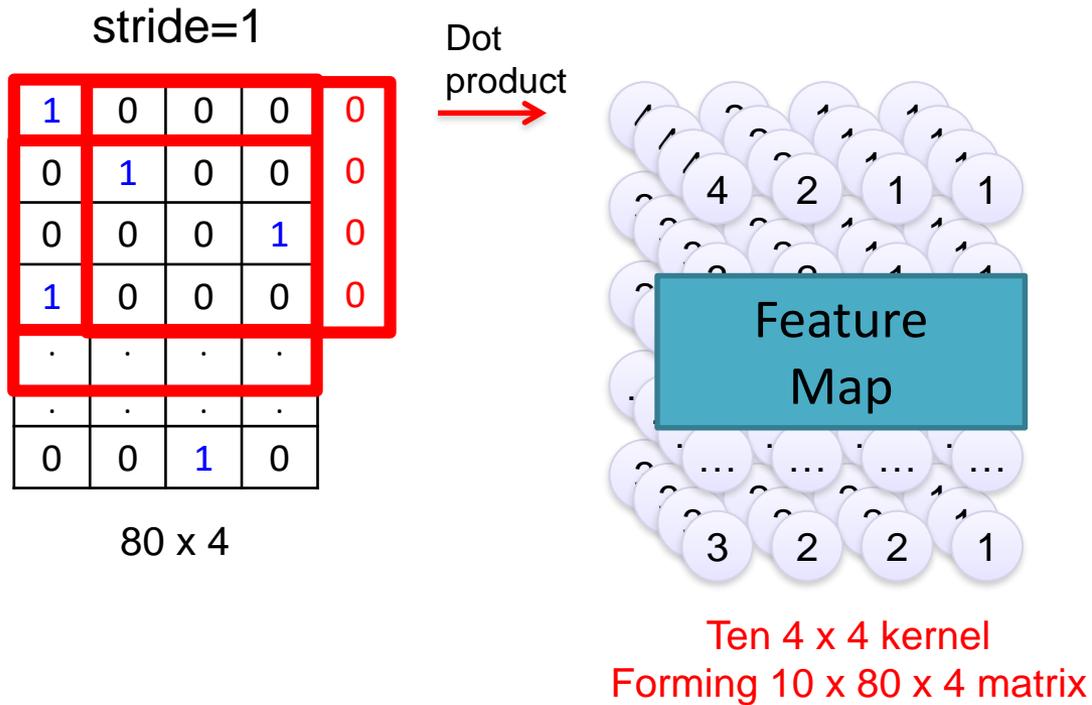
$$\text{ReLU}(x) = \max(0, x)$$

A convolutional layer



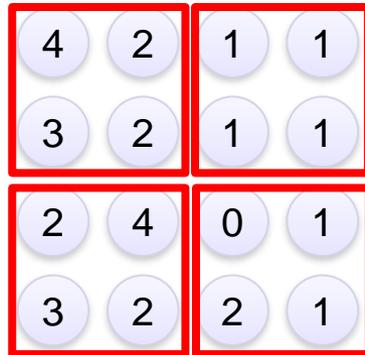
Repeat this for each filter

A convolutional layer



A convolutional layer

$$p_{nfi} = \max_{|k| < P/2} (a_{nf, i+k})$$

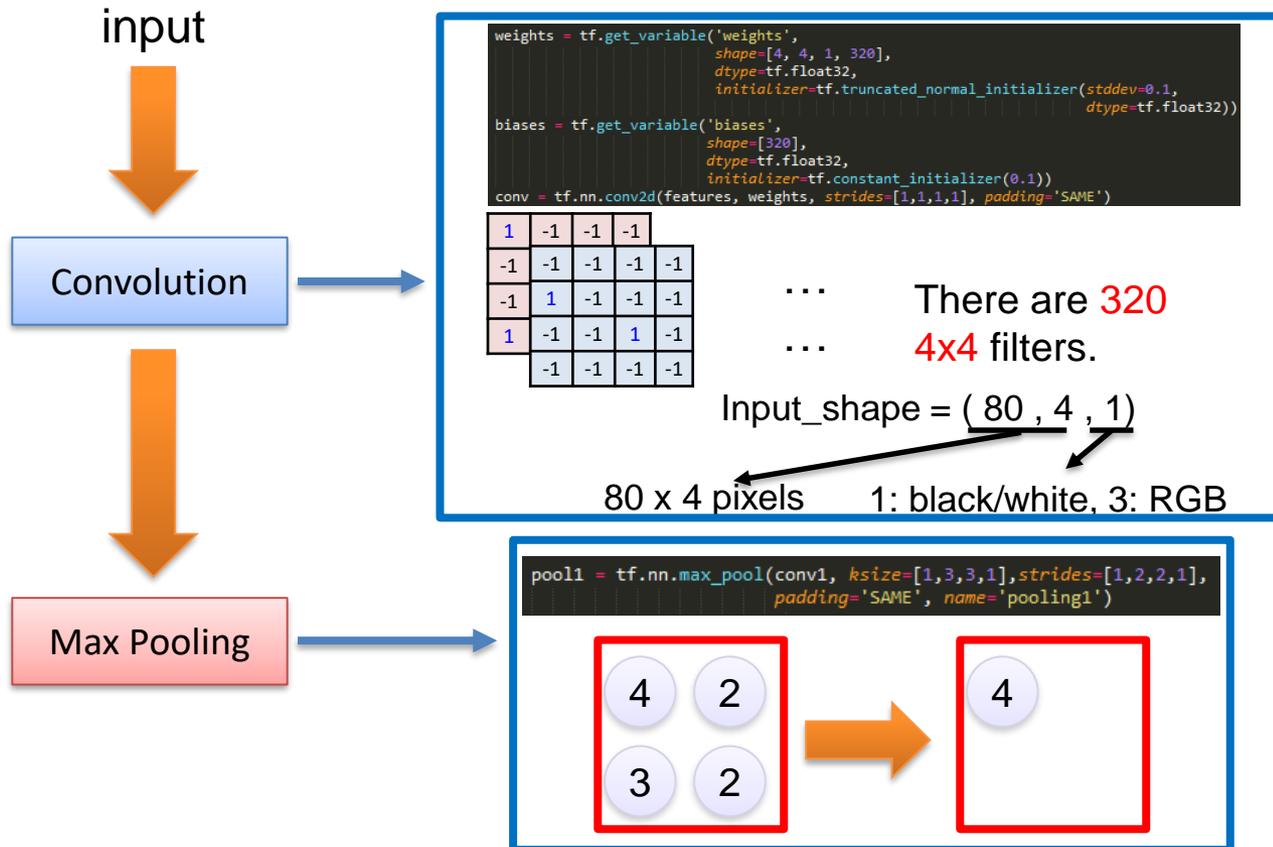


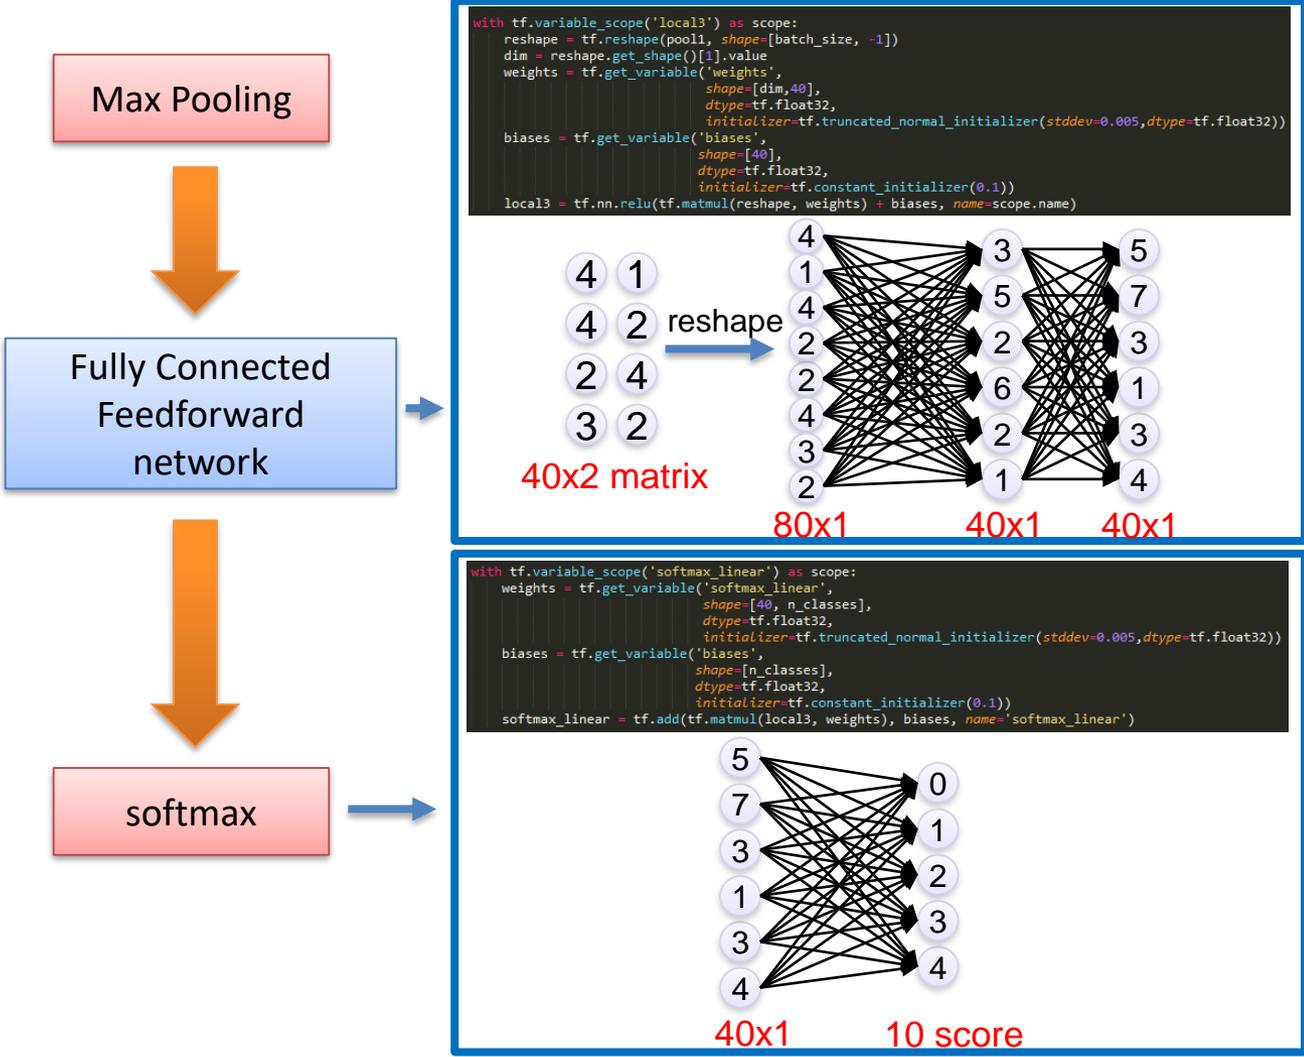
Max-Pooling



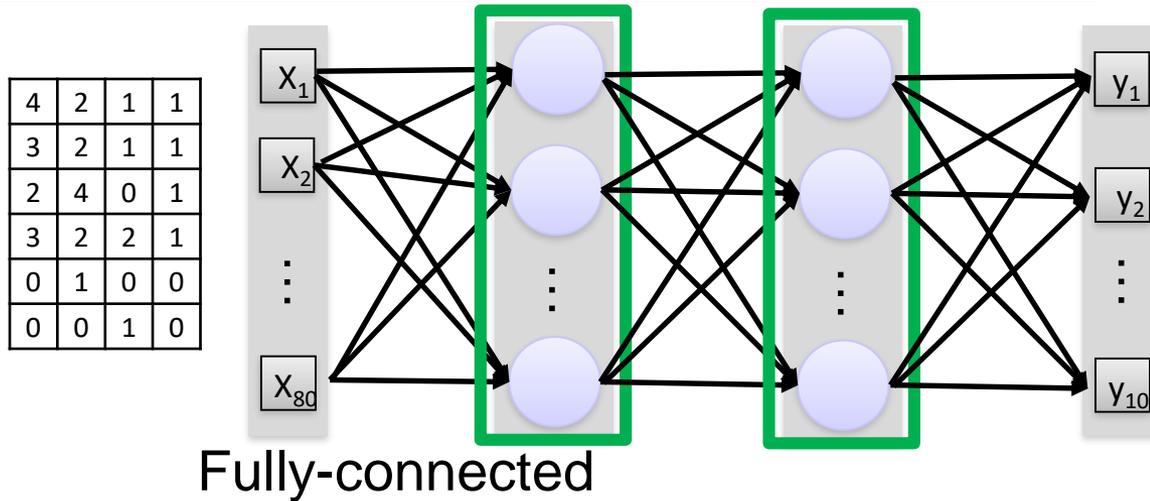
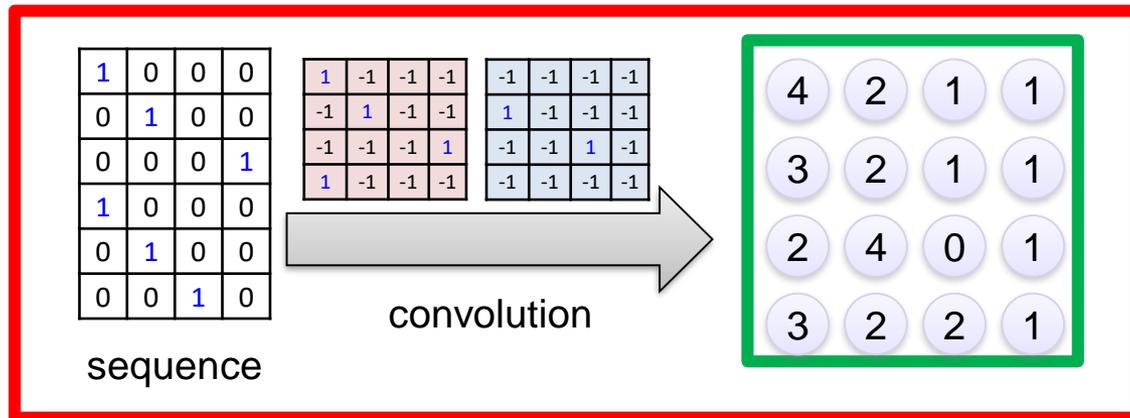
80 x 4 matrix
Forming 40 x 2 matrix

CNN model construction in Tensorflow



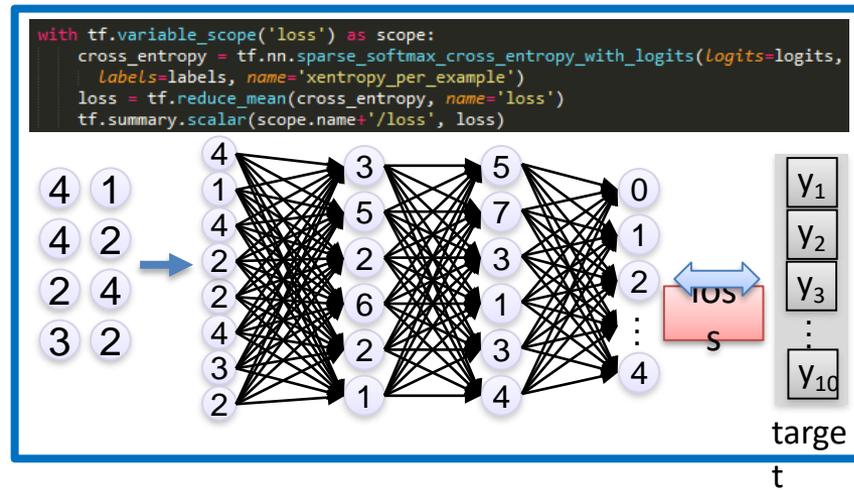


Convolution v.s. Fully Connected



Pick the best model

loss



Randomly initialize network parameters

Pick the 1st batch

$$L' = l_1 + l_2 + \dots$$

Update parameters once

Pick the 2nd batch

$$L'' = l_1 + l_2 + \dots$$

Update parameters once

...

Until all mini-batches have been picked

training

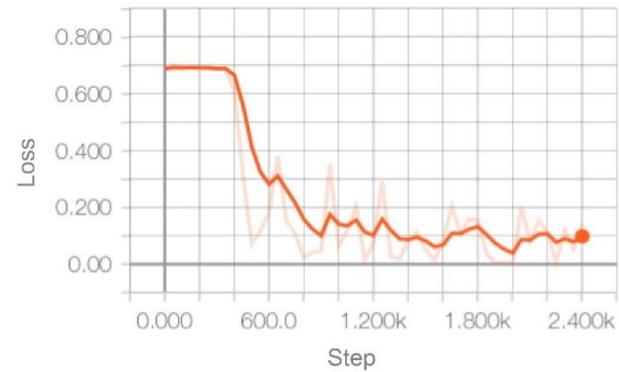
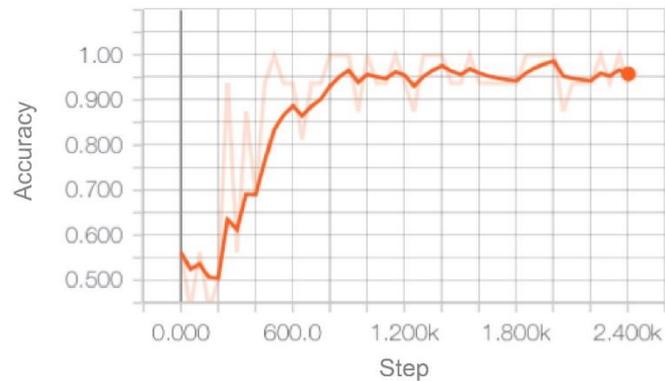
```
N_CLASSES = 2
seq_W = 20 # resize the seq, if the input seq is too large, training will be very slow.
seq_H = 80
BATCH_SIZE = 16
CAPACITY = 2000
MAX_STEP = 10000 # with current parameters, it is suggested to use MAX_STEP>10k
learning_rate = 0.0001 # with current parameters, it is suggested to use learning rate<0.0001
n = 1601
```

Output:

```
Step 0, train loss = 0.69, train accuracy = 62.50%
Step 50, train loss = 0.69, train accuracy = 56.25%
Step 100, train loss = 0.69, train accuracy = 43.75%
Step 150, train loss = 0.69, train accuracy = 56.25%
Step 200, train loss = 0.68, train accuracy = 93.75%
Step 250, train loss = 0.68, train accuracy = 43.75%
Step 300, train loss = 0.57, train accuracy = 68.75%
Step 350, train loss = 0.45, train accuracy = 68.75%
Step 400, train loss = 0.36, train accuracy = 75.00%
Step 450, train loss = 0.41, train accuracy = 87.50%
```

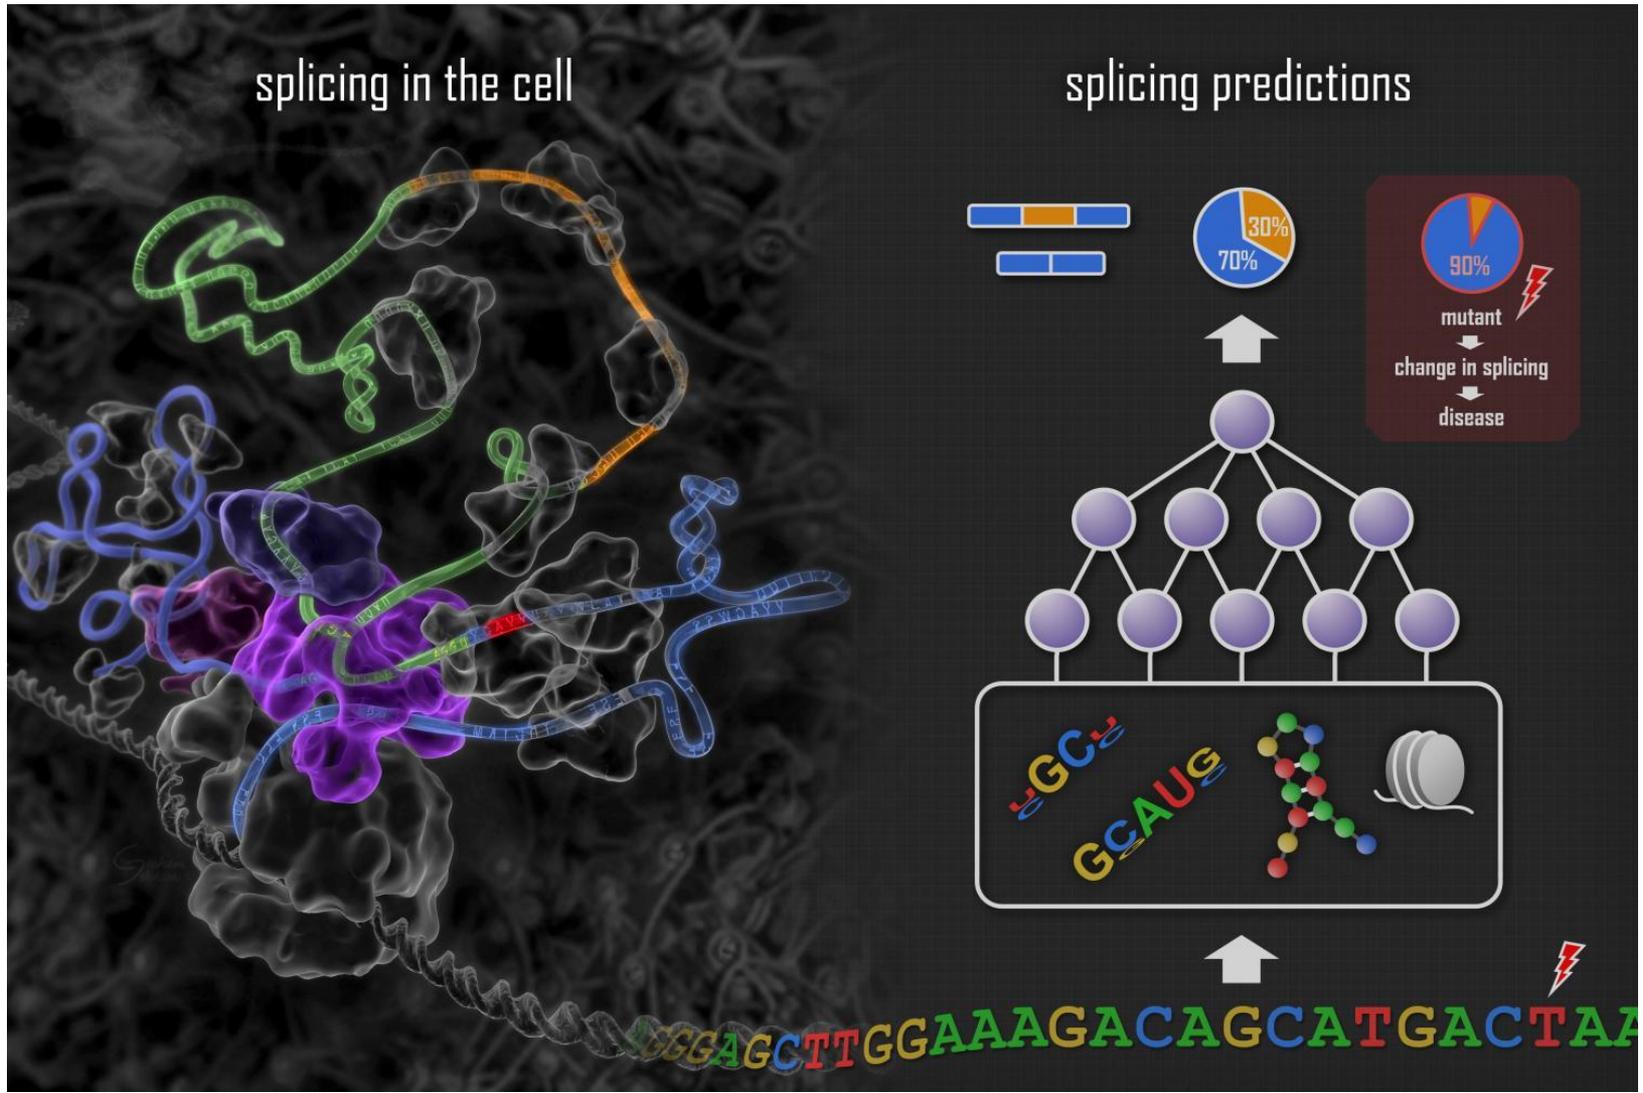
Training Module1

- Data: 5000 Cas9 sequences with domain sites and 5000 without domain sites

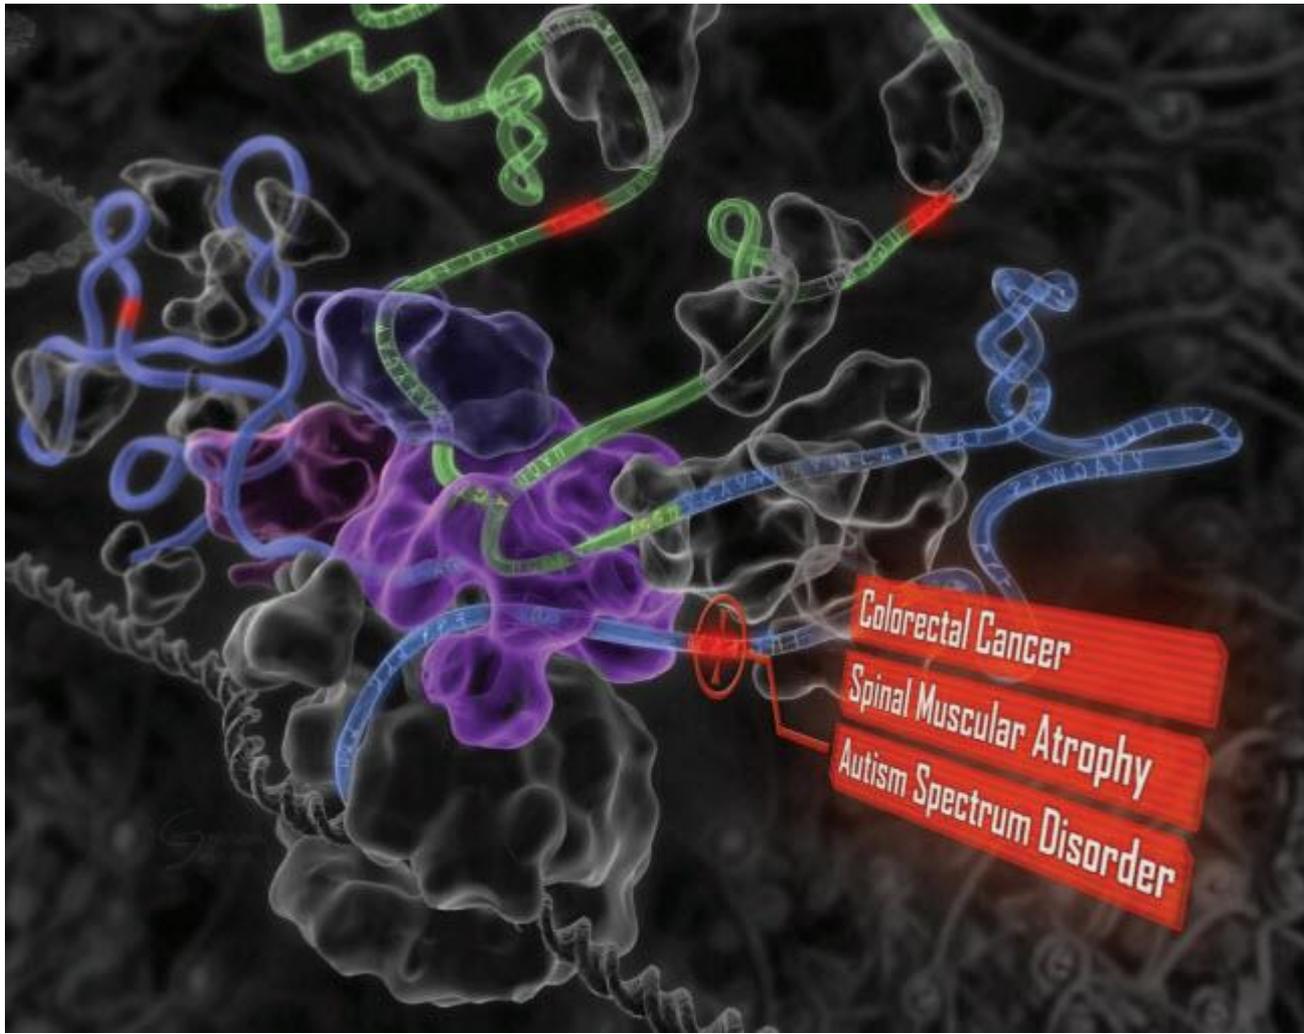


Accuracy: 0.80225

New variant discovery

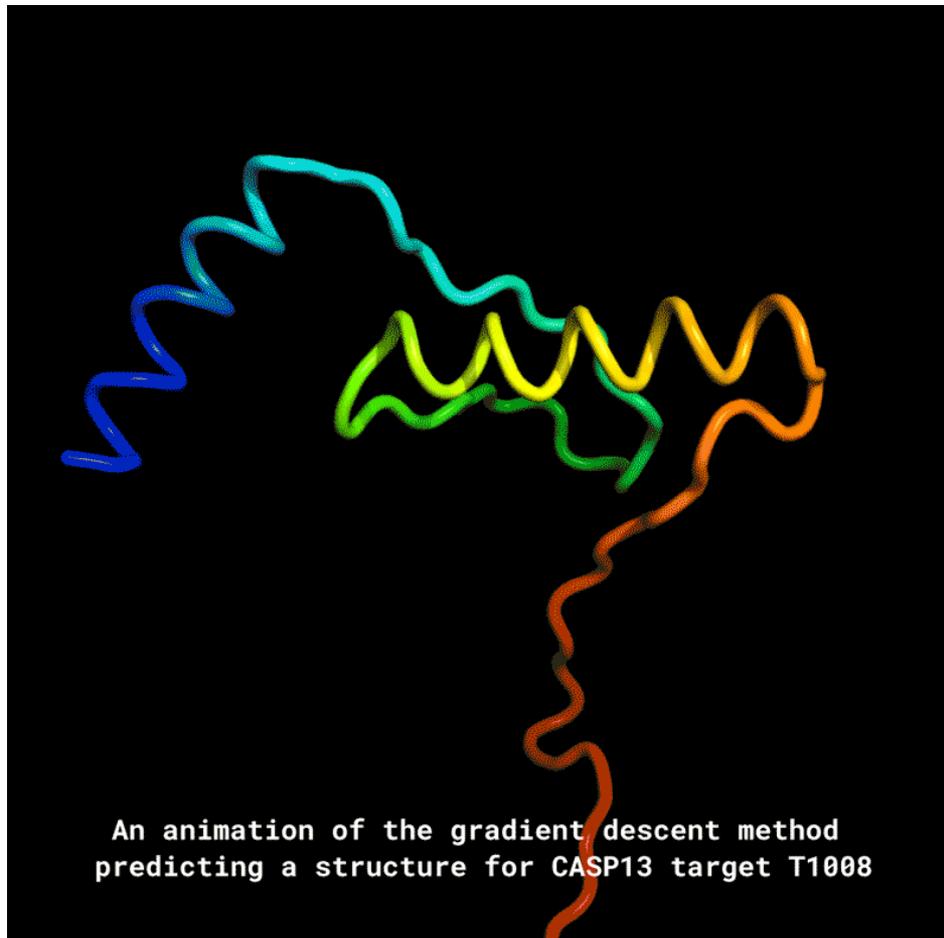


New variant discovery



CASP

(Critical Assessment of Techniques for Protein Structure Prediction)



CASP

(Critical Assessment of Techniques for Protein Structure Prediction)

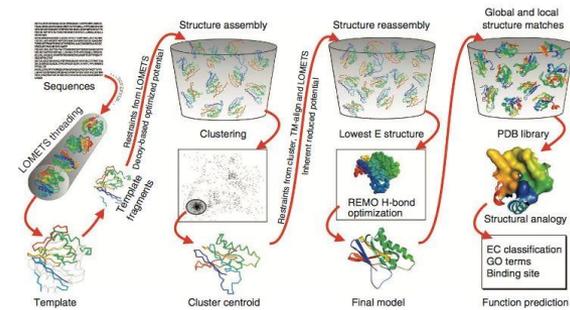
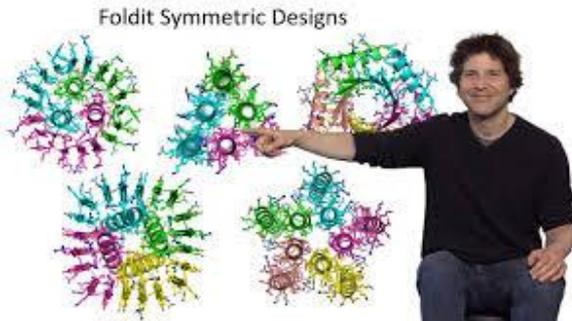


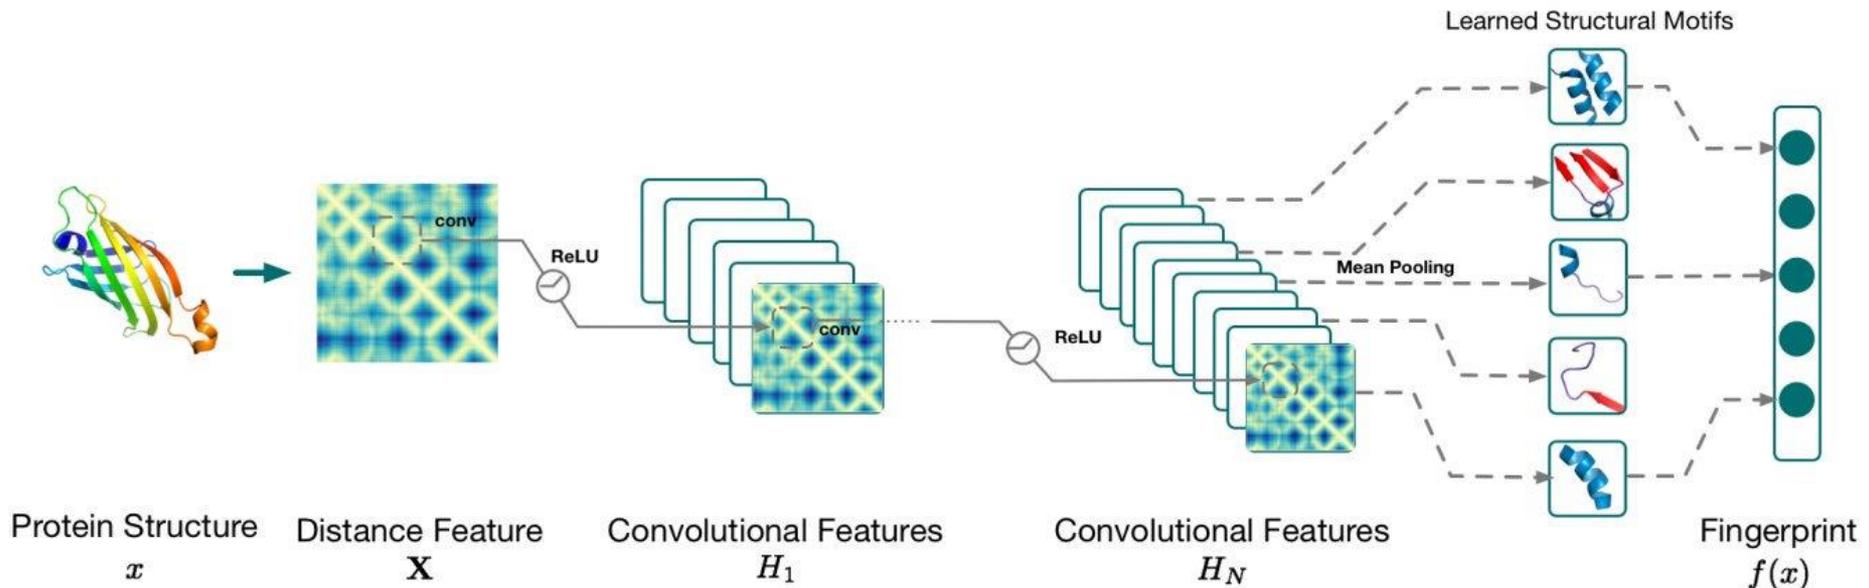
Figure 1 | A schematic representation of the I-TASSER protocol for protein structure and function predictions. The protein chains are colored from blue at the N-terminus to red at the C-terminus. 驻波



CASP

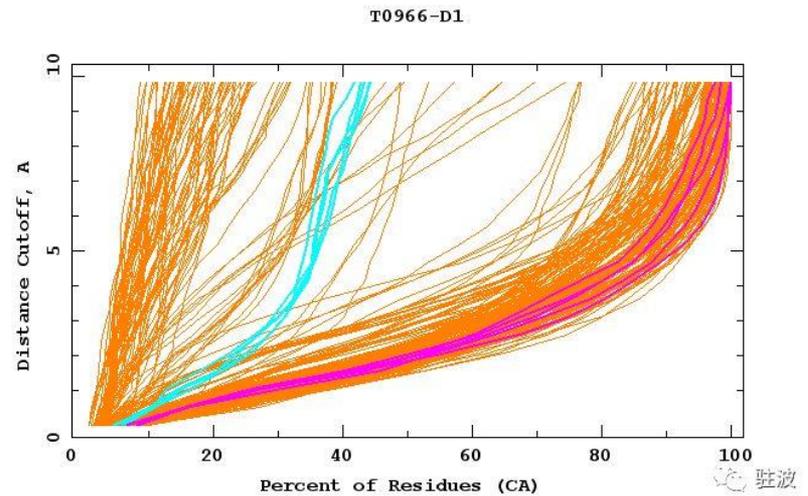
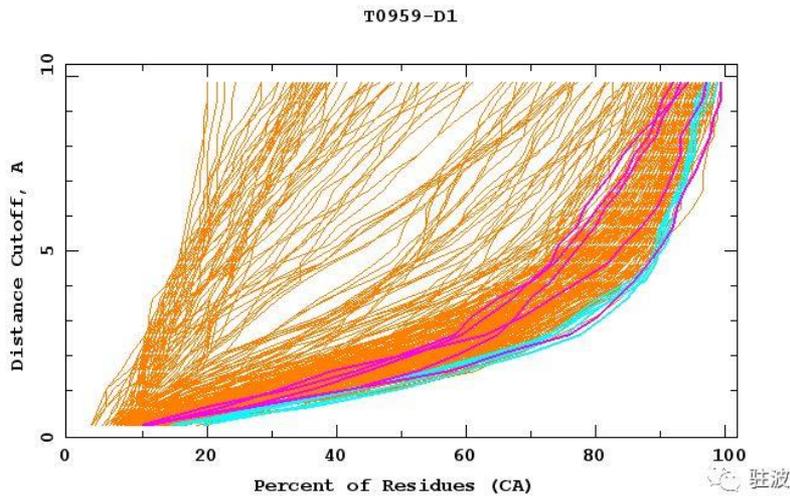
(Critical Assessment of Techniques for Protein Structure Prediction)

DeepFold: overall winner of CASP18



CASP

(Critical Assessment of Techniques for Protein Structure Prediction)



Recap (知识点总结)

- 深度学习：
 - 概念和区别
 - 模型和方法
 - 应用领域
- 生物大数据的深度学习：
 - 序列和特征
 - 特征学习和模型构建
 - 应用案例

References

- Angermueller, C., et al., Deep learning for computational biology. *Molecular systems biology*, 2016. 12(7): p. 878.
- Jurtz, V.I., et al., An introduction to deep learning on biological sequence data: examples and solutions. *Bioinformatics*, 2017. 33(22): p. 3685-3690.
- Wang, S., et al., Protein secondary structure prediction using deep convolutional neural fields. *Scientific reports*, 2016. 6.
- Angermueller, C., et al., DeepCpG: accurate prediction of single-cell DNA methylation states using deep learning. *Genome biology*, 2017. 18(1): p. 67.
- Yuan, L., et al., Applications of Deep Learning in Biological and Medical Data Analysis. *PROGRESS IN BIOCHEMISTRY AND BIOPHYSICS*, 2016. 43(5): p. 472-483.
- Chai, G., et al., HMMCAS: a web tool for the identification and domain annotations of Cas proteins. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 2017.

深度学习的问题。。。

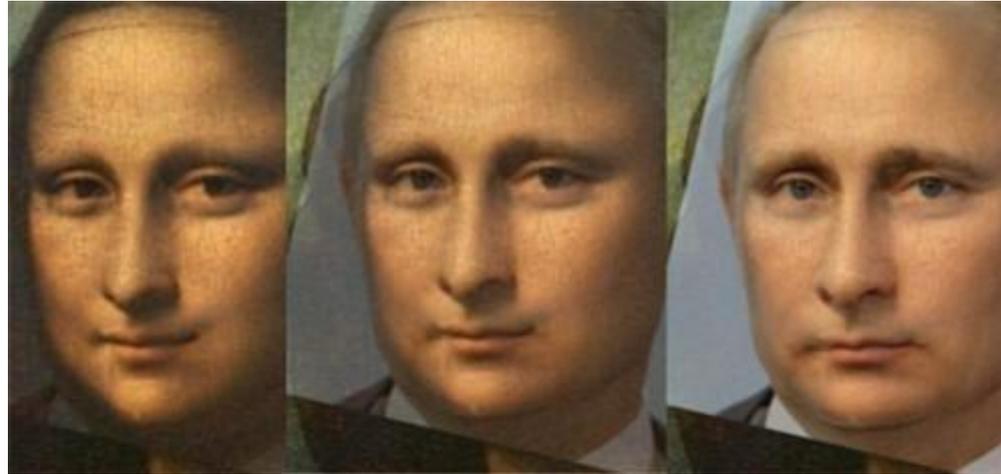


Feature extraction? Or D&C

Case 1

为何对基因的分析结果难以达到共识？“其中一个原因是遗传变异太常见。一个典型的基因包含了数以百计的变异，有些变异可能对健康没有任何影响。有些疾病可能是多个基因变异相互作用的结果，这种作用可能经常被算法所遗漏”，威廉信托基金会桑格研究所遗传学家 **Matthew Hurles**说。Hurles目前正引领一个关于破解发育障碍的项目，该项目分析了1400个家庭中未确诊重症儿的外显子。他说，“即使当一个单一变异能解释疾病，也还需要进行统计调查、分析数据以及开展临床试验，以提供最终的确诊。”

深度学习的问题。。。。



Majority vote? Or rely on expert?

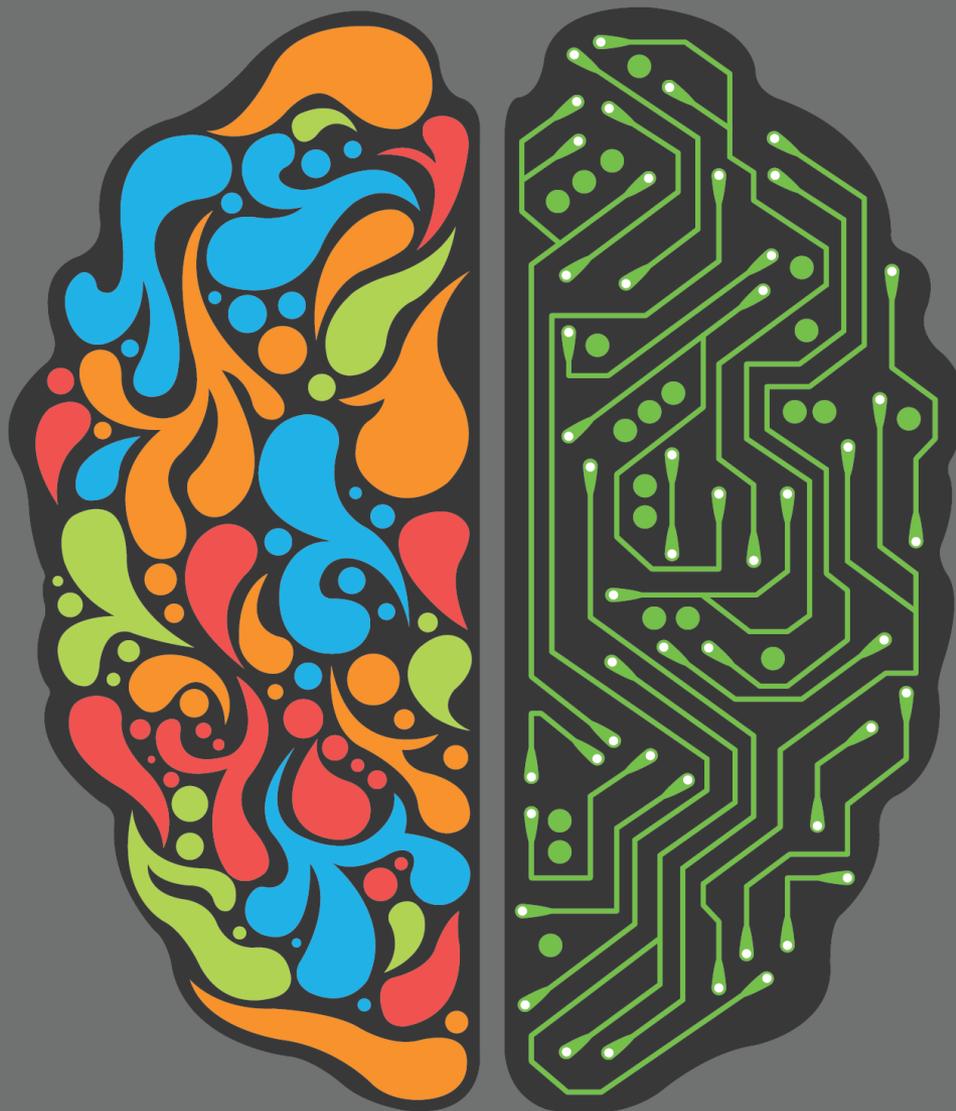
Case 1

一名40岁女子从小被诊断为重型地中海贫血及血小板过多症，后来并被切除脾脏及长期服用抑制血小板功能的药物，进一步做全外显子基因体定序，才发现真正的致病机转是因细胞内质网相关的基因突变引起「先天性红血球生成异常性贫血」，输血治疗地中海贫血反受其害。

Case 2

2007年，美国最大的分子诊断公司Quest Diagnostics 子公司——Athena Diagnostics 给患儿Christian Millare 进行了SCN1A基因突变检测。2008年5月1日，两岁的Christian疾病发作并不幸去世。此后死亡患儿母亲Amy Williams 通过咨询专家和查阅发表的文献，再结合Christian的病历，她坚信儿子的生命不应该这样过早的结束。8年后，Williams对Quest 诊断公司和Athena诊断公司提出了诉讼。

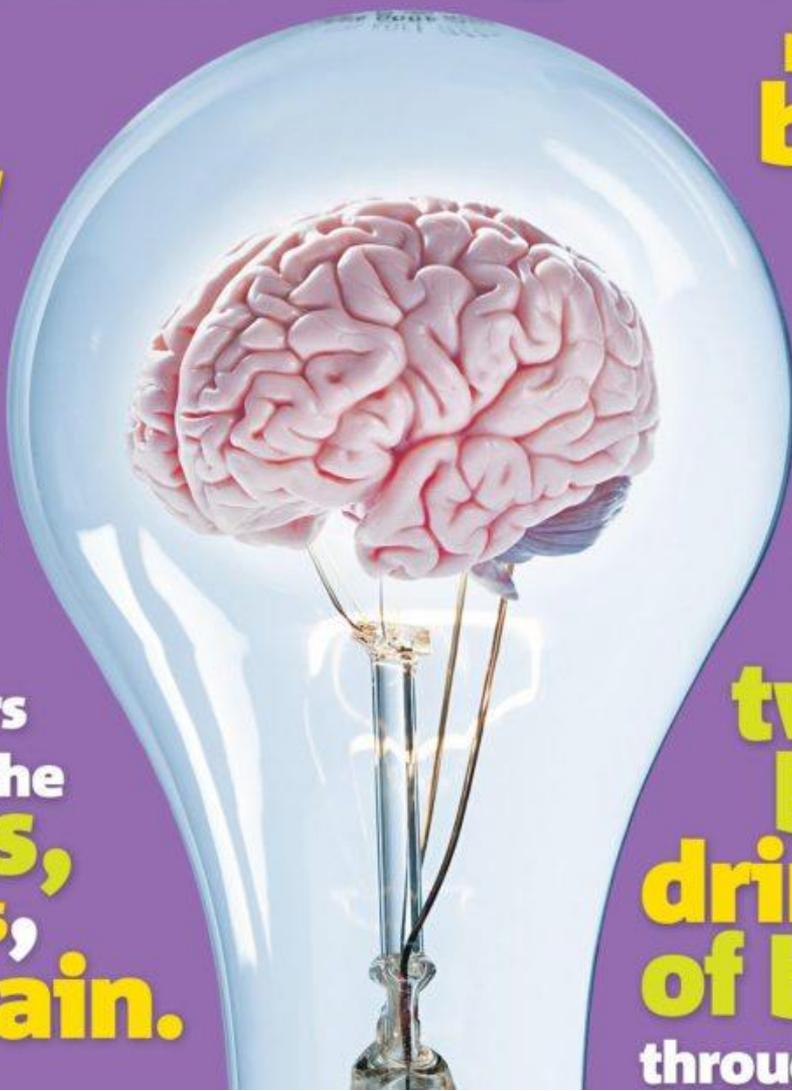
深度学习的问题：脑科学。。。



深度学习的问题：脑科学。。。

1 Your brain generates enough electricity to power a light bulb.

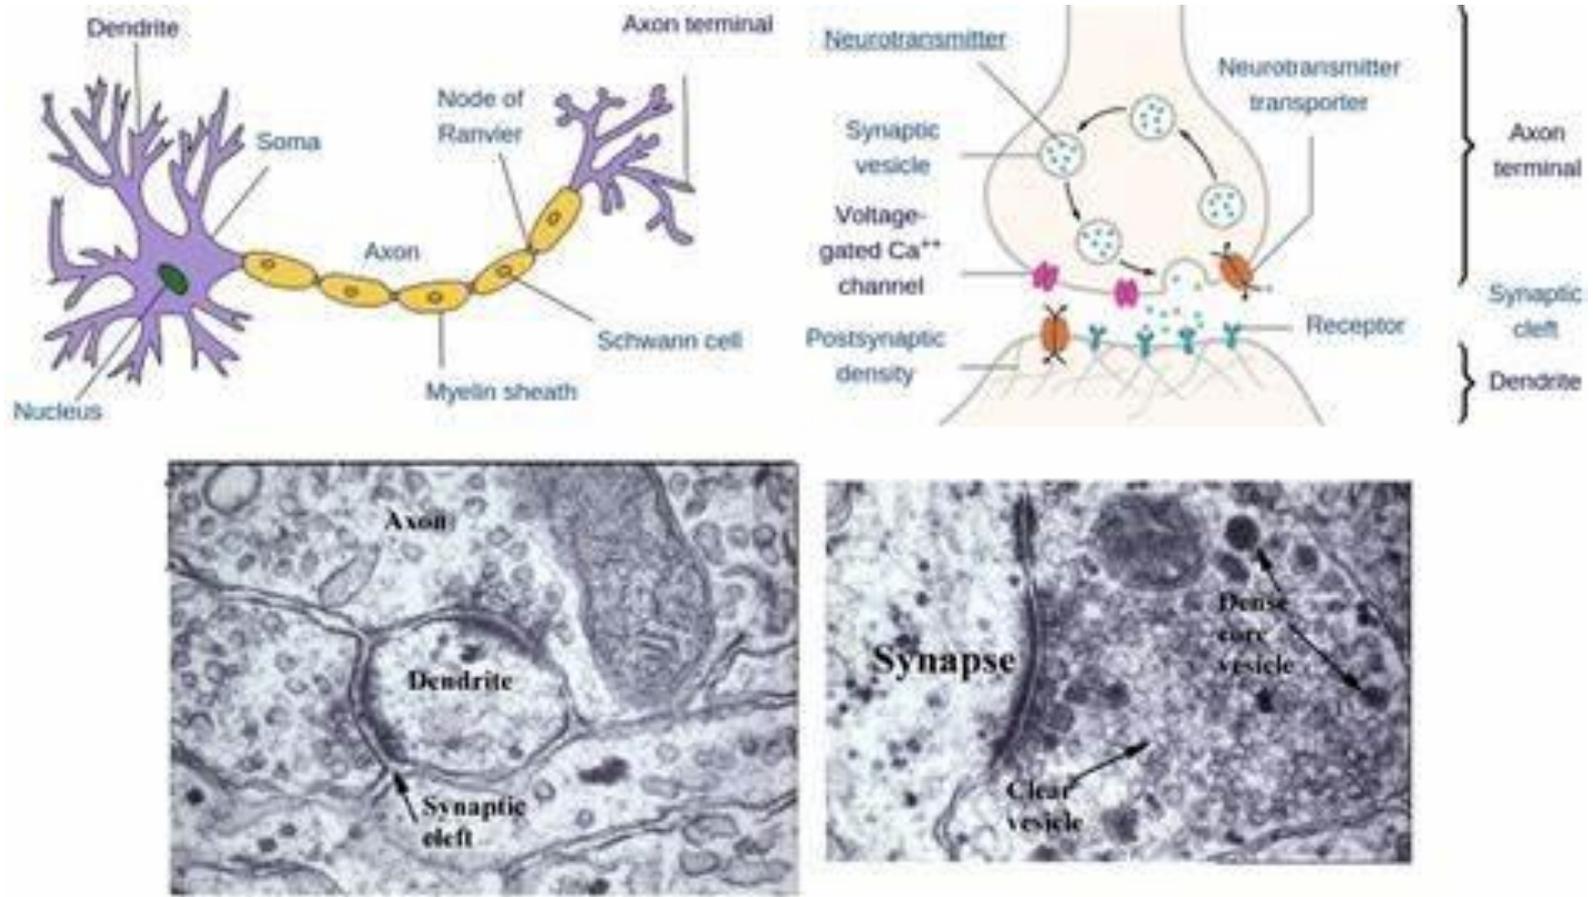
2 It would take close to 3,000 years to count the neurons, or nerve cells, in your brain.



4 Exercise can make your brain work better.

5 Each minute, about 750 millilitres – or two and a bit fizzy drinks cans – of blood travels through the brain.

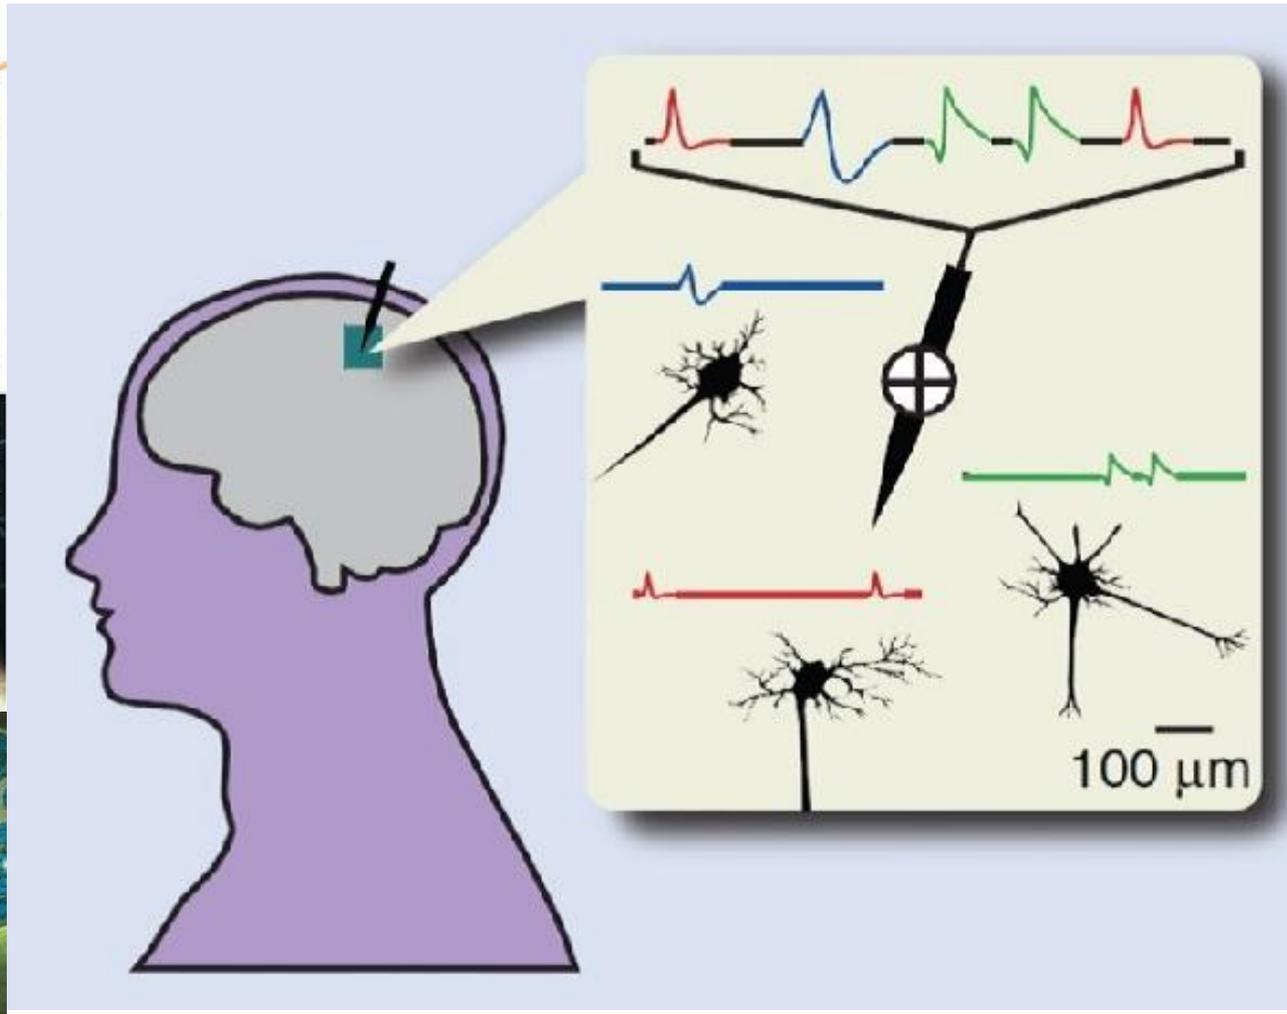
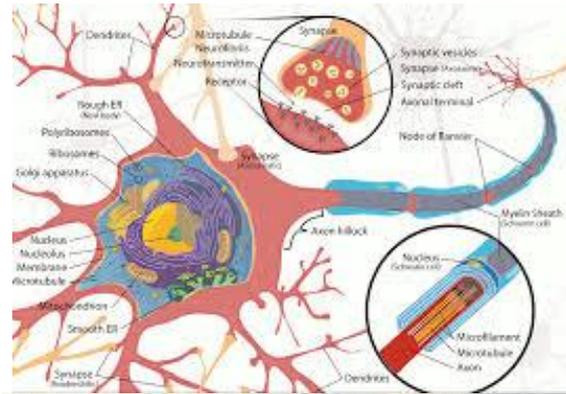
深度学习的问题：脑科学。。。



The Brain vs. Deep Learning vs. Singularity

深度学习的问题：脑科学。。。

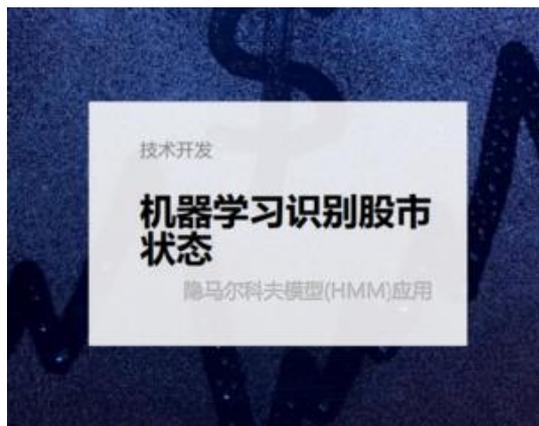
We are actually HMM (or Deep learning) animal...



深度学习的问题。。。

More problems that you can think of?

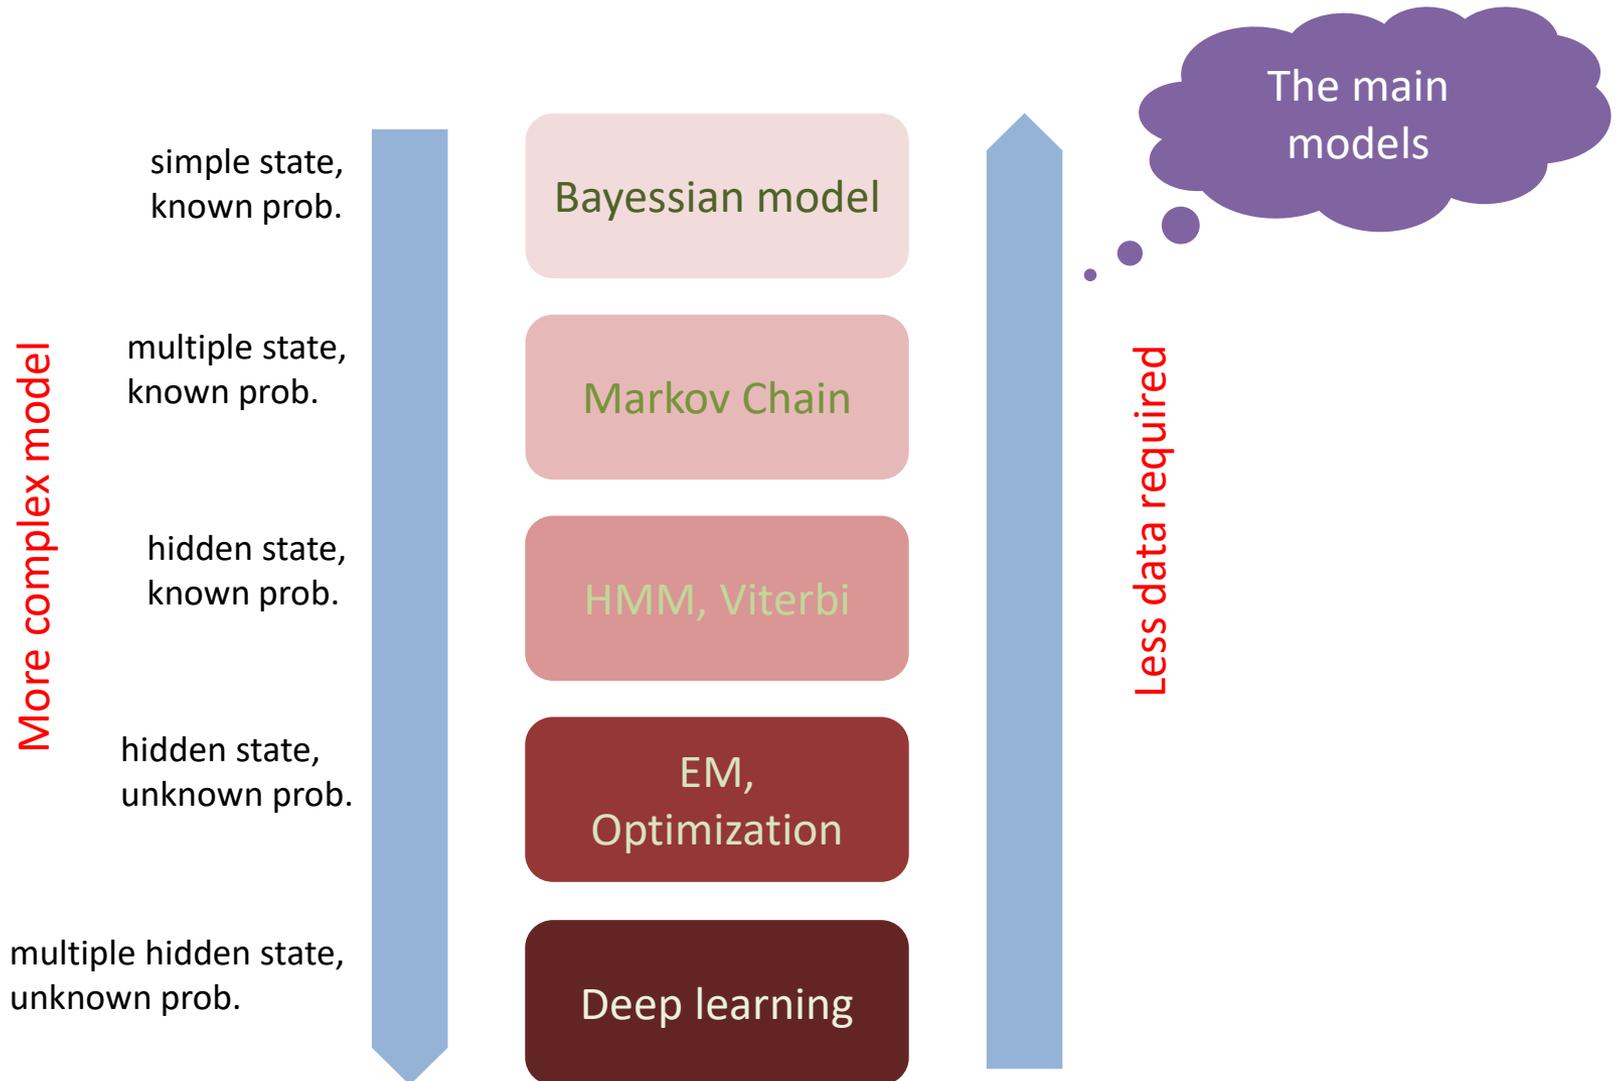
不建议用简单的深度学习来解决的问题。。。



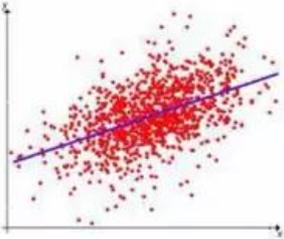
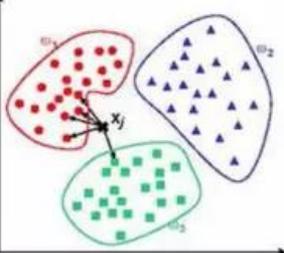
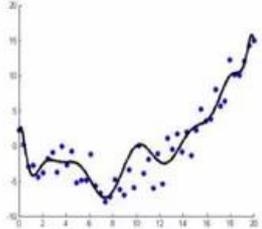
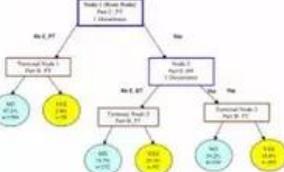
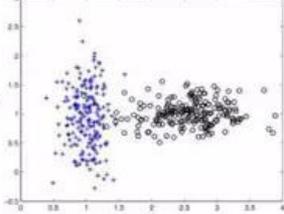
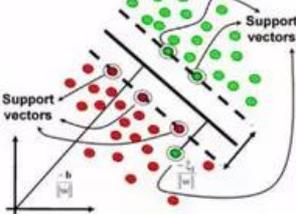
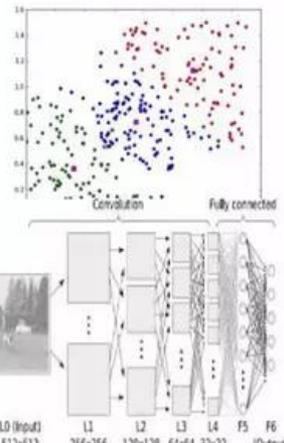
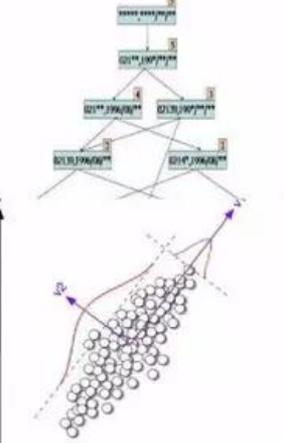
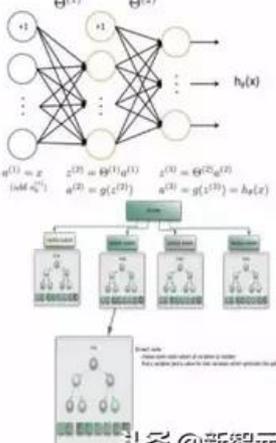
Statistical modeling

Distribution, Modeling, Prediction

Statistical modeling



Statistical modeling

<p>回归算法</p> 	<p>基于实例的算法</p> 	<p>正则化方法</p> 
<p>决策树学习</p> 	<p>贝叶斯方法</p> 	<p>基于核的算法</p> 
<p>聚类算法</p>  <p>LO (Input) 512x512 L1 256x256 L2 128x128 L3 64x64 L4 32x32 F5 F6 (Output)</p>	<p>关联规则学习</p> 	<p>人工神经网络</p>  <p> $a^{(1)} = x$ $z^{(2)} = \theta^{(1)T} a^{(1)}$ $a^{(2)} = \sigma(z^{(2)})$ $(\text{and } a_0^{(1)} = 1)$ $a^{(2)} = g(z^{(2)})$ $a^{(2)} = g(z^{(2)}) = h_{\theta}(x)$ </p>

课程安排

- 生物背景和课程简介
- 传统生物统计学及其应用
- 生物统计学和生物大数据挖掘
 - Hidden Markov Model (HMM)及其应用
 - Markov Chain
 - HMM理论
 - HMM和基因识别 (Topic I)
 - HMM和序列比对 (Topic II)
 - 进化树的概率模型 (Topic III)
 - Motif finding中的概率模型 (Topic IV)
 - EM algorithm
 - Markov Chain Monte Carlo (MCMC)
 - 基因表达数据分析 (Topic V)
 - 聚类分析-Mixture model
 - Classification-Lasso Based variable selection
 - 基因网络推断 (Topic VI)
 - Bayesian网络
 - Gaussian Graphical Model
 - 基因网络分析 (Topic VII)
 - Network clustering
 - Network Motif
 - Markov random field (MRF)
 - Dimension reduction及其应用 (Topic VIII)
- 面向生物大数据挖掘的深度学习

研究对象：
生物序列，
进化树，
生物网络，
基因表达

...

方法：
生物计算与生物统计

Good luck in biostatistics!



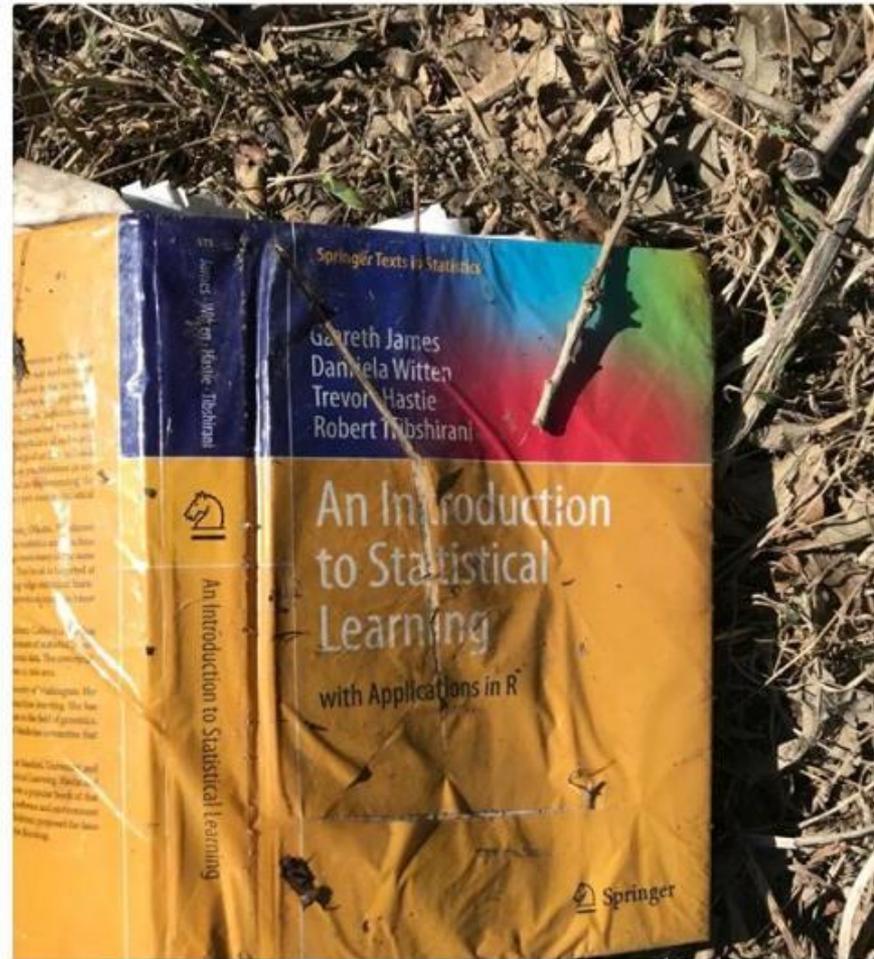
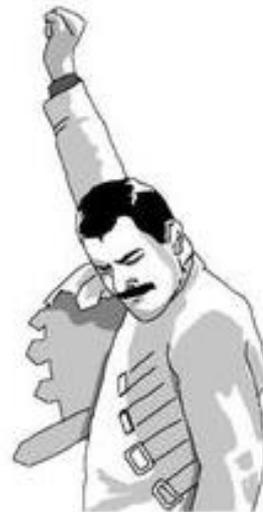
Noah Williams
@Bellmanequation

Found on the roadside:
[@SpringerStats](#) yellow book,
beside empty vodka bottles &
cigarette packs

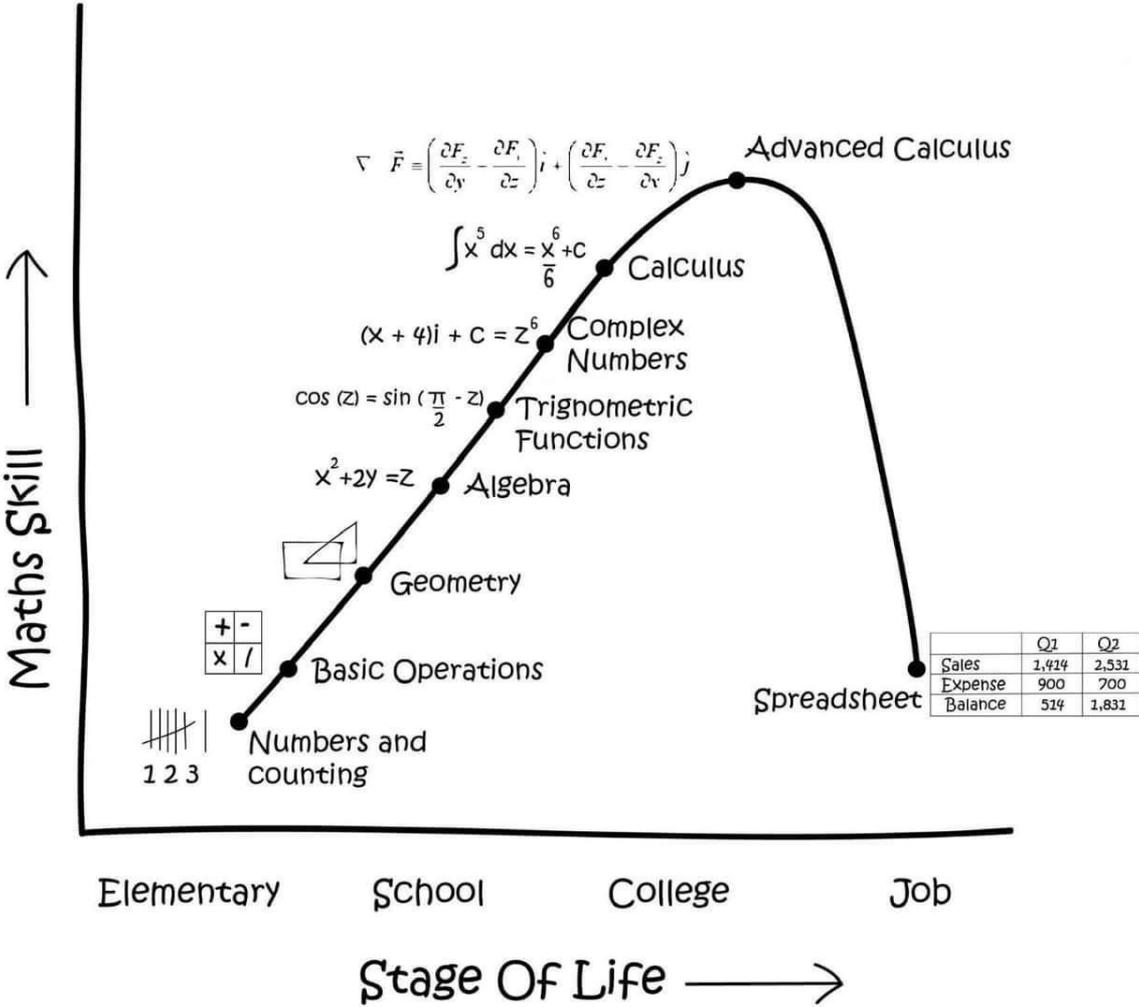


P > .05

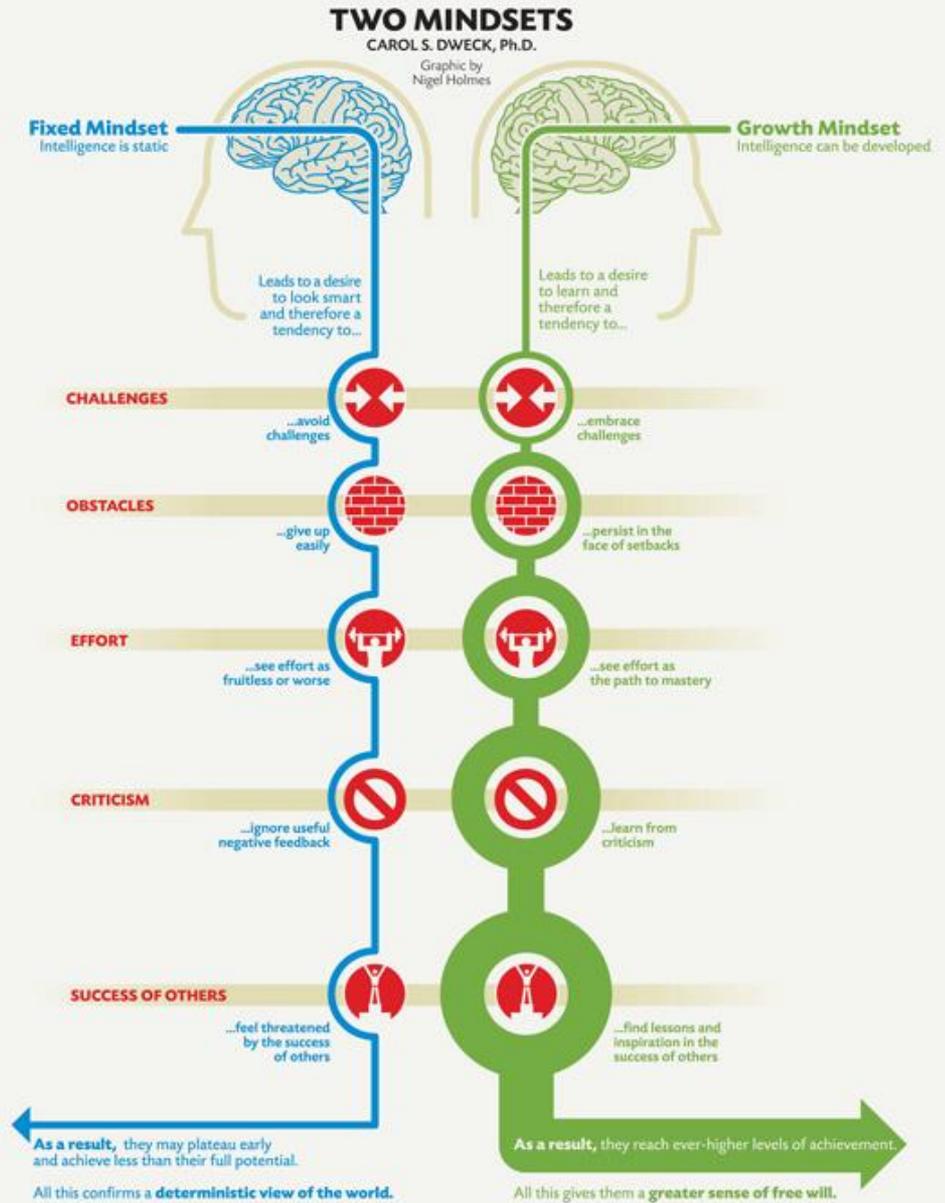
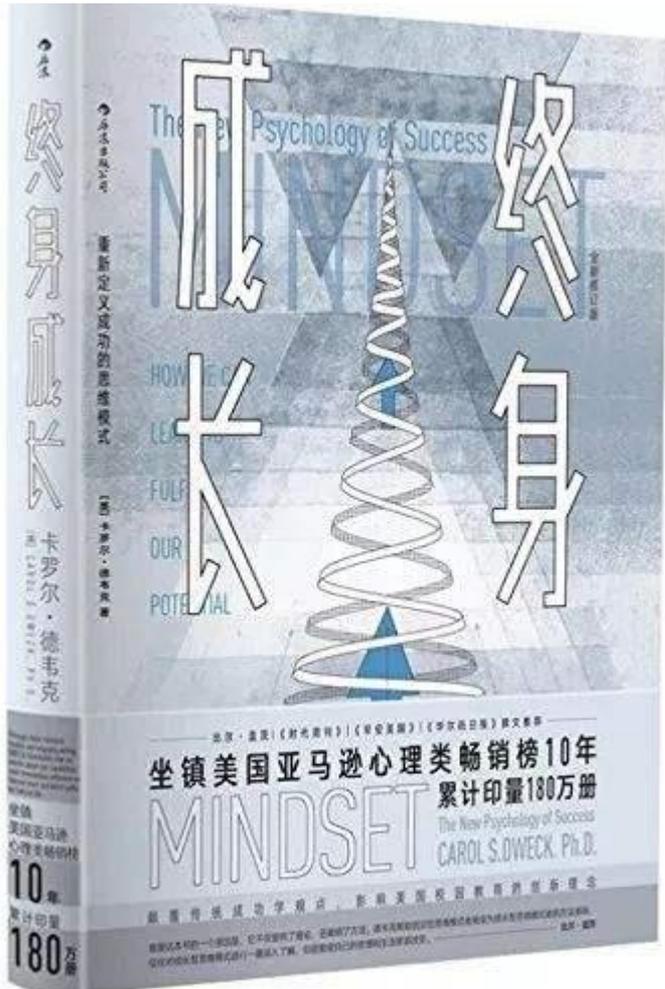
$p < 0.05$



Maybe you will not face these many statistical problems...



But you have the knowledge



Statistical modeling

More to learn:

- **Fast Fourier Transformation:**
<https://zhuanlan.zhihu.com/p/19763358>
- **Causality analysis:**
<https://zhuanlan.zhihu.com/p/33860572>
- ...

Statistical modeling

More to learn:

- 生物统计与生物信息的区别与联系? (<https://www.zhihu.com/question/30284194>)
 - 生物学的生物信息学: 首先是讲进化, 做序列比对, blast, 构建进化树; 然后, 讲基因功能, 基因功能富集分析等。侧重的是, 将进化中的生物学原理, 参数如何设置, 软件如何使用, 如何将这此软件应用到生物学问题。
 - 计算机系的生物信息学: 首先讲的也是进化, 但是, 讲的是算法设计; 然后讲了很多马尔科夫链, HMM模型, 序列比对中的blosum62矩阵是怎么来的, 如何加快计算效率, 如何降低空间存储等等。侧重的是, 如何设计合理的算法, 给生物学的人使用。
 - 数学系的系统生物学: 讲了回归中的penalty function, 然后我学了LASSO回归; 后来学了SVM中的核函数; 还有Gibbs sampling, MCMC; 还有时间序列中的Granger因果推断等等。
 - 生统对数学要求更高, 生信基本不需要。国内生统大多是挂XXX卖XXX
- Yoshua Bengio:
 - **你死我活那是邪教, 开放包容才是正道。**
- **What do you think?**

Biostatistics in the market...

[公司文件] 加强与国内大学合作，吸纳全球优秀人才，共同推动中国基础研究
——任总与中国科学技术大学包信和校长座谈的讲话 

2018-12-13 17:59  6799  113

[只看楼主](#)

总裁办电子邮件

电邮讲话【2018】128号

签发人：任正非

加强与国内大学合作，吸纳全球优秀人才，共同推动中国基础研究
——任总与中国科学技术大学包信和校长座谈的讲话
2018年11月19日

在高校学科设置上，我特别支持你们重视**统计学**。计算机科学不仅仅是技术，还应该以**统计学**为基础。大数据需要**统计学**，信息科学需要**统计学**，**生命科学也需要统计学**。国家要搞人工智能，更要重视**统计学**。**统计学**不是一个纯粹的学科，而是每一个学科都要以**统计学**为基础。

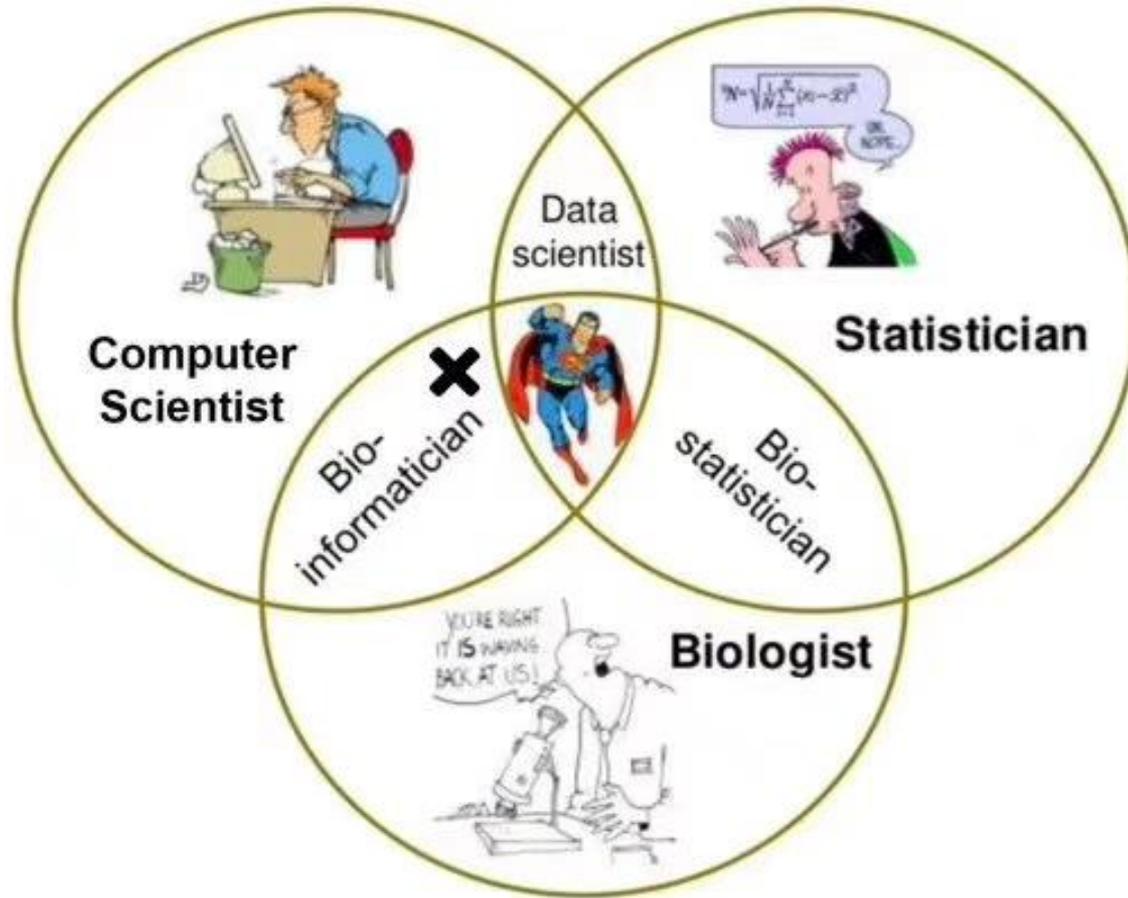


Or



Good luck in final exam!

Superman/Wonder woman



Good luck in your life!

谢谢大家!

- 
- I thank you from the bottom of my heart
 - That's so kind of you
 - I am very thankful
 - You are great
 - I'm really grateful
 - Thanks a million



OTHER WAYS TO SAY
"THANK YOU"

www.englishstudypage.com



Please accept my deepest thanks •

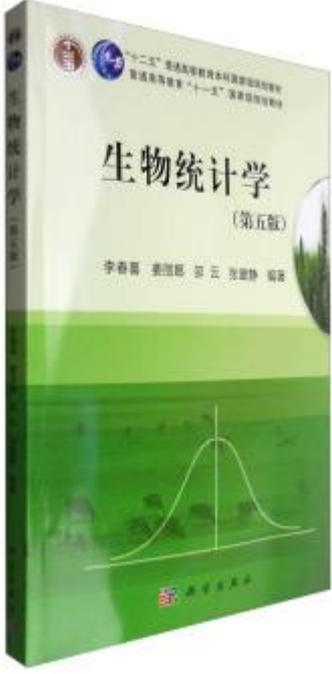
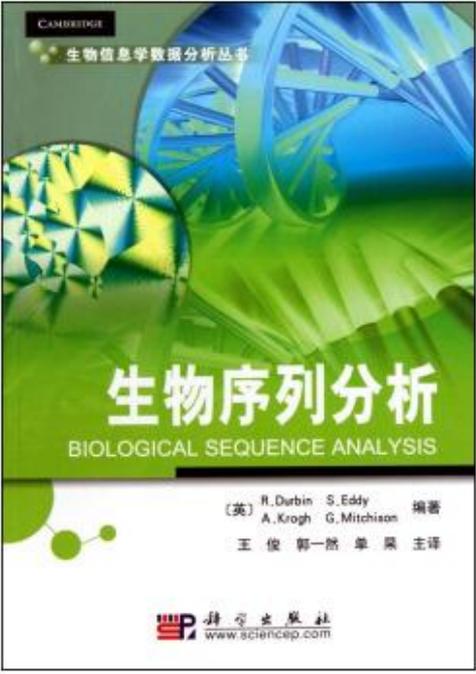
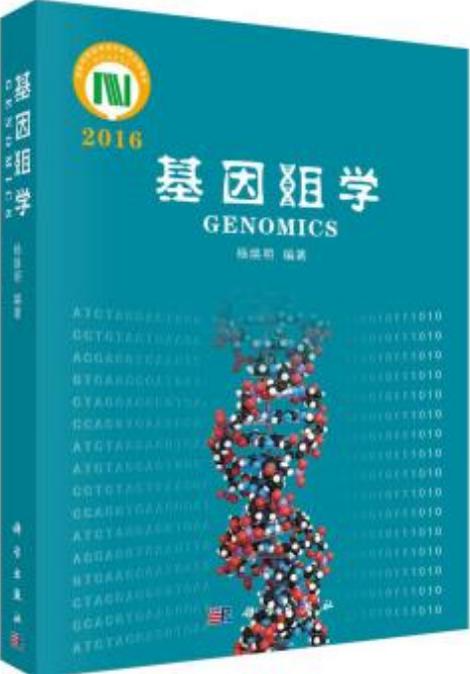
- I'm in your debt
- You are the best
- I owe you one
- Thank you so much
- I really appreciate it

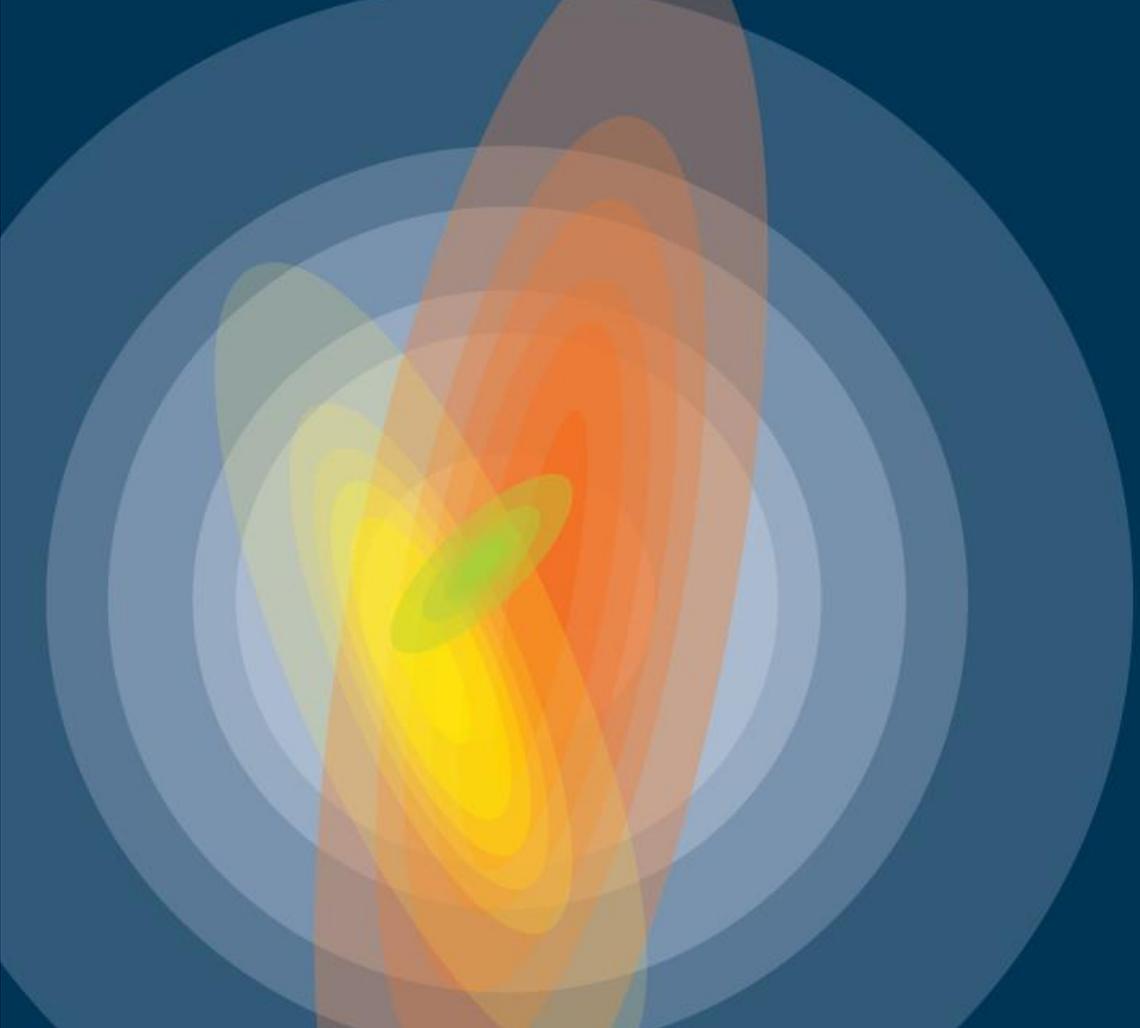


希望大家和我一起，
终身学习！



References

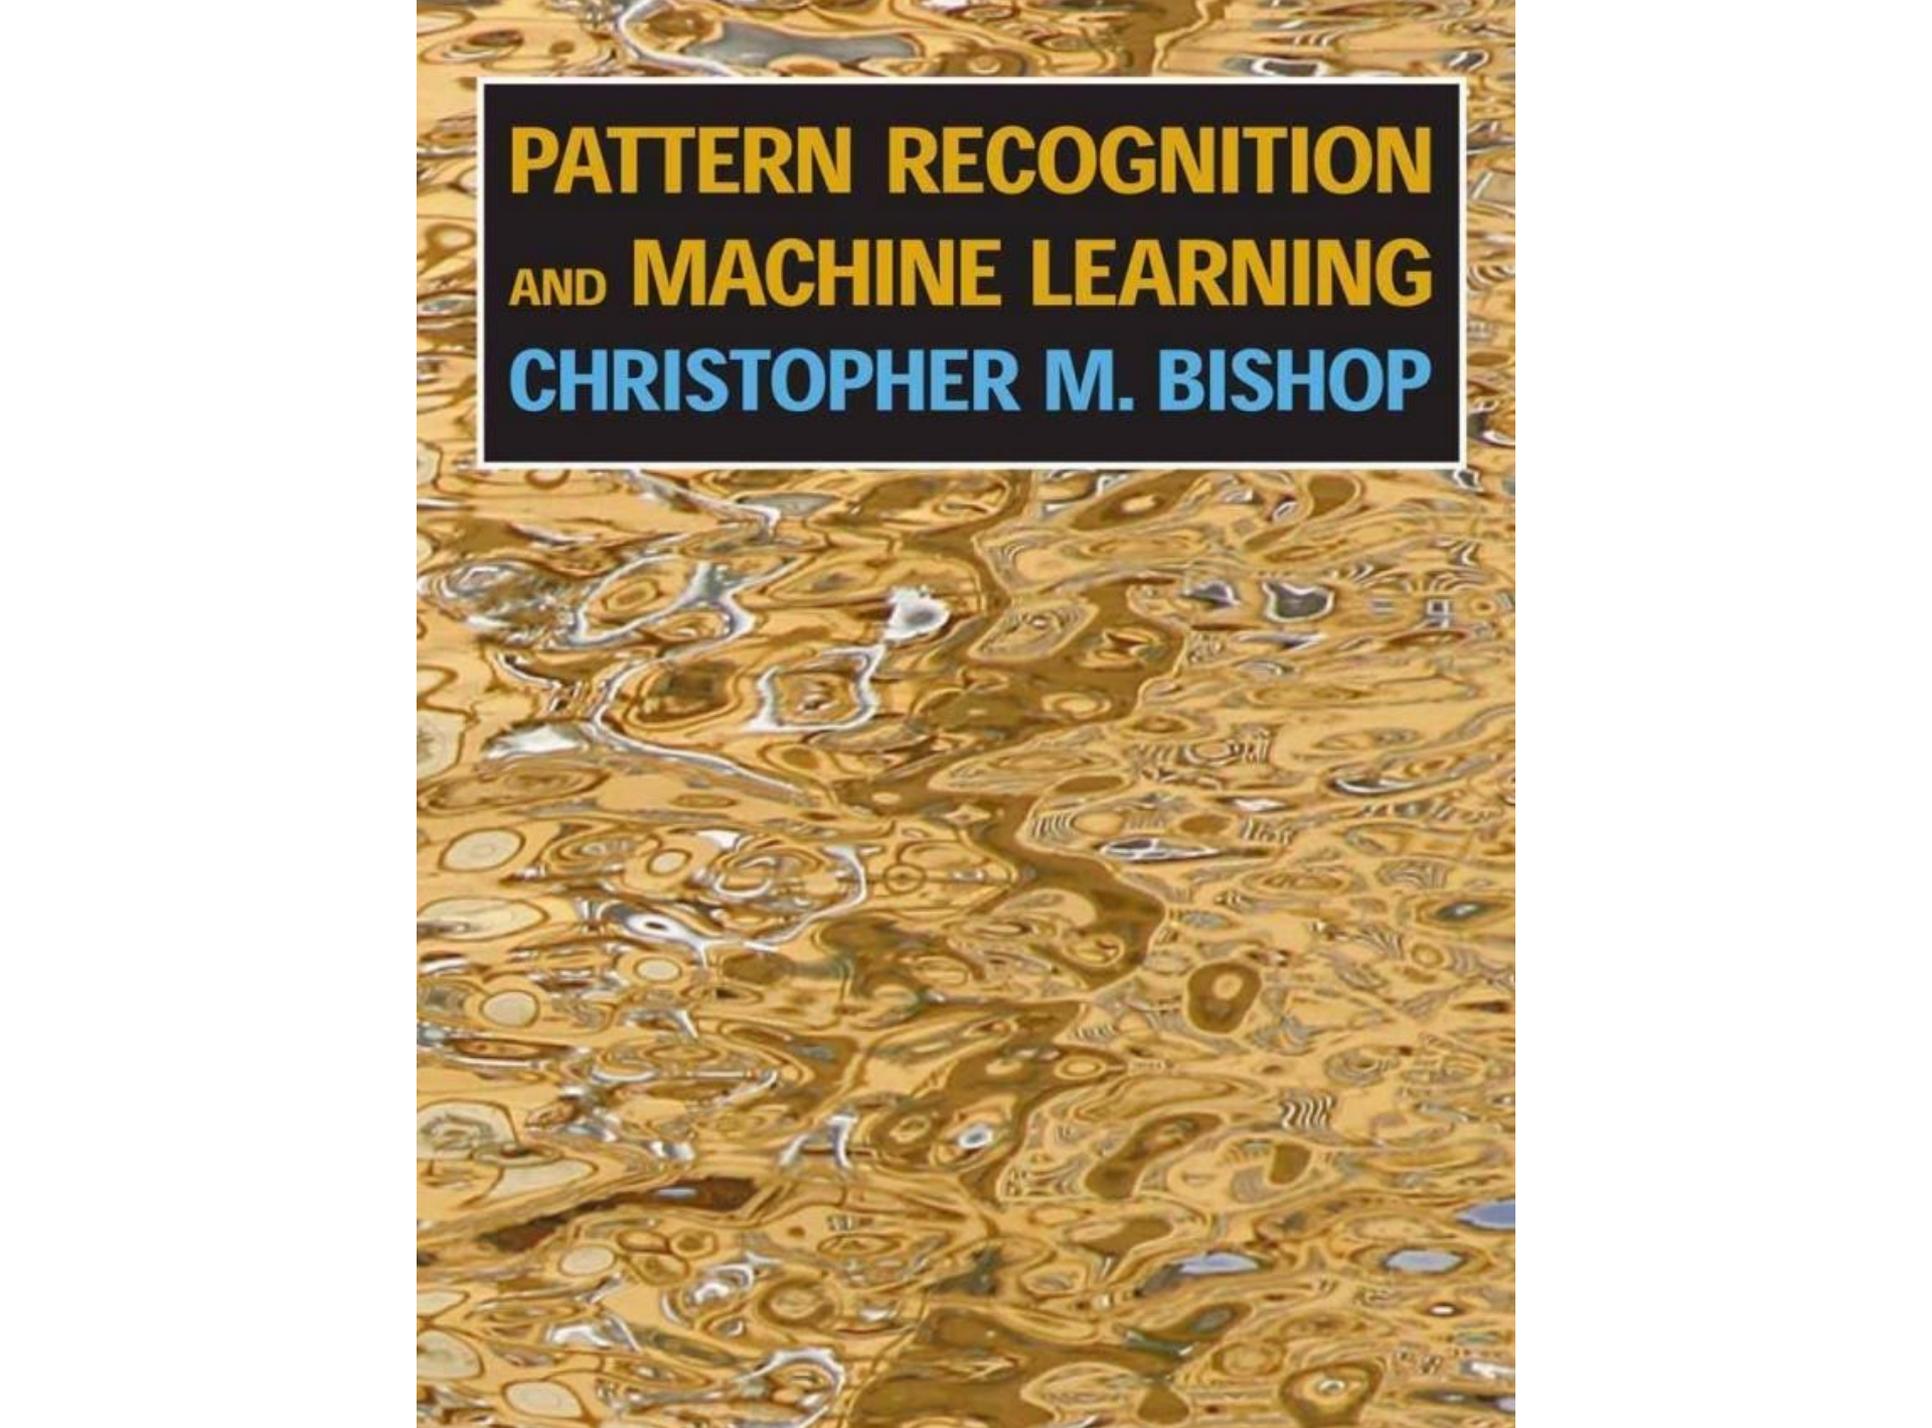




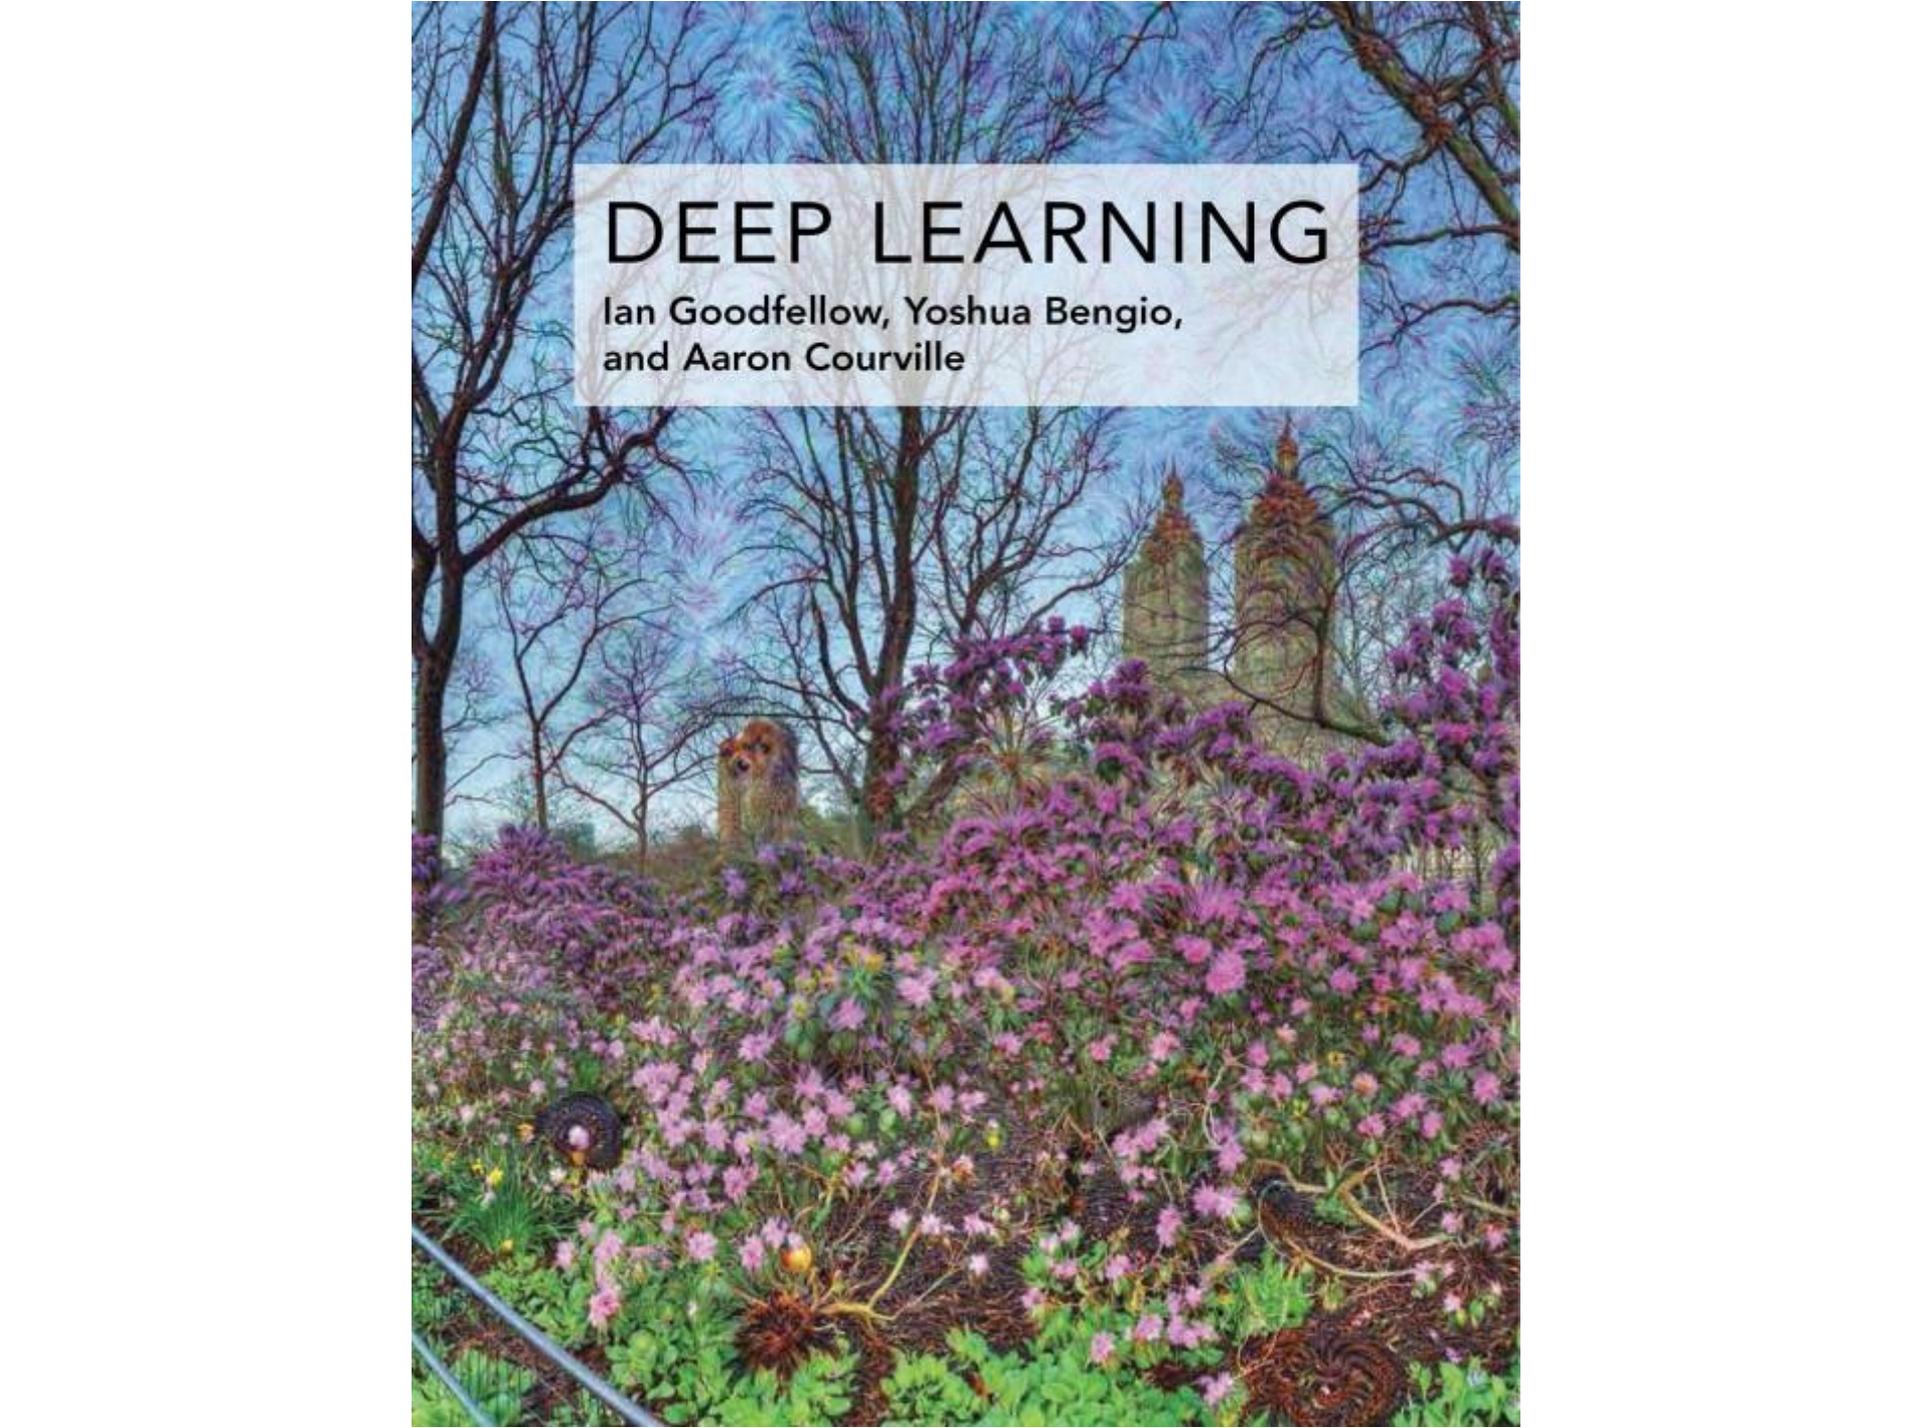
Machine Learning

A Probabilistic Perspective

Kevin P. Murphy



**PATTERN RECOGNITION
AND MACHINE LEARNING
CHRISTOPHER M. BISHOP**

A photograph of a garden scene. In the foreground, there is a dense patch of purple flowers, possibly azaleas, with some green foliage and brown mulch. In the background, several trees with bare branches are visible against a blue sky. Two stone structures, possibly pagodas or towers, are partially obscured by the trees. A white rectangular text box is overlaid on the upper part of the image.

DEEP LEARNING

Ian Goodfellow, Yoshua Bengio,
and Aaron Courville

