

Chapter 1

What is Artificial Intelligence

Text A

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

For years, it was thought that computers would never be more powerful than the human brain, but as development has *accelerated* in modern times, this has proven to be not the case.

AI as a concept refers to computing hardware being able to *essentially* think for itself, and make decisions based on the data it is being fed. AI systems are often hugely complex and powerful, with the ability to process *unfathomable* depths of information in an extremely quick time in order to come to an effective conclusion.

Thanks to detailed *algorithms*, AI systems are now able to perform *mammoth* computing tasks much faster and more efficiently than human minds, helping making big *strides* in research and development areas around the world.

Some of the most notable real-world applications of AI are **IBM's Watson¹**, which is being used to power research in a huge

New Words and Expressions

artificial /ɑ:tɪ'fɪʃ(ə)l/ adj.

人造的; 仿造的

demonstrate /'demənstreɪt/ vt.

证明; 展示; 论证

agent /'eɪdʒənt/ n.

代理; 主体

perceive /pə'si:v/ vt.

察觉, 感觉; 理解; 认知

maximize /'mæksɪmaɪz/ vt.

取...最大值; 对...极为重视

mimic /'mɪmɪk/ vt.

模仿, 摹拟

cognitive /'kɒɡnɪtɪv/ adj.

认知的, 认识的

associate /ə'səʊʃieɪt/ vt.

使联合; 使发生联系

accelerate /ək'seləreɪt/ vi.

加速; 促进; 增加

essentially /ɪ'senʃ(ə)li/ adv.

本质上; 本来

unfathomable /ʌn'fæð(ə)məb(ə)l/ adj.

深不可测的; 无底的; 高深莫测的

algorithm /'ælgərɪð(ə)m/ n.

【计】【数】 算法, 运算法则

mammoth /'mæməθ/ adj.

巨大的, 庞大的

stride /straɪd/ n.

进展

range of fields, with Microsoft's **Azure**² Machine Learning and **TensorFlow**³ also making headlines around the world.

But AI-powered smart assistants are becoming a common presence on mobile devices too, with the likes of Siri, Cortana and Alexa all being welcomed into many people's lives.

There seems no limit to the applications of AI technologies, and perhaps the most exciting aspect of the ecosystem is that there's no telling where it can go next, and what problems it may *ultimately* be able to solve.

Why is Artificial Intelligence Important?

AI *automates* repetitive learning and discovery through data. But AI is different from hardware-driven, robotic automation. Instead of automating manual tasks, AI performs frequent, *high-volume*, computerized tasks *reliably* and without fatigue. For this type of automation, human inquiry is still essential to set up the system and ask the right questions.

AI adds intelligence to existing products. In most cases, AI will not be sold as an individual application. Rather, products you already use will be improved with AI capabilities, much like Siri was added as a feature to a new generation of Apple products. Automation, conversational platforms, *bots* and smart machines can be combined with large amounts of data to improve many technologies at home and in the workplace, from security intelligence to investment analysis.

AI adapts through progressive learning algorithms to let the data do the programming. AI finds structure and *regularities* in data so that the algorithm acquires a skill: The algorithm becomes a classifier or a predictor. So, just as the algorithm can teach itself how to play chess, it can teach itself what product to recommend next online. And the models adapt when given new data. **Back propagation**⁴ is an AI technique that allows the model to adjust, through training and added data, when the first answer is not quite right.

AI analyzes more and deeper data using *neural* networks that have many hidden layers. Building a **fraud detection**⁵ system with five hidden layers was almost impossible a few years ago. All that has changed with incredible computer power and big data. You need lots of data to train deep learning models because they learn directly from the data. The more data you can feed them, the more accurate they become.

New Words and Expressions

ultimately /'ʌltɪmətli/ adv.

最后; 根本; 基本上

automate /'ɔ:təmet/ vt.

使自动化, 使自动操作

high-volume adj.

大容量

reliably /ri'laɪəbli/ adv.

可靠地; 确实地

bots /bɒts/ n.

机器人

regularity /regjʊ'lærəti/ n.

规律性; 整齐; 正规

propagation /ˌprɒpə'geɪʃən/ n.

传播; 繁殖; 增殖

neural /'njʊərə(ə)l/ adj.

神经的; 神经系统的

AI achieves incredible accuracy through deep neural networks—which was previously impossible. For example, your interactions with Alexa, Google Search and Google Photos are all based on deep learning—and they keep getting more accurate the more we use them. In the medical field, AI techniques from deep learning, image classification and object recognition can now be used to find cancer on *MRIs* with the same accuracy as highly trained *radiologists*.

AI gets the most out of data. When algorithms are self-learning, the data itself can become intellectual property. The answers are in the data; you just have to apply AI to get them out. Since the role of the data is now more important than ever before, it can create a competitive advantage. If you have the best data in a competitive industry, even if everyone is applying similar techniques, the best data will win.

New Words and Expressions

MRI abbr.

(Magnetic Resonance Imaging)

核磁共振成像

radiologist /ˈredɪˈɒlədʒɪst/ n.

放射科医生

Terms

1. IBM Watson

IBM Watson 是认知计算系统的杰出代表，也是一个技术平台。认知计算代表一种全新的计算模式，它包含信息分析、自然语言处理和机器学习领域的大量技术创新，能够助力决策者从大量非结构化数据中提取和分析重要信息，提升洞察能力。

2. Azure

Azure 机器学习服务是一项云服务，可以使用它来训练、部署、自动执行以及管理机器学习模型，所有这些都是在云提供的广泛范围内进行的。

3. TensorFlow

TensorFlow 是一个基于数据流编程 (dataflow programming) 的符号数学系统，被广泛应用于各类机器学习 (machine learning) 算法的编程实现，其前身是谷歌的神经网络算法库 DistBelief。TensorFlow 拥有多层级结构，可部署于各类服务器、PC 终端和网页并支持 GPU 和 TPU 高性能数值计算，被广泛应用于谷歌内部的产品开发和各领域的科学研究。

4. Back Propagation

Back Propagation (反向传播算法) 是目前用来训练人工神经网络 (Artificial Neural Network, ANN) 的最常用且最有效的算法。其主要思想是：

- (1) 将训练集数据输入 ANN 的输入层，经过隐藏层，最后达到输出层并输出结果，这是 ANN 的前向传播过程；
- (2) 由于 ANN 的输出结果与实际结果有误差，则计算估计值与实际值之间的误差，将被从输出层向隐藏层反向传播，直至传播到输入层；
- (3) 在反向传播的过程中，根据误差调整各种参数的值；不断迭代上述过程，直至收敛。

5. Fraud Detection

Fraud Detection 为反欺诈中所用到的机器学习模型。

反欺诈应用的机器模型算法，多为二分类算法：

(1) 梯度提升决策树（Gradient Boosting Decision Tree, GBDT）算法，该算法的性能高，且在各类数据挖掘中应用广泛，表现优秀，应用场景较多。

(2) Logistic 回归又称 Logistic 回归分析，是一种广义的线性回归分析模型，常用于数据挖掘、疾病自动诊断、经济预测等领域，在有标注样本下训练模型对不同的欺诈情况进行二元判别。

(3) 非监督的异常检测的方法，主要是从数据中找出异常的点，这些异常往往与欺诈有关联。

Comprehension

Blank Filling

1. Artificial Intelligence (AI), sometimes called _____, is intelligence demonstrated by machines, in contrast to the _____ intelligence displayed by humans and other animals.
2. AI research is defined as the study of “intelligent agents”: any device that perceives its _____ and takes actions that _____ its chance of successfully achieving its goals.
3. The term “artificial intelligence” is applied when a machine mimics “_____” functions that humans associate with other _____, such as “learning” and “problem solving”.
4. AI as a concept refers to computing _____ being able to essentially think for itself, and make decisions based on the _____ it is being fed.
5. Thanks to detailed _____, AI systems are now able to perform mammoth _____ tasks much faster and more efficiently than human minds.
6. AI-powered smart assistants are becoming a common presence on _____.
7. AI automates repetitive learning and discovery through _____.
8. AI adapts through _____ to let the data do the programming.
9. AI finds _____ and _____ in data so that the algorithm acquires a skill.

Content Questions

1. What is artificial intelligence?
2. In computer science AI research is defined as the study of “intelligent agents”. What does “intelligent agents” refer to?
3. How AI is applied in the real world?
4. What is “back propagation”?

Answers

Blank Filling

1. machine intelligence; natural

2. environment; maximize
3. cognitive; human minds
4. hardware; data
5. algorithms; computing
6. mobile devices
7. data
8. progressive learning algorithms
9. structure; regularities

Content Questions

1. Artificial Intelligence (AI), sometimes called machine intelligence, is intelligence demonstrated by machines, in contrast to the natural intelligence displayed by humans and other animals.
2. It refers to any device that perceives its environment and takes actions that maximize its chance of successfully achieving its goals.
3. Some of the most notable real-world applications of AI are IBM's Watson, which is being used to power research in a huge range of fields, with Microsoft's Azure Machine Learning and TensorFlow also making headlines around the world.
4. Back propagation is an AI technique that allows the model to adjust, through training and added data, when the first answer is not quite right.

参考译文 A

人工智能 (AI)，有时被称为机器智能。与人类和其他动物展示的天生智能不同，人工智能是由机器展示的智能。在计算机科学中，人工智能研究被定义为对“智能主体”的研究：能够感知环境，并采取行动以成功实现目标的智能机器。一般来说，“人工智能”一词是指机器模仿人类的“认知”功能和“智能”行为，如“学习”和“解决问题”等。

多年来，人们一直认为计算机永远不会比人脑强大，但随着现代科技的发展，事实证明并非如此。

人工智能的概念是指计算硬件能够独立思考，并根据输入的数据做出决策。人工智能系统通常非常复杂和强大，能够在极短的时间内处理深不可测的信息，以便得出有效的结论。

得益于周密的算法，人工智能系统现在能够比人脑更快、更有效地执行庞大的计算任务，从而有助于世界各地的研究和开发取得重大进展。

人工智能最引人注目的一些应用有：IBM 的沃森 (Watson) 被用于许多领域的研究，微软的 Azure 机器学习 (Azure Machine Learning) 和 TensorFlow 也都引起了大家的关注。

但是，随着 Siri、Cortana 和 Alexa 等智能助手进入人们的生活，人们的移动设备也应用了人工智能。

人工智能技术的应用似乎没有限制，这正是它令人兴奋的地方，人们不知道它下一步将走向何方，也不知道最终可能解决什么问题。

为什么人工智能至关重要?

人工智能通过渐进式学习算法进行调整,让数据进行编程。人工智能在数据中发现结构和规律,使算法获得一种技能:算法成为分类器或预测器。因此,就像算法可以自学如何下棋一样,它也可以自学下一步应该推荐的产品类型。当得到新的数据时,模型就会适应。反向传播是一种人工智能技术,当第一个答案不太正确时,它允许模型通过训练和添加数据来调整。

人工智能通过数据来实现自动化的重复学习和发现。但人工智能不同于硬件驱动的机器人自动化。人工智能不再自动完成手工任务,而是可靠地、不知疲倦地执行频繁的、大批量的计算机任务。对于这种类型的自动化,人工查询对于建立系统和提出正确的问题仍然是必不可少的。

人工智能提高了现有产品的智能化。在大多数情况下,人工智能不会作为一个单独的应用程序出售,而是使在用产品得到改进,就像 Siri 被添加到新一代苹果产品中一样。自动化、对话平台、机器人和智能机器可以与大量数据相结合,使安全智能、投资分析等许多家庭和工作场所的技术得以改进。

人工智能使用具有许多隐藏层的神经网络分析更多和更深层次的数据。多年前,建立一个包含五个隐藏层的欺诈检测系统几乎是不可能的。这一切都因计算机的强大功能和大数据而改变。人们需要大量的数据来训练深度学习模型,因为它们直接从数据中学习。你能提供给它们的数据越多,它们就越准确。

人工智能通过深层神经网络实现了难以置信的准确性,这在以前是不可能的。例如,你与 Alexa、谷歌搜索和谷歌照片的互动都是基于深度学习的,我们越使用它们,它们就越准确。在医学领域,来自深度学习、图像分类和目标识别的人工智能技术,现在可以像训练有素的放射科医生一样,在核磁共振成像上准确地发现癌症。

人工智能得益于数据。当算法是自我学习时,数据本身可以成为知识产权,答案就在数据中。人们只需要应用人工智能就可以把它们计算出来。现在数据的作用比以往任何时候都重要,它可以创造竞争优势。如果你在竞争激烈的行业中拥有最好的数据,即使每个人都在应用类似的技术,谁拥有最好的数据,谁就会胜出。

Text B

Artificial Intelligence is everywhere, from Apple's iPhone keyboard to Zillow's home price estimates. There's also a lot of stuff out there that marketers are calling AI, but really isn't.

Perhaps things reached a new high point last month when AlphaGo, a virtual player of the ancient Chinese board game Go developed by Alphabet's DeepMind AI research group, *trounced* the top human player in the world, China's Ke Jie.

A moment of drama *encapsulates* the achievement: After Jie *resigned* in the second of three matches, the 19-year-old lingered in his chair, staring down at the board for several minutes, *fidgiting*

New Words and Expressions

trounce/traʊns/vt.

痛打; 严责; 打败

encapsulate /ɪn'kæpsjʊleɪt/ vt.

将...封进内部; 概述

resign /rɪ'zain/ vi.

辞职; 放弃

fidget /'fidʒɪt/vi.

烦躁; 坐立不安

with game pieces and *scratching* his head. Aja Huang, the DeepMind senior research scientist who was tasked with moving game pieces on behalf of AlphaGo, eventually got up from his chair and walked offstage, leaving Jie alone for a moment.

Still, it's generally true that a human being like Jie has more brainpower than a computer. That's because a person can perform a wide range of tasks better than machines, while a given computer program enhanced with AI like AlphaGo might be able to *edge out* a person at just a few things.

But the prospect of AI becoming smarter than people at most tasks is the single biggest thing that drives debates about effects on employment, creativity and even human existence.

Here's an overview of what AI really is, and what the biggest companies are doing with it.

So what is AI, really?

Given that everybody's talking about AI now, you would think it's new. But the underlying techniques are not. The field got its start in the mid-twentieth century, and one of its most popular methods came about in the 1980s.

AI first took hold in the 1950s. While some of its underlying concepts predate it, the term itself dates to 1956, when John McCarthy, a math professor at Dartmouth College, proposed a summer research project based on the idea that "every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to *simulate* it."

In the next few years AI research labs popped up at the Massachusetts Institute of Technology (MIT) and Stanford University. Research touched on computer chess, robotics and natural-language communication.

Interest in the field *fluctuated* over time. AI winters came in the 1970s and 1980s as public interest *waned* and outside funding dried up. Startups boasting promising capabilities and venture capital backing in the mid-1980s abruptly disappeared, as John Markoff detailed in his 2015 book "Machines of Loving Grace."

There are several other terms you often hear in connection to AI.

Machine learning generally *entails* teaching a machine how to do a particular thing, like recognizing a number, by feeding it a bunch of data and then directing it to make predictions on new data.

New Words and Expressions

scratch /skrætʃ/ vt.

抓; 刮

edge out

替代; 微微胜过

simulate /'sɪmjʊleɪt/ vt.

模仿

fluctuate /'flʌktʃueɪt/ vt.

使波动; 使动摇

wane /weɪn/ vi.

衰落; 变小

entail /en'teɪl/ vt.

使需要, 必需; 承担

The big deal about machine learning now is that it's getting easier to invent software that can learn over time and get smarter as it accumulates more and more data. Machine learning often requires people to hand-engineer certain features for the machine to look for, which can be complex and time-consuming.

Deep learning is one type of machine learning that demands less hand-engineering of features. Often the approach involves artificial neural networks, a mathematical system loosely inspired by the way neurons work together in the human brain. Neuroscientist Warren McCulloch and mathematician Walter Pitts came up with the first such system in 1943. Through the years, researchers advanced the concept with various techniques, including adding multiple layers. With each *successive* layer, higher-level features could be detected in the original data to make a better prediction. The layers pick out features in the data themselves. But using more layers demands more computing power.

Why is it suddenly so hot?

Through the years, hardware has gotten more powerful, and chipmakers including Nvidia have refined their products to better suit AI computations. Larger data sets in many domains have become available to train models more *extensively* as well.

In 2012, Google made headlines when it trained a neural network with 16,000 central processing unit (CPU) chips on 10 million images from YouTube videos and taught it to recognize cats. But later that year, the world of image recognition was rocked when an eight-layer neural network trained on two graphics processing units (GPUs) outdid all others in a competition to accurately classify images based on their content. Months later, Google acquired DNNresearch, the University of Toronto team behind the breakthrough.

Since then, AI activity has only accelerated, with the world's top technology companies leading the way.

Meanwhile, the world's most valuable companies — technology companies! — continue to publish research on their latest gains, which only adds to the *fascination*.

Google and its parent company Alphabet have made several AI. Beyond that, perhaps in a few decades, an AI system with superhuman capabilities in most domains — sometimes referred to as artificial general intelligence — could emerge. Depending on whom you ask, that could be either very good or very bad. In an extreme

New Words and Expressions

successive /sək'sesɪv/ adj.

连续的; 继承的; 依次

extensively /ɪk'stensɪvli/ adv.

广阔地; 广大地

fascination /fæsi'neiʃ(ə)n/ n.

魅力; 魔力; 入迷

case, an AGI system could end up making humans *extinct*. But if things go well, perhaps AGI will be something that will *supercharge* humans and help them live longer. The prospect of either of these two *scenarios* is perhaps what draws so much attention to AI development today, and what has inspired so much science fiction in the past.

But for now, what people generally see is narrow AI — intelligence applied to a small number of domains — and it doesn't always work the way it should. Look at Alexa, Cortana, the Google Assistant or Siri — they misunderstand spoken words all the time.

The thing is, the biggest companies in the world are now investing in AI like never before. And that trend is not about to let up.

New Words and Expressions

extinct/ɪk'stɪŋkt/ adj.

灭绝的，绝种的；熄灭的

supercharge/'su:pətʃɑ:dʒ/ vt.

对…增压；使…超负荷

scenarios/sɪ'neiəriəz/ n.

情节；脚本

参考译文 B

从苹果的 iPhone 键盘到 Zillow 的房价估算，人工智能无处不在。市场上也有很多营销人员称之为人工智能的东西，但实际上并非如此。

上个月，由 Alphabet 旗下 DeepMind AI 研究集团开发的围棋虚拟玩家 AlphaGo 击败了来自中国的世界顶级棋手柯洁，让人工智能达到了一个新的高潮。

在这背后有一个戏剧性的时刻：在三场比赛的第二场比赛中，19 岁的柯洁几乎崩溃。他坐在椅子上，低头盯着棋盘看了几分钟，一边摆弄棋子，一边挠头。负责替 AlphaGo 移动棋子的 DeepMind 高级研究科学家阿贾·黄最终从椅子上站起来，走下赛场，让柯洁独自待了一会儿。

尽管如此，像柯洁这样的人确实比计算机更聪明。因为人类可以比机器更好地执行各种各样的任务，而具有人工智能的计算机程序，例如 AlphaGo，只能在某些事情上胜过人类。

但是，在一些任务中人工智能表现得比人类聪明，因此引发了关于就业、创造力乃至人类生存影响的争论。

以下是关于人工智能的概述，以及大公司正在用它做什么。

那么人工智能到底是什么呢？

因为现在每个人都在谈论人工智能，你可能会认为它是新的。但它的基本技术并不新。这个领域在 20 世纪中期开始发展，最流行的方法之一出现在 20 世纪 80 年代。

人工智能的首次提出是在 20 世纪 50 年代。虽然它的一些基础概念在此之前就已经出现，但这个词本身可以追溯到 1956 年，由达特茅斯学院的数学教授约翰·麦卡锡（John McCarthy）在一个暑期研究项目中提出，该项目基于这样一个理念：学习的方方面面以及智能的其他特征都可以精确描述，从而能够进行机器仿真。

在接下来的几年里，麻省理工学院（MIT）和斯坦福大学（Stanford University）成立了人工智能研究实验室，研究涉及计算机象棋、机器人和自然语言交流。

人们对这一领域的兴趣随着时间而发生了变化。随着公众兴趣的减退和外部资金的枯竭，20 世纪 70 年代和 80 年代成了人工智能的寒冬。正如约翰·马尔科夫（John Markoff）在 2015 年出版的 *Machines of Loving Grace* 一书中所描述的那样，在 20 世纪 80 年代中期，那些拥有良好能力和风险资本支持的初创企业突然消失了。

下面是几个人们经常听到的与人工智能有关的术语。

机器学习通常需要教会机器如何做一件特定的事情，例如识别一个数字，给它输入一组数据，然后引导它对新数据进行预测。

现在机器学习的重要之处在于，随着时间的推移，积累越来越多的数据，开发出能够学习，以及变得更智能的软件，这变得越来越容易。机器学习通常需要人们手工设计某些特性，以便机器寻找，这可能是复杂和耗时的。

深度学习是一种机器学习，它对手动操作要求较少。这种方法通常涉及人工神经网络，这是一种数学系统，其灵感来自于人类大脑中神经元协同工作的方式，神经学家沃伦·麦卡洛克和数学家沃尔特·皮茨在 1943 年第一次提出。多年来，研究人员通过各种技术发展了这一概念，其中包括多层感知技术。对于每一个连续的层，可以在原始数据中检测到更高层次的特征，从而做出更好地预测。这些层在数据中挑选，但是使用更多的层需要更强的计算能力。

人工智能为什么突然备受关注？

这些年，硬件变得越来越强大，包括英伟达（Nvidia）在内的芯片制造商已经改进了产品，以更好地适应人工智能计算。在许多领域中，更大的数据集也可用于广泛的训练模型。

2012 年，谷歌用 16,000 个中央处理器（CPU）芯片对 YouTube 视频中的 1000 万张图片进行神经网络训练，并教会它识别猫，这让它成为头条新闻。但同年晚些时候，在一场根据图像内容进行精确分类的竞赛中，使用两个图形单元（GPU）训练的八层神经网络击败了其他对手，震惊了图像识别领域。几个月后，谷歌收购了取得这一突破性进展的多伦多大学 DNNresearch 团队。

从那以后，人工智能发展更为迅速，全球顶尖科技公司走在了前面。同时，世界上最有价值的科技公司继续发表它们的最新研究成果，这让人对人工智能更加着迷。

谷歌及其母公司 Alphabet 已经开发了数个人工智能系统，或许在几十年内，一个在大多数领域都具备超人能力的人工智能系统将会出现。这可能是好的，也可能是坏的。在极端情况下，AGI 系统可能最终导致人类灭绝。但如果一切进展顺利，AGI 系统或许能使人能力更强，或帮助人们活得更久。这两种情况中任何一种，都引起了人们对人工智能的关注，也赋予了科幻小说许多灵感。

就目前而言，人们通常看到的是狭隘的人工智能，应用于少数领域的智能，它并不总是以它应该的方式工作。看看 Alexa、Cortana、谷歌助理或 Siri，它们总是误解别人说的话。

世界上的许多大公司正在前所未有地投资人工智能，这种趋势不会停止。