The State of The Art

School of Electronic and Computer Engineering
Peking University

Wang Wenmin

PoAI

Contents

Four Categories of AI

	Humanly	Rationally
Acting	Acting humanly	Acting rationally
Thinking	Thinking humanly	Thinking rationally

- Humanly

  to measure success in terms of fidelity to *human* performance.

  类人地：以对人类表现的逼真度来衡量。

- Rationally

  to measure against an *ideal* performance measure.

  理性地：用理想的性能表现来衡量。

  - A system is rational if it does the right thing, given what it knows.

    一个系统如果对已知的知识做出正确的动作，则被称为理性。

Eight Definitions for four Categories of AI

Acting humanly

- Kurzweil,1990: To perform functions that require intelligence performed by people.
  完成需要人类智能所能完成的功能。
- Rich and Knight,1991: To make computers do things at which, at the moment, people are better.
  使计算机去做此时此地人类才能做好的事情。

Acting rationally

- Poole et al.,1998: Computational Intelligence is the study to design intelligent agents.
  计算智能是研究如何设计智能体。
- Nilsson,1998: AI is concerned with intelligent behavior in artifacts.
  AI是关注于用人工手段去实现智能行为。

Eight Definitions for four Categories of AI

Thinking humanly

- Bellman,1978: The automation of activities that we associate with human thinking …
  我们与人类思维相关活动的自动化 …
- Haugeland,1985: The new effort to make computers think … machines with minds …
  新的努力使计算机思考 … 机器具有智力 …

Thinking rationally

- Charniak and McDermott,1985: The study of mental faculties through the use of computational models.
  通过使用计算模型进行心智能力的研究。
- Winston,1992: To make computer possible to perceive, reason, and act.
  使计算机能够感知、推理、以及动作。

Weak AI vs. Strong AI vs. Super AI

Weak AI

- Also called Artificial Narrow Intelligence (ANI).
  弱人工智能：也被称为人工狭义智能(ANI)。
- It is non-sentient AI that is focused on one narrow task (just a specific problem).
  它是无意识的AI，专注于一个具体的任务（仅针对一个特定的问题）。

Strong AI

- Also called Artificial General Intelligence (AGI).
  强人工智能：也被称为人工广义智能（AGI)。
- It means a machine with the ability to apply intelligence to any problem. It is a primary goal of artificial intelligence research.
  意味着机器具有将智能用于处理任何问题的能力。它是人工智能研究的主要目标。

Weak AI vs. Strong AI vs. Super AI

Super AI

- Also called Artificial Super Intelligence (ASI).
  超人工智能：亦称人工超级智能(ASI)。
- It is a hypothetical agent that possesses intelligence far surpassing that of the brightest and most gifted human minds.
  是一个假定的智能体，拥有远远超过聪明和最有天赋的人类大脑的智能。
- Also refer to a property of problem-solving systems, e.g., super intelligent language translators or engineering assistants.
  也指的是问题求解系统的特性，例如，超级智能语言翻译器或工程助理。

Typical Problems to Which AI is Applied

Computer vision	计算机视觉
Image processing	图像处理
VR, AR and MR	VR, AR 和 MR
Pattern recognition	模式识别
Intelligent Diagnosis	智能诊断
Game theory and Strategic planning	博弈论和策略规划
Game AI and Gamebot	AI游戏和游戏机器人
Machine Translation	机器翻译
Natural language processing, and Chatbot	自然语言处理和聊天机器人
Nonlinear control, and Robotics	非线性控制和机器人技术

Other Fields in Which AI is Implemented

Artificial life	■	智能生活
Automated reasoning	■	自动推理
Automation	■	自动化
Biological computing	■	生物计算
Concept mining	■	概念计算
Data mining	■	数据挖掘
Knowledge representation	■	知识表示
Semantic Web	■	语义Web
E-mail spam filtering	■	垃圾邮件过滤
Litigation	■	诉讼

Other Fields in Which AI is Implemented

- Robotics — 机器人学
- Behavior-based robotics
- Cognitive
- Cybernetics
- Development robotics
- Evolutionary robotics
- Hybrid intelligent system — 混合人工智能
- Intelligent agent — 智能代理
- Intelligent Control — 智能控制

Joshua Tenenbaum, Vin Silva, John Langford.

"A Global Geometric Framework for Nonlinear Dimensionality Reduction".
SCIENCE, Vol. 290, Dec. 2000.

"一种用于非线性降维的全局几何框架"

"Here we describe an approach to solving dimensionality reduction problems that uses easily measured local metric information to learn the underlying global geometry of a data set."

"我们在此描述一种解决降维问题的方法，使用易测的局部度量信息来学习数据集潜在的全局几何结构。"

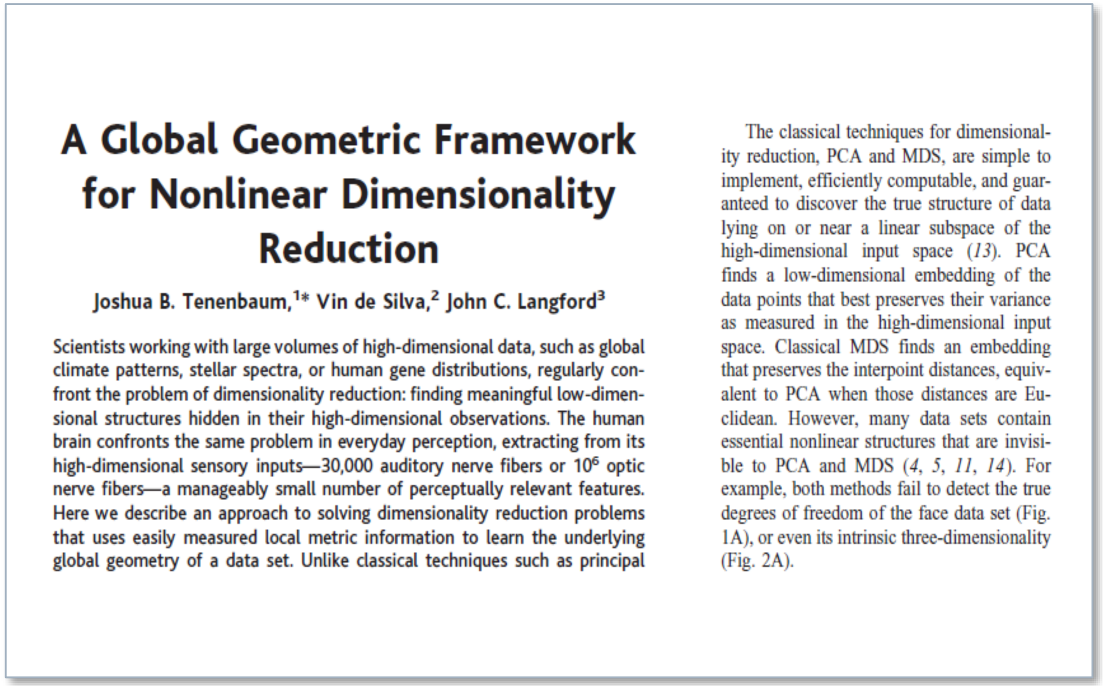

A Global Geometric Framework for Nonlinear Dimensionality Reduction

Joshua B. Tenenbaum,[1*] Vin de Silva,[2] John C. Langford[3]

Scientists working with large volumes of high-dimensional data, such as global climate patterns, stellar spectra, or human gene distributions, regularly confront the problem of dimensionality reduction: finding meaningful low-dimensional structures hidden in their high-dimensional observations. The human brain confronts the same problem in everyday perception, extracting from its high-dimensional sensory inputs—30,000 auditory nerve fibers or 10^6 optic nerve fibers—a manageably small number of perceptually relevant features. Here we describe an approach to solving dimensionality reduction problems that uses easily measured local metric information to learn the underlying global geometry of a data set. Unlike classical techniques such as principal

The classical techniques for dimensionality reduction, PCA and MDS, are simple to implement, efficiently computable, and guaranteed to discover the true structure of data lying on or near a linear subspace of the high-dimensional input space (*13*). PCA finds a low-dimensional embedding of the data points that best preserves their variance as measured in the high-dimensional input space. Classical MDS finds an embedding that preserves the interpoint distances, equivalent to PCA when those distances are Euclidean. However, many data sets contain essential nonlinear structures that are invisible to PCA and MDS (*4*, *5*, *11*, *14*). For example, both methods fail to detect the true degrees of freedom of the face data set (Fig. 1A), or even its intrinsic three-dimensionality (Fig. 2A).

Isomap（等距特征映射）

Isomap (Isometric Feature Mapping)

Sam Roweis and Lawrence Saul.

"Nonlinear Dimensionality Reduction by Locally Linear Embedding".
SCIENCE, Vol. 290, Dec. 2000.

"通过局部线性嵌入进行非线性降维"

"Here, we introduce locally linear embedding (LLE), an unsupervised learning algorithm that computes low-dimensional, neighborhood-preserving embeddings of high-dimensional inputs."

"这里，我们提出局部线性嵌入（LLE），一种计算高维输入数据中低维、邻域保护嵌入的非监督学习算法。"

Nonlinear Dimensionality Reduction by Locally Linear Embedding

Sam T. Roweis[1] and Lawrence K. Saul[2]

Many areas of science depend on exploratory data analysis and visualization. The need to analyze large amounts of multivariate data raises the fundamental problem of dimensionality reduction: how to discover compact representations of high-dimensional data. Here, we introduce locally linear embedding (LLE), an unsupervised learning algorithm that computes low-dimensional, neighborhood-preserving embeddings of high-dimensional inputs. Unlike clustering methods for local dimensionality reduction, LLE maps its inputs into a single global coordinate system of lower dimensionality, and its optimizations do not involve local minima. By exploiting the local symmetries of linear reconstructions, LLE is able to learn the global structure of nonlinear manifolds, such as those generated by images of faces or documents of text.

along shortest paths confined to the manifold of observed inputs. Here, we take a different approach, called locally linear embedding (LLE), that eliminates the need to estimate pairwise distances between widely separated data points. Unlike previous methods, LLE recovers global nonlinear structure from locally linear fits.

The LLE algorithm, summarized in Fig. 2, is based on simple geometric intuitions. Suppose the data consist of N real-valued vectors $\vec{X}_i$, each of dimensionality D, sampled from some underlying manifold. Provided there is sufficient data (such that the manifold is well-sampled), we expect each data point and its neighbors to lie on or close to a locally linear patch of the manifold. We characterize the local geometry of these patches by linear coefficients that reconstruct each data point from its neighbors. Reconstruction errors are measured

Geoffrey Hinton and Ruslan Salakhutdinov.

"Reducing the Dimensionality of Data with Neural Networks".
SCIENCE, Vol. 313, Jul. 2006.

"利用神经元网络降低数据的维度"

"We describe an effective way of initializing the weights that allows deep autoencoder networks to learn low-dimensional codes that work much better than principal components analysis as a tool to reduce the dimensionality of data."

"我们描述一种初始化权重的有效方法，可让深度自编码网络学习低维代码，作为一种降低数据维度的工具，远远好于主成分分析方法。"

materials are identical for all configurations. The blue bars in Fig. 1 summarize the measured SHG signals. For excitation of the *LC* resonance in Fig. 1A (horizontal incident polarization), we find an SHG signal that is 500 times above the noise level. As expected for SHG, this signal closely scales with the square of the incident power (Fig. 2A). The polarization of the SHG emission is nearly vertical (Fig. 2B). The small angle with respect to the vertical is due to deviations from perfect mirror symmetry of the SRRs (see electron micrographs in Fig. 1). Small detuning of the *LC* resonance toward smaller wavelength (i.e., to 1.3-μm wavelength) reduces the SHG signal strength from 100% to 20%. For excitation of the Mie resonance with vertical incident polarization in Fig. 1D, we find a small signal just above the noise level. For excitation of the Mie resonance with horizontal incident polarization in Fig. 1C, a small but significant SHG emission is found, which is again po-

Reducing the Dimensionality of Data with Neural Networks

G. E. Hinton* and R. R. Salakhutdinov

High-dimensional data can be converted to low-dimensional codes by training a multilayer neural network with a small central layer to reconstruct high-dimensional input vectors. Gradient descent can be used for fine-tuning the weights in such "autoencoder" networks, but this works well only if the initial weights are close to a good solution. We describe an effective way of initializing the weights that allows deep autoencoder networks to learn low-dimensional codes that work much better than principal components analysis as a tool to reduce the dimensionality of data.

Dimensionality reduction facilitates the classification, visualization, communication, and storage of high-dimensional data. A simple and widely used method is principal components analysis (PCA), which finds the directions of greatest variance in the data set and represents each data point by its coordinates along each of these directions. We describe a nonlinear generalization of PCA that uses an adaptive, multilayer "encoder" network

504 28 JULY 2006 VOL 313 SCIENCE www.sciencemag.org

code: http://www.cs.toronto.edu/~hinton/MatlabForSciencePaper.html

Alex Rodriguez and Alessandro Laio.

"Clustering by fast search and find of density peaks".
SCIENCE, Vol. 344, Jun. 2014.

"通过快速查找和发现密度峰值进行聚类"

"We propose an approach based on the idea that cluster centers are characterized by a higher density than their neighbors and by a relatively large distance from points with higher densities."

"我们提出一种基于如下思想的方法：聚类中心点具有密度高于相邻点、距离相对大于次高密度点的特性。"

MACHINE LEARNING

Clustering by fast search and find of density peaks

Alex Rodriguez and Alessandro Laio

Cluster analysis is aimed at classifying elements into categories on the basis of their similarity. Its applications range from astronomy to bioinformatics, bibliometrics, and pattern recognition. We propose an approach based on the idea that cluster centers are characterized by a higher density than their neighbors and by a relatively large distance from points with higher densities. This idea forms the basis of a clustering procedure in which the number of clusters arises intuitively, outliers are automatically spotted and excluded from the analysis, and clusters are recognized regardless of their shape and of the dimensionality of the space in which they are embedded. We demonstrate the power of the algorithm on several test cases.

Brenden Lake, Ruslan Salakhutdinov, Joshua Tenenbaum.

“Human-level concept learning through probabilistic program induction”.
SCIENCE, Vol. 350, Dec. 2015.

“凭借概率规划归纳法进行人类层级的概念学习”

“We see the one-shot learning capacities studied here as a challenge for these neural models: one we expect they might rise to by incorporating the principles of compositionality, causality, and learning to learn that BPL instantiates.”

“我们看到本文研究的一次性学习能力作为对那些神经模型的一种挑战：通过将组合性、因果性和学会学习BPL实例化的原则相结合，成为一个我们期待它们会崛起的方向。”

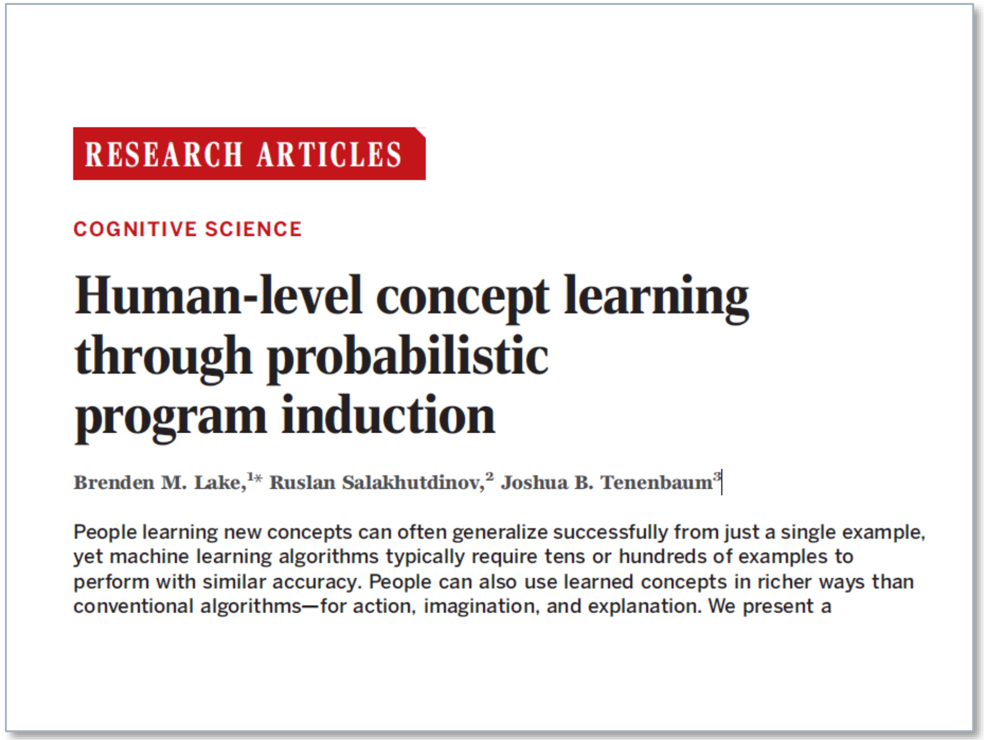

RESEARCH ARTICLES

COGNITIVE SCIENCE

Human-level concept learning through probabilistic program induction

Brenden M. Lake,[1*] Ruslan Salakhutdinov,[2] Joshua B. Tenenbaum[3]

People learning new concepts can often generalize successfully from just a single example, yet machine learning algorithms typically require tens or hundreds of examples to perform with similar accuracy. People can also use learned concepts in richer ways than conventional algorithms—for action, imagination, and explanation. We present a

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, *et al.*

"Human-level control through deep reinforcement learning".
NATURE, Vol. 518, Feb. 2015.

"凭借深度强化学习达到人类水平的操控"

"Here we use recent advances in training deep neural networks to develop a novel artificial agent, termed a deep Q-network, that can learn successful policies directly from high-dimensional sensory inputs using end-to-end reinforcement learning."

"这里我们采用训练深度网络的最新进展开发一种新颖的人造智能体，称为深度Q网络，应用端到端的强化学习，能直接从高维感知输入中学习成功的策略。"

Human-level control through deep reinforcement learning

Volodymyr Mnih[1]*, Koray Kavukcuoglu[1]*, David Silver[1]*, Andrei A. Rusu[1], Joel Veness[1], Marc G. Bellemare[1], Alex Graves[1], Martin Riedmiller[1], Andreas K. Fidjeland[1], Georg Ostrovski[1], Stig Petersen[1], Charles Beattie[1], Amir Sadik[1], Ioannis Antonoglou[1], Helen King[1], Dharshan Kumaran[1], Daan Wierstra[1], Shane Legg[1] & Demis Hassabis[1]

The theory of reinforcement learning provides a normative account[1], deeply rooted in psychological[2] and neuroscientific[3] perspectives on animal behaviour, of how agents may optimize their control of an environment. To use reinforcement learning successfully in situations approaching real-world complexity, however, agents are confronted with a difficult task: they must derive efficient representations of the environment from high-dimensional sensory inputs, and use these to generalize past experience to new situations. Remarkably, humans and other animals seem to solve this problem through a harmonious combination of reinforcement learning and hierarchical sensory processing systems[4,5], the former evidenced by a wealth of neural data revealing notable parallels between the phasic signals emitted by dopa-

agent is to select actions in a fashion that maximizes cumulative future reward. More formally, we use a deep convolutional neural network to approximate the optimal action-value function

$$Q^*(s,a) = \max_\pi \mathbb{E}\left[r_t + \gamma r_{t+1} + \gamma^2 r_{t+2} + \ldots \middle| s_t = s, a_t = a, \pi\right],$$

which is the maximum sum of rewards r_t discounted by γ at each time-step t, achievable by a behaviour policy $\pi = P(a|s)$, after making an observation (s) and taking an action (a) (see Methods)[19].

Reinforcement learning is known to be unstable or even to diverge when a nonlinear function approximator such as a neural network is used to represent the action-value (also known as Q) function[20]. This instability has several causes: the correlations present in the sequence

Y. LeCun, Y. Bengio and G. Hinton.

"Deep learning". *NATURE*, Vol. 521, May. 2015.

"深度学习"

"Deep learning discovers intricate structure in large data sets by using the backpropagation algorithm to indicate how a machine should change its internal parameters that are used to compute the representation in each layer from the representation in the previous layer. "

"深度学习通过利用反向传播算法发现大型数据集中复杂的结构表明，一台机器如何改变其内部参数被用于从前一层表征中计算出每层的表征。"

Deep learning

Yann LeCun[1,2], Yoshua Bengio[3] & Geoffrey Hinton[4,5]

Deep learning allows computational models that are composed of multiple processing layers to learn representations of data with multiple levels of abstraction. These methods have dramatically improved the state-of-the-art in speech recognition, visual object recognition, object detection and many other domains such as drug discovery and genomics. Deep learning discovers intricate structure in large data sets by using the backpropagation algorithm to indicate how a machine should change its internal parameters that are used to compute the representation in each layer from the representation in the previous layer. Deep convolutional nets have brought about breakthroughs in processing images, video, speech and audio, whereas recurrent nets have shone light on sequential data such as text and speech.

Machine-learning technology powers many aspects of modern society: from web searches to content filtering on social networks to recommendations on e-commerce websites, and it is increasingly present in consumer products such as cameras and smartphones. Machine-learning systems are used to identify objects in images, transcribe speech into text, match news items, posts or intricate structures in high-dimensional data and is therefore applicable to many domains of science, business and government. In addition to beating records in image recognition[1–4] and speech recognition[5–7], it has beaten other machine-learning techniques at predicting the activity of potential drug molecules[8], analysing particle accelerator data[9,10], reconstructing brain circuits[11], and predicting the effects of mutations

David Silver, Aja Huang, Chris Maddison, *et al.*

"Mastering the game of Go with deep neural networks and tree search".
NATURE, Vol. 529, Jan. 2016.

"利用深度神经网络和树搜索征服围棋游戏"

Here we introduce a new approach to computer Go that uses 'value networks' to evaluate board positions and 'policy networks' to select moves. … Without any look ahead search, the neural networks play Go at the level of state-of-the-art Monte Carlo tree search programs that simulate thousands of random games of self-play.

"我们在此提出一种计算机围棋的新方法，使用 '价值网络' 评价棋盘位置、使用 '策略网络' 选择走子。… 没有任何前向搜索，该神经网络以先进水平的蒙特卡洛树搜索程序博弈围棋，模拟成千上万次随机自我对弈。"

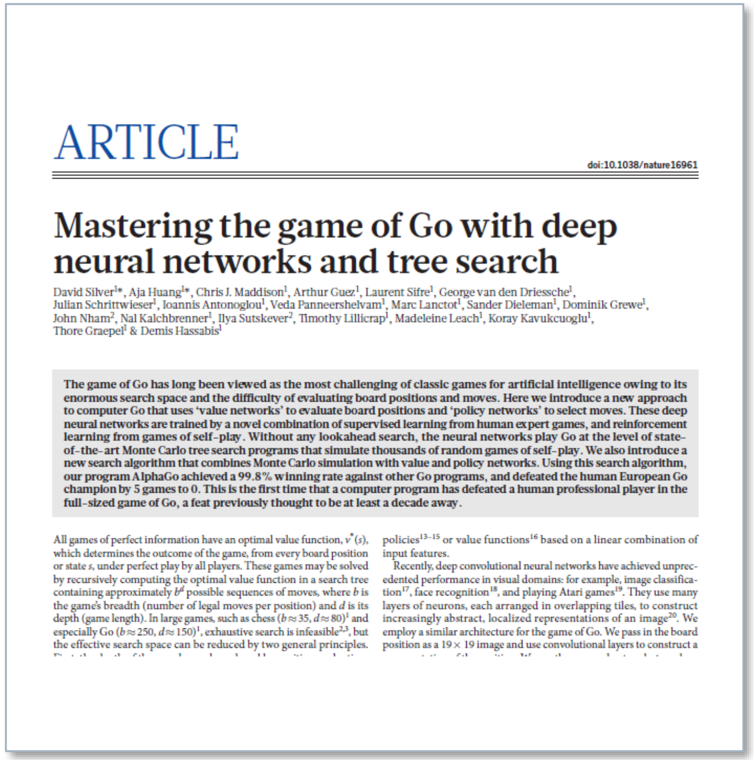
ARTICLE

doi:10.1038/nature16961

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

David Silver[1*], Aja Huang[1*], Chris J. Maddison[1], Arthur Guez[1], Laurent Sifre[1], George van den Driessche[1], Julian Schrittwieser[1], Ioannis Antonoglou[1], Veda Panneershelvam[1], Marc Lanctot[1], Sander Dieleman[1], Dominik Grewe[1], John Nham[2], Nal Kalchbrenner[1], Ilya Sutskever[2], Timothy Lillicrap[1], Madeleine Leach[1], Koray Kavukcuoglu[1], Thore Graepel[1] & Demis Hassabis[1]

The game of Go has long been viewed as the most challenging of classic games for artificial intelligence owing to its enormous search space and the difficulty of evaluating board positions and moves. Here we introduce a new approach to computer Go that uses 'value networks' to evaluate board positions and 'policy networks' to select moves. These deep neural networks are trained by a novel combination of supervised learning from human expert games, and reinforcement learning from games of self-play. Without any lookahead search, the neural networks play Go at the level of state-of-the-art Monte Carlo tree search programs that simulate thousands of random games of self-play. We also introduce a new search algorithm that combines Monte Carlo simulation with value and policy networks. Using this search algorithm, our program AlphaGo achieved a 99.8% winning rate against other Go programs, and defeated the human European Go champion by 5 games to 0. This is the first time that a computer program has defeated a human professional player in the full-sized game of Go, a feat previously thought to be at least a decade away.

All games of perfect information have an optimal value function, $v^*(s)$, which determines the outcome of the game, from every board position or state s, under perfect play by all players. These games may be solved by recursively computing the optimal value function in a search tree containing approximately b^d possible sequences of moves, where b is the game's breadth (number of legal moves per position) and d is its depth (game length). In large games, such as chess ($b \approx 35$, $d \approx 80$)[1] and especially Go ($b \approx 250$, $d \approx 150$)[1], exhaustive search is infeasible[2,3], but the effective search space can be reduced by two general principles.

policies[13–15] or value functions[16] based on a linear combination of input features.

Recently, deep convolutional neural networks have achieved unprecedented performance in visual domains: for example, image classification[17], face recognition[18], and playing Atari games[19]. They use many layers of neurons, each arranged in overlapping tiles, to construct increasingly abstract, localized representations of an image[20]. We employ a similar architecture for the game of Go. We pass in the board position as a 19×19 image and use convolutional layers to construct a

Research Area of AI

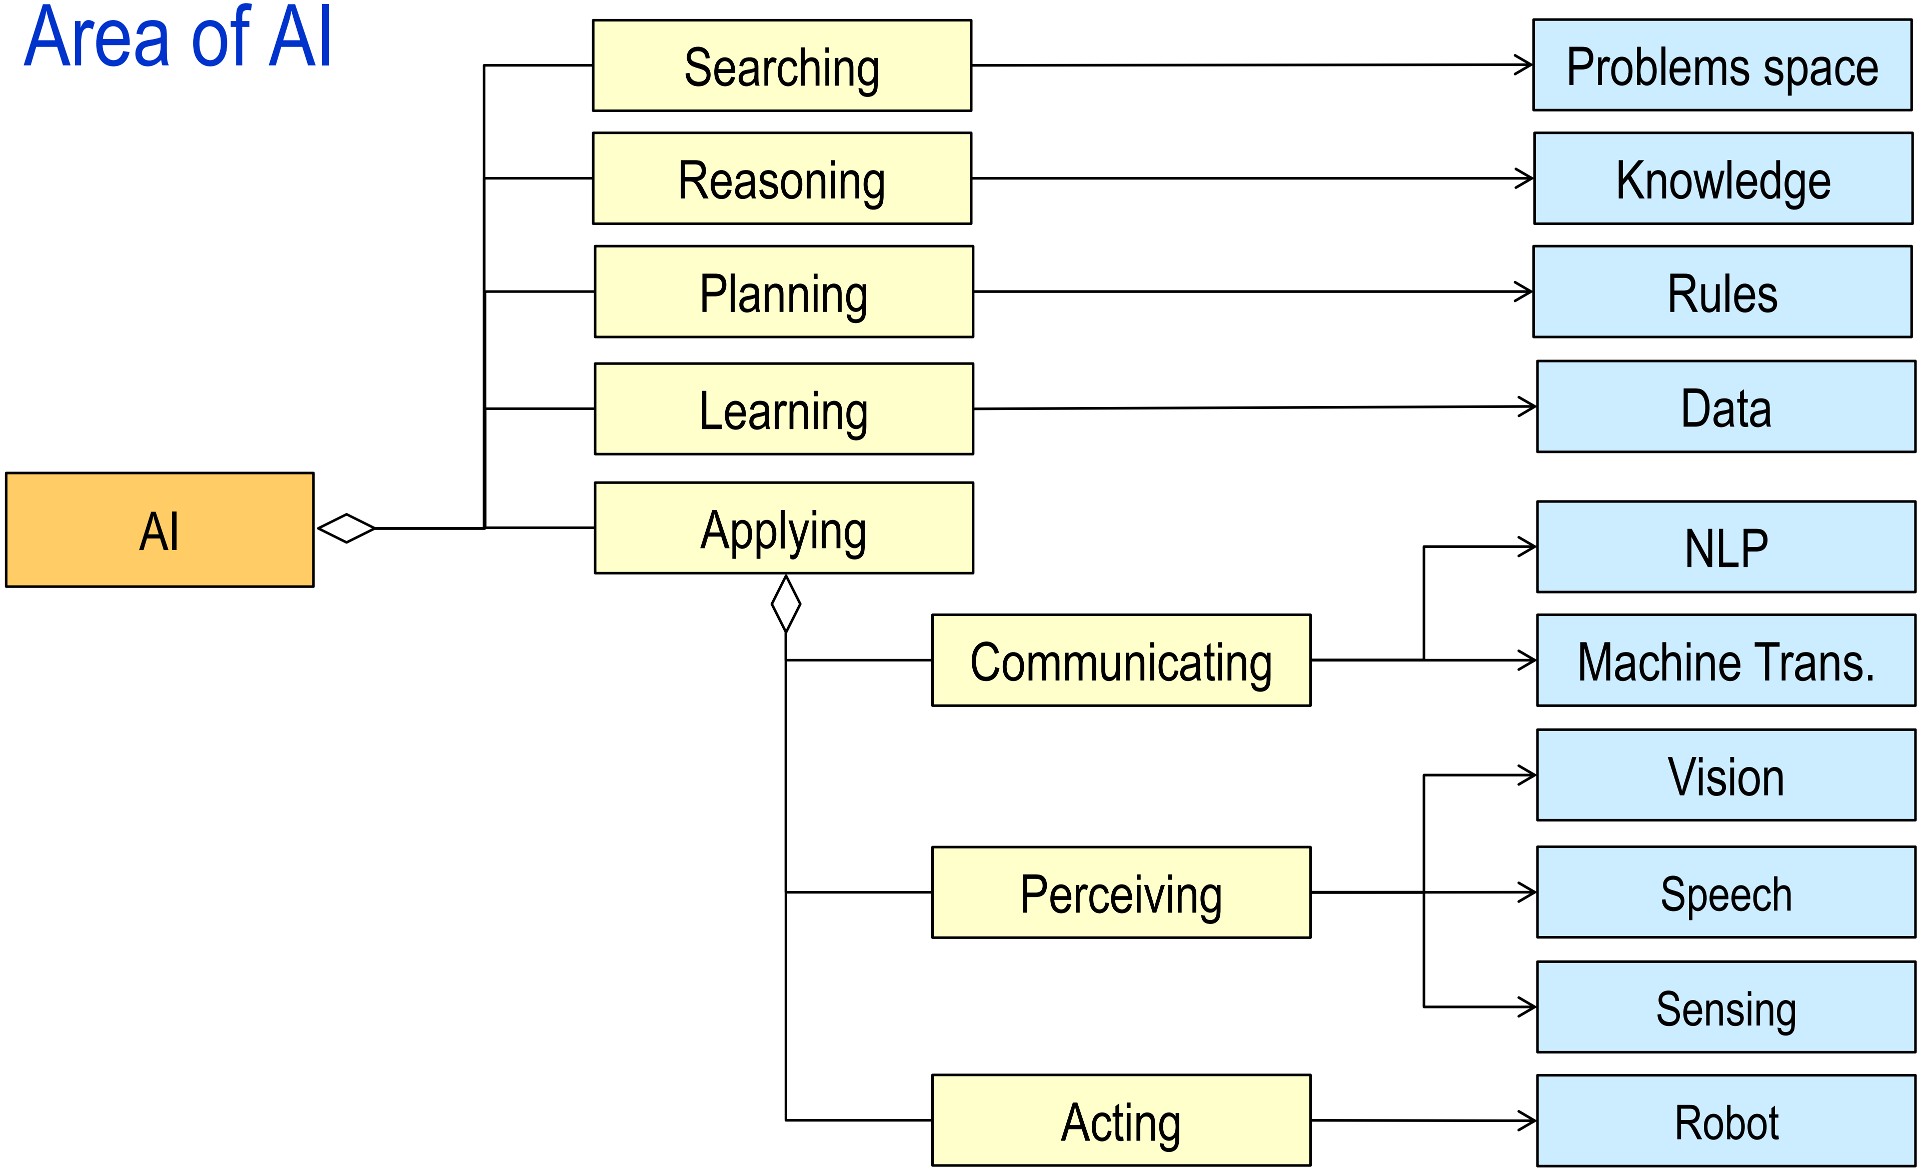

This Course

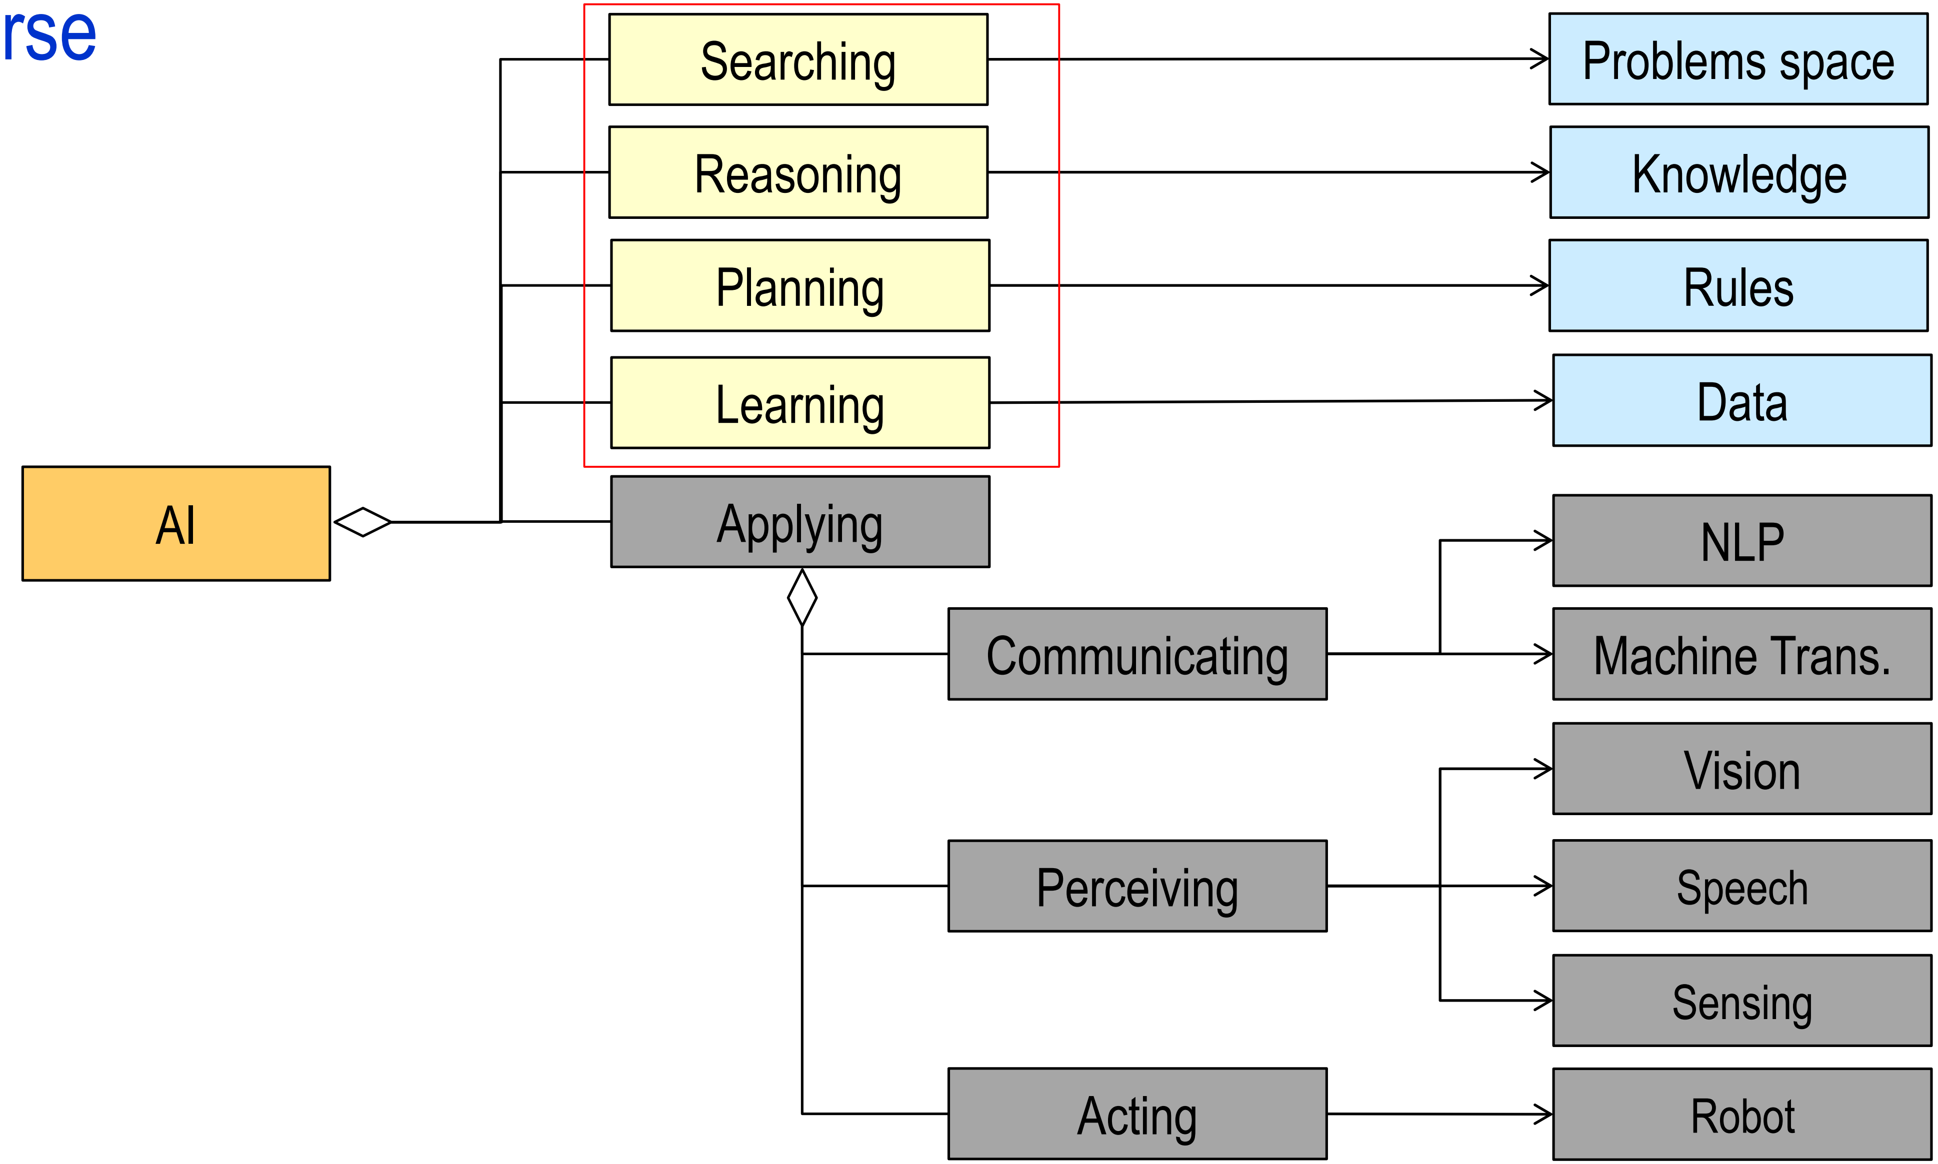
Searching
Reasoning
Planning
Learning
Problems space
Knowledge
Rules
Data
AI
Applying
Communicating
Perceiving
Acting
NLP
Machine Trans.
Vision
Speech
Sensing
Robot

What does it feel like to stand here?

- It seems like a pretty intense place to be standing — but then you have to remember something about what it's like to stand on a time graph: you can't see what's to your right.

  站在看起来好像是令人非常紧张的地方——然后你要记住站在时间曲线上是什么感觉：你看不到你的右侧是什么。

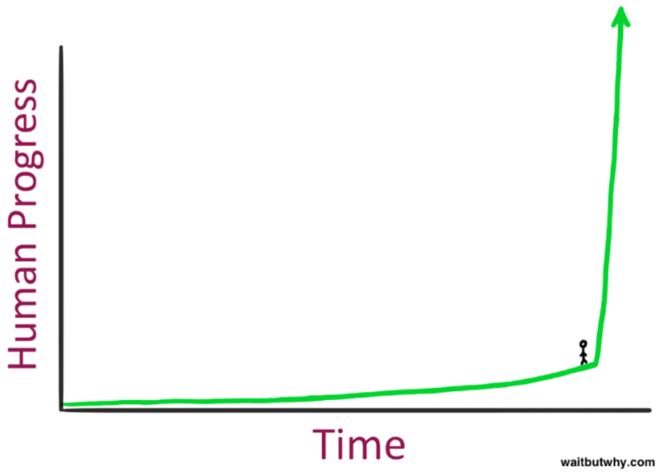

- So here's how it actually feels to stand there: which probably feels pretty normal…

  而这里是要站立的实际感觉如何的地方：大概感觉相当平常…

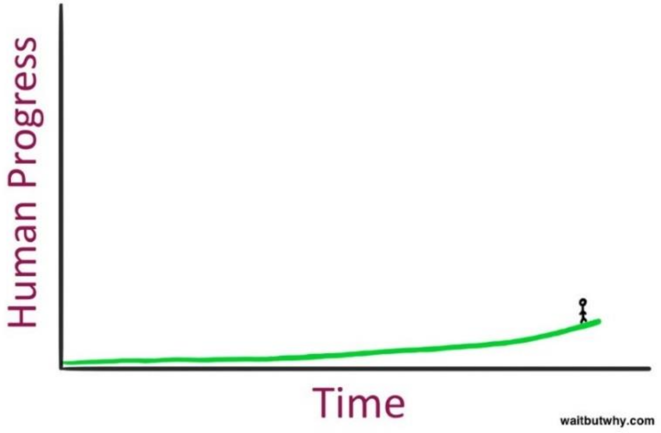

Source: http://waitbutwhy.com/

Thank you for your attention!

PoAI