

论文推荐

清华大学戴琼海院士：从脑科学到人工智能

临菲信息技术港

清华大学戴琼海院士团队在 *Engineering* 发表综述论文：“From Brain Science to Artificial Intelligence”，第一作者：Jingtao Fan，通信作者：Qionghai Dai（戴琼海）。（J. Fan, L. Fang, J. Wu et al., From Brain Science to Artificial Intelligence, *Engineering*, <https://doi.org/10.1016/j.eng.2019.11.012>）

模拟人类的大脑，是人工智能的一个重要研究方向。而对于脑科学的认知，是人工智能模拟人脑的基础。该文综述了当前脑科学与人工智能的研究成果，讨论了脑科学适应人工智能的机遇和挑战。

为了在脑科学与 AI 研究之间建立起一座桥梁，首先需要通过研究新型脑成像技术来探索脑科学的秘密，建立大脑的动态连接图谱以及将神经科学实验与理论、模型和统计学相结合等，从而建立起从机器感知和机器学习到机器思维和机器决策的颠覆性模型和工作模式。

脑科学与 AI 的联系

脑科学与 AI 之间联系紧密，许多 AI 的先驱科学家也是脑科学家，脑科学的研究成果揭示了与 AI 原理相关的一些重要问题，并进一步使 AI 在理论和技术上取得了重大突破，现在的深度学习更是直接受到了脑科学的启发，越来越多的脑科学研究成果可以激发新的深度学习模型。

AI 在一定程度上受脑科学的启发，研究 AI 的一种简单方法是将 AI 与脑科学及其相关领域相结合，许多 AI 研究的先驱做出的贡献也归功于他们在脑科学方面的扎实背景，包括现代计算机、Hebbian 学习算法、生物大脑的工作记忆功能、连续学习、强化学习等。

人类大脑的研究

许多国家和地区都开展了大脑研究项目以加速脑科学研究，共同目标是基于脑科学的发现进行下一代 AI 开发。例如，光遗传学能够以单细胞分辨率精确地控制神经活动，先进的波束

调制方法可以进一步实现大规模控制,大规模神经活动可以用多种方式以 3D 形式记录; BRAIN Initiative 旨在通过神经科学革新机器学习。通过这些研究,我们能够以前所未有的时空分辨率系统地观察神经活动,更好地了解人类的大脑,启发下一代 AI 的设计。

仪器的地位和作用

仪器是脑科学和 AI 之间的桥梁,利用仪器对大脑进行的观测对 AI 的出现和发展贡献巨大。当前,人类对大脑的理解得益于神经技术的进步,用于高质量成像采集的新技术和新仪器的研发有望让我们了解大脑记录、处理、应用、存储和检索大量信息的过程;神经科学研究中,非侵入性的光学方法是更好的选择,空间分辨率、时间分辨率以及灵敏度高,能够从神经元个体、神经活动及其相互作用中获取动态和静态信息,把对神经系统的分析从亚细胞水平扩展到整个大脑,利用光学方法通过光遗传学方法可以控制高时空分辨率下的神经活动。开发具备大视场和高时空分辨率的技术和仪器,以及实现大视场、高时空分辨率和大景深将成为未来十年微观成像技术发展面临的巨大挑战。最后,对微观结构维度的探索可能会产生一种新型的神经计算单元,显微成像仪器的发展目标是实现从像素到体速、从静态到动态的更宽、更高、更快和更深的成像,弥补 AI 与人类智能之间的差距。

附：原文



临菲信息技术港



临菲信息技术港公众号



临菲学堂



Contents lists available at ScienceDirect

Engineering

journal homepage: www.elsevier.com/locate/eng

Research
Artificial Intelligence—Review

From Brain Science to Artificial Intelligence

Jingtao Fan^a, Lu Fang^b, Jiamin Wu^a, Yuchen Guo^a, Qionghai Dai^{a,*}

^a Department of Automation, Tsinghua University, Beijing 100084, China

^b Tsinghua-Berkeley Shenzhen Institute, Tsinghua University, Shenzhen 518055, China

ARTICLE INFO

Article history:

Received 11 July 2019

Revised 22 October 2019

Accepted 4 November 2019

Keywords:

Artificial intelligence
Brain science

ABSTRACT

Reviewing the history of the development of artificial intelligence (AI) clearly reveals that brain science has resulted in breakthroughs in AI, such as deep learning. At present, although the developmental trend in AI and its applications has surpassed expectations, an insurmountable gap remains between AI and human intelligence. It is urgent to establish a bridge between brain science and AI research, including a link from brain science to AI, and a connection from knowing the brain to simulating the brain. The first steps toward this goal are to explore the secrets of brain science by studying new brain-imaging technology; to establish a dynamic connection diagram of the brain; and to integrate neuroscience experiments with theory, models, and statistics. Based on these steps, a new generation of AI theory and methods can be studied, and a subversive model and working mode from machine perception and learning to machine thinking and decision-making can be established. This article discusses the opportunities and challenges of adapting brain science to AI.

© 2020 THE AUTHORS. Published by Elsevier LTD on behalf of Chinese Academy of Engineering and Higher Education Press Limited Company. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

The history of artificial intelligence (AI) clearly reveals the connections between brain science and AI. Many pioneer AI scientists are also brain scientists. The neural connections in the human brain that were discovered using microscopes inspired the artificial neural network [7]. The brain's convolution property and multi-layer structure, which were discovered using electronic detectors, inspired the convolutional neural network and deep learning [10,11]. The attention mechanism that was discovered using a positron emission tomography (PET) imaging system inspired the attention module [15]. The working memory that was discovered from functional magnetic resonance imaging (fMRI) results inspired the memory module in machine learning models that led to the development of long short-term memory (LSTM) [25]. The changes in the spine that occur during learning, which were discovered using two-photon imaging systems, inspired the elastic weight consolidation (EWC) model for continual learning [31]. Although the AI community and the brain science community currently appear to be unconnected, the results from brain science reveal important issues related to the principles of intelligence, which lead to significant theoretical and technological break-

throughs in AI. We are now in the deep learning era, which was directly inspired by brain science. It can be seen that the increasing research findings in brain science can inspire new deep learning modes. Furthermore, the next breakthrough in AI is likely to come from brain science.

2. Artificial intelligence inspired by brain science

The goal of AI is to investigate theories and develop computer systems that are able to conduct tasks that require biological or human intelligence, with functions such as perceptrons, recognition, decision-making, and control [1]. Conversely, the goal of brain science, which is also termed neuroscience, is to study the structures, functions, and operating mechanisms of biological brains, such as how the brain processes information, makes decisions, and interacts with the environment [2]. It is easy to see that AI can be regarded as the simulation of brain intelligence. Therefore, a straightforward way to develop AI is to combine it with brain science and related fields, such as cognition science and psychology. In fact, many pioneers of AI, such as Alan Turing [3], Marvin Minsky [4], John McCarthy [5], and Geoffrey Hinton [6], were interested in both fields and contributed a great deal to AI thanks to their solid backgrounds in brain science.

Research on AI began directly after the emergence of modern computers, with the goal of building intelligent “thinking”

* Corresponding author.

E-mail address: qh dai@mail.tsinghua.edu.cn (Q. Dai).

<https://doi.org/10.1016/j.eng.2019.11.012>

2095-8099/© 2020 THE AUTHORS. Published by Elsevier LTD on behalf of Chinese Academy of Engineering and Higher Education Press Limited Company. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

machines. Since the birth of AI, there have been interactions between it and brain science. At the beginning of the 20th century, researchers were able to observe the connections between neurons in the neural system, including brains, due to the development of microscopy. Inspired by the connections between neurons, computer scientists developed the artificial neural network, which is one of the earliest and most successful models in the history of AI. In 1949, Hebbian learning was proposed [7]. This is one of the oldest learning algorithms. Hebbian learning was directly inspired by the dynamics of biological neural systems. In particular, based on the observation that a synapse between two neurons is strengthened when the neurons on either side of the synapse (input and output) have highly correlated outputs, the Hebbian learning algorithm increases the connection weight between two neurons if they are highly correlated. After this development, artificial neural networks received considerable research attention from researchers. A representative work was the perceptron [8], which directly modeled the information storage and organization in the brain. The perceptron is a single-layer artificial neural network with a multidimensional input, which laid the foundation for the multilayer network.

In 1959, Hubel and Wiesel—the recipients of the 1981 Nobel Prize in Physiology or Medicine—utilized electronic signal detectors to capture the responses of neurons when a visual system saw different images [9]. The single-cell recordings from the mammalian visual cortex revealed how visual inputs are filtered and pooled in simple and complex cells in the V1 area. This research demonstrated that the visual processing system in the brain conducted convolutional operations and had a multilayered structure. It indicated that biological systems utilized successive layers with nonlinear computations to transform raw visual inputs into an increasingly complex set of features, thereby making the vision system invariant to the transformations, such as pose and scale, in the visual inputs during the recognition task. These observations directly inspired the convolutional neural network [10,11], which was the fundamental model for the recent, ground-breaking deep learning technique [12]. Another key component of artificial neural networks and deep learning is the back-propagation algorithm [13], which addresses the problem of how to tune the parameters or weights in a network. Interestingly, the basic idea of back propagation was first proposed in the 1980s by neuroscientists and cognitive scientists [14], rather than by computer scientists or machine learning researchers. The scientists observed that the microstructures of neural systems and the neural system of the biological brain were gradually tuned using a learning procedure with the purpose of minimizing the error and maximizing the reward of the output. The attention mechanism was first introduced in the 1890s as a psychological concept, and was designed such that an intelligent agent selectively concentrated on certain important parts of the information—instead of concentrating on all of the information—in order to improve the cognition process [15]. In the 1990s, studies began using new medical imaging technologies, such as PET, to investigate the attention mechanism in the brain. In 1999, PET was utilized to study selective attention in the brain [16]. Then, using other imaging technologies, researchers discovered more about the attention mechanism in a biological brain [17]. Inspired by the attention mechanism in a biological brain, AI researchers began incorporating attention modules into artificial neural networks in temporal [18] or spatial [19] ways, which improved the performance of deep neural networks for natural language processing and computer vision, respectively. With an attention module, the network is able to selectively focus on important objects or words and ignore irrelevant ones, thereby making the training and inferential processes more efficient than those of a conventional deep network.

A machine learning model usually forgets the information in the data that it has processed, whereas biological intelligence is able to maintain such information for a period of time. It is believed that there is working memory in a biological brain that remembers past data. The concept of working memory was first introduced in the 1970s and was summarized from cognition experiments [20,21]. Since 1990, researchers have used PET and fMRI to study the working memory in biological brains, and have found that the prefrontal cortex in the brain is a key part [22–24]. Inspired by the working memory research from brain science, AI researchers have attempted to incorporate a memory module into machine learning models. One representative method is LSTM [25], which laid the foundation for many sequential processing tasks, such as natural language processing, video understanding, and time-series analysis. A recent study also showed that with a working memory module, a model can perform complicated reasoning and inference tasks, such as finding the shortest path between specific points and inferring the missing links in randomly generated graphs [26]. By remembering previous knowledge, it is also possible to perform one-shot learning, which requires just a few labeled samples to learn a new concept [27].

Continual learning is a basic skill in biological intelligence that is used to learn a new task without forgetting previous ones. How a biological neural system learns multiple tasks at different times is a challenging research topic. In 1990, the two-photon microscopy technique [28] made it possible to observe the *in vivo* structures and functions of dendritic spines during learning at the spatial scale of single synapses [29]. With this imaging system, researchers in the 2010s studied neocortical plasticity in the brain during continual learning. The results revealed how neural systems remember previous tasks when learning new tasks by controlling the growth of neurons [30]. Inspired by the observation of biological neural systems, a learning algorithm termed EWC was proposed for deep neural networks. This algorithm controlled the changes in the network parameters when learning a new task, such that older knowledge was preserved, thereby making continual learning in deep learning possible [31].

Reinforcement learning (RL) is a widely used machine learning framework that has been utilized in many applications, such as AlphaGo. It relates to how AI agents take action and interact with the environment. In fact, RL is also strongly related to the biological learning process [32]. One important RL method—which was also one of the earliest methods—is temporal-difference learning (TDL). TDL learns by bootstrapping from the current estimate of the value function. This strategy is similar to the concept of second-order conditioning that has been observed in animal systems [33].

3. Brain projects

Many countries and regions have conducted projects to accelerate brain science research, as shown in Table 1 [34–39]. Despite different emphases and routes, the development of the next generation of AI based on discoveries in brain science is a common objective of all brain research projects. Governments and most scientists seem to have reached a consensus that advancing neural imaging and manipulating techniques can help us explore the working principles of the brain, which will allow us to design a better AI architecture, including both hardware and software. During such studies, mutual collaboration between multiple disciplines including biology, physics, informatics, and chemistry are necessary to enable new discoveries in different aspects.

During the past five years, important achievements in brain research have been made with the support of brain research projects. The development of optogenetics has made it possible to

Table 1
Overview of brain science research projects around the world.

Project	Timeline	Funding amount (billion USD)	Main objectives
United States NIH BRAIN Initiative [34]	2013–2025	~4.5	<ul style="list-style-type: none"> • Discovering diversity: Identify different brain cell types and determine their roles in health and disease. • Maps at multiple scales: Generate circuit diagrams with varying resolutions from the synapses to the whole brain. • The brain in action: Produce a dynamic picture of the functioning brain through large-scale monitoring of neural activity. • Demonstrating causality: Link brain activity to behavior with precise interventional tools that change neural circuit dynamics. • Identifying fundamental principles: Produce conceptual foundations for understanding mental processes by developing new theoretical and analytical tools. • Advancing human neuroscience: Develop innovative technologies to understand the human brain and treat its disorders, and create and support human brain research networks. • From BRAIN Initiative to the brain: Apply new technological/conceptual approaches to discover how neural activity patterns transform into cognitions, emotions, perceptions, and actions.
Korea Korea Brain Initiative [35]	2018–2027	>1.2	<ul style="list-style-type: none"> • Decipher the brain functions and mechanisms that mediate the integration and control of brain functions underlying decision-making. • Map a functional connectome with searchable, multidimensional, and information-integrated features. • Develop novel technologies and neuro-tools for integrated brain mapping. • Enable socioeconomic ramifications that not only facilitate global collaboration in the neuroscience community, but also develop various brain science-related industrial and medical innovations.
Europe Human Brain Project (HBP) [36]	2013–2023	>1	<ul style="list-style-type: none"> • Develop a scientific infrastructure for brain research and cognitive neuroscience. • Gather and disseminate data describing the brain and related diseases. • Simulate the brain. • Build theories and models of the brain. • Develop brain-inspired computing, data analytics, and robots.
Japan Brain/MINDS [37]	2014–2024	>0.3	<ul style="list-style-type: none"> • Use the marmoset, a small primate with a short life cycle, for functional and structural brain mapping and genetic studies. • Develop innovative tools to monitor and manipulate different aspects of neural activity. • Establish biomarkers for brain disorders.
Canada Brain Canada [38]	2006–Unknown	>0.24	<ul style="list-style-type: none"> • Understand the brain in health and illness, improve lives, and achieve societal impact. • Increase the scale and scope of funding to accelerate the pace of Canadian brain research. • Create a collective commitment to brain research across the public, private, and voluntary sectors. • Deliver transformative, original, and outstanding research programs.
Australia Australian Brain Initiative [39]	2016–2026	>0.2	<ul style="list-style-type: none"> • Health: Develop new treatments by revealing the mechanisms of brain abnormalities in neuropsychiatric diseases. • Education: Help improve brain growth by coding the cognitive functions of neural circuits and brain networks. • New industry: Develop new drugs, medical devices, and wearable technologies by promoting the combination of industrial collaborators and brain research.

manipulate neural activities at a single-cell resolution [40]. Large-scale manipulation can be further accomplished using advanced beam-modulation techniques [41,42]. In the meantime, various methods have been proposed to record large-scale neural activities in three dimensions (3D) [43–45]. The number of neurons that can be recorded at the same time has increased rapidly from tens to thousands, and may be increased to millions in the near future with the increasing technological developments in wide-field high-resolution imaging [46–48]. Such significant improvements in the field of neurophotonics provide a basis for important discoveries in neuroscience [49,50]. For example, the emphasis in the BRAIN Initiative will be gradually moved to discovery-driven science.

One typical case in the BRAIN Initiative, which aims to revolutionize machine learning through neuroscience, is MICrONS (machine intelligence from cortical networks). With serial-section electron microscopy, complicated neural structures can be reconstructed in 3D at unprecedented resolutions [51]. In combination with high-throughput data analysis techniques for multiscale data [52,53], novel scientific questions can be developed to explore fundamental neuroscience problems [54]. With this improved understanding, researchers have proposed novel architectures for deep neural networks, and have tried to understand the working principles of current architectures [55,56]. In addition, the current deep learning techniques can help to accelerate the massive amount of data processing that is necessary in such research, thus forming a virtuous circle.

Thanks to technological developments in recent years, it is now possible to observe neural activities in a systematic view at unprecedented spatial-temporal resolutions. Many large-scale data analysis techniques have been proposed in the meantime to solve the challenges that result from the massive amount of data produced by such technologies. Following this route, various brain projects can exponentially accelerate brain research. By achieving an increasing number of discoveries, we can develop a better picture of the human brain. There is no doubt that the working principles of the brain will inspire the design of the next generation of AI, just as past discoveries in brain research have inspired today's AI achievements.

4. Instrumental bridges between brain science and artificial intelligence

Instrumental observations of the brain have made enormous contributions to the emergence and advancement of AI. Modern neurobiology started from the information acquisition of microstructures across the subcellular to tissue levels, and benefited from the inventions of microscopy and the biased staining of substances in cells and tissues. The renowned neuroanatomist Santiago Ramón y Cajal was the first to use Golgi staining to observe a large number of tissue specimens of the nervous system, and put forward the fundamental theories on neurons and neural signal transduction. Cajal and Golgi shared the Nobel Prize in

Physiology or Medicine in 1906. Cajal is now widely known as the father of modern neurobiology.

Our ever-growing understanding of the human brain has benefited from countless advances in neurotechnology, including the manipulation, processing, and information acquisition of neurons, neural systems, and brains; and cognitive and behavioral learning. Among these advances, the development of new technologies and instruments for high-quality imaging acquisition has been the focus of the past era and is expected to attract the most attention in the future. For example, the BRAIN Initiative, which was launched in the United States in 2013, aims to map dynamic brain images that exhibit the rapid and complex interactions between brain cells and their surrounding nerve circuits, and to unveil the multidimensional intertwined relationships between neural organizations and brain functions. Such advances are also expected to make it possible for us to understand the processes of recording, processing, applying, storing, and retrieving large amounts of information in the brain. In 2017, the BRAIN Initiative sponsored a number of interdisciplinary scientists at Harvard, who undertook to research the understanding of the relationship between neural circuits and behavior, mainly by acquiring and processing large datasets of neural systems under various conditions using high-quality imaging.

Traditional neuroscience research mostly uses electrophysiological methods, such as the use of metal electrodes for nerve excitation and signal acquisition, which have the advantages of high sensitivity and high temporal resolution. However, electrophysiology is invasive and is not suitable for long-term observation. In addition, it has a low spatial resolution and limited expansion ability for the parallel observations that are required to extract the global neural activities at a single neuron resolution of the brain. In contrast, optical methods are noninvasive and have high spatial and temporal resolution and high sensitivity. These methods are capable of acquiring dynamic and static information from individual neurons, nerve activities, and interactions and expanding our analyses of the nervous system from the subcellular level to—potentially—the whole brain. Furthermore, optical methods have been developed as manipulating tools to control nerve activities at high spatial–temporal resolutions by using optogenetics.

It is very urgent to develop technology and instruments with large fields of view and high spatial–temporal resolutions. On the spatial scale, imaging must span from submicron synapses and neurons that are tens of microns in size to brains that are a few millimeters across. On the temporal scale, the rate of frame acquisition should be higher than the response rate of the probing fluorescent proteins that are used. However, due to the intrinsic diffraction limit of optical imaging, there is an inherent contradiction among large fields of view, high resolution, and large depths of view. High-resolution imaging of single neurons or even smaller features usually cannot see brain tissue features that are larger than a few millimeters, and dynamic imaging is often accompanied by higher noise. Live and noninvasive imaging for real-time and long-term acquisition is, however, limited to the superficial layer due to tissue granules that scatter light. How to break through the above bottlenecks and realize a wide field of view, high spatiotemporal resolution, and large depth of view will be the biggest challenge of microscopic imaging in the coming decade.

It is conclusive that exploring from the microstructure dimension may lead to a new type of neurocomputing unit, whereas exploring from the macrostructure dimension in real time may enable an understanding of trans-brain operations and reveal the comprehensive decision-making mechanisms of the brain using multiple information sources (auditory, visual, olfactory, tactile, etc.) in complex environments. The binary ability of the whole brain to explore both the micro- and macro-dimensions in real time will, beyond any doubt, promote the development of the next

generation of AI. Therefore, the developmental goal of a microscopic imaging instrument is to possess broader, higher, faster, and deeper imaging from pixels to voxels and from static to dynamic. Such an instrument could establish a direct link between biological macro-cognitive decision-making and the structure and function of a neural network, lay a foundation for revealing the computational essence of cognition and intelligence, and ultimately promote human self-recognition, thereby filling the research gap between AI and human intelligence.

Acknowledgements

This work is supported by the Consulting Research Project of the Chinese Academy of Engineering (2019-XZ-9), the National Natural Science Foundation of China (61327902), and the Beijing Municipal Science & Technology Commission (Z181100003118014).

Compliance with ethics guidelines

Jingtao Fan, Lu Fang, Jiamin Wu, Yuchen Guo, and Qionghai Dai declare that they have no conflicts of interest or financial conflicts to disclose.

References

- [1] Russell SJ, Norvig P. *Artificial intelligence: a modern approach*. 3rd ed. New York: Pearson Education; 2010.
- [2] Miller GA. The cognitive revolution: a historical perspective. *Trends Cogn Sci* 2003;7(3):141–4.
- [3] Turing A. Computing machinery and intelligence. *Mind* 1950;236:433–60.
- [4] Minsky M, Papert S. *Perceptrons: an introduction to computational geometry*. Cambridge, MA: MIT Press; 1987.
- [5] McCarthy J. *Defending AI research: a collection of essays and reviews*. Stanford: CSLI Publications; 1996.
- [6] Hinton GE, Rumelhart DE, McClelland JL. Distributed representations. In: *Parallel distributed processing: explorations in the microstructure of cognition: foundations*. Cambridge, MA: MIT Press; 1986. p. 77–109.
- [7] Hebb DO. *The organization of behavior*. Hoboken: John Wiley & Sons; 1949.
- [8] Rosenblatt F. The perceptron: a probabilistic model for information storage and organization in the brain. *Psychol Rev* 1958;65(6):386–408.
- [9] Hubel DH, Wiesel TN. Receptive fields of single neurons in the cat's striate cortex. *J Physiol* 1959;148(3):574–91.
- [10] LeCun Y, Boser B, Denker JS, Henderson D, Howard RE, Hubbard W, et al. Backpropagation applied to handwritten zip code recognition. *Neural Comput* 1989;1(4):541–51.
- [11] Krizhevsky A, Sutskever I, Hinton G. ImageNet classification with deep convolutional neural networks. In: Pereira F, Burges CJC, Bottou L, Weinberger KQ, editors. *Proceedings of the Neural Information Processing Systems 2012*; 2012 Dec 3–6; Lake Tahoe, NV, USA. p. 1097–105.
- [12] LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature* 2015;521(7553):436–44.
- [13] Rumelhart DE, McClelland JL. Learning internal representations by error propagation. In: *Parallel distributed processing: explorations in the microstructure of cognition: foundations*. Cambridge, MA: MIT Press; 1986. p. 318–62.
- [14] Rumelhart DE, McClelland JL. *Parallel distributed processing: explorations in the microstructures of cognition: foundations*. Cambridge, MA: MIT Press; 1986.
- [15] James W, Burkhardt F, Bowers F, Skrupskelis IK. *The principles of psychology*. New York: Henry Holt; 1890.
- [16] Raichle ME. Positron emission tomography. In: Wilson RA, Keil LC, editors. *The MIT encyclopedia of the cognitive sciences*. Cambridge, MA: MIT Press; 1999. p. 656–8.
- [17] Scolaro M, Seidl-Rathkopf KN, Kastner S. Functions of the human frontoparietal attention network: evidence from neuroimaging. *Curr Opin Behav Sci* 2015;1:32–9.
- [18] Bahdanau D, Cho K, Bengio Y. Neural machine translation by jointly learning to align and translate. 2014. arXiv:1409.0473.
- [19] Reed S, Zhang Y, Zhang Y, Lee H. Deep visual analogy-making. In: Cortes C, Lawrence ND, Lee DD, Sugiyama M, Garnett R, editors. *Proceedings of the Neural Information Processing Systems 2015*; 2015 Dec 7–12; Montreal, QC, Canada. p. 1252–60.
- [20] Atkinson RC, Shiffrin RM. Human memory: a proposed system and its control processes. In: Spence KW, Spence JT, editors. *Psychology of learning and motivation (volume 2)*. New York: Academic Press; 1968. p. 89–195.
- [21] Baddeley AD, Hitch G. Working memory. In: Bower GH, editor. *Psychology of learning and motivation (volume 8)*. New York: Academic Press; 1974. p. 47–89.

- [22] Goldman-Rakic PS. Cellular and circuit basis of working memory in prefrontal cortex of nonhuman primates. *Prog Brain Res* 1990;85:325–35.
- [23] McCarthy G, Puce A, Constable RT, Krystal JH, Gore JC, Goldman-Rakic P. Activation of human prefrontal cortex during spatial and nonspatial working memory tasks measured by functional MRI. *Cereb Cortex* 1996;6(4):600–11.
- [24] Jonides J, Smith EE, Koeppel RA, Awh E, Minoshima S, Mintun MA. Spatial working memory in humans as revealed by PET. *Nature* 1993;363(6430):623–5.
- [25] Hochreiter S, Schmidhuber J. Long short-term memory. *Neural Comput* 1997;9(8):1735–80.
- [26] Graves A, Wayne G, Reynolds M, Harley T, Danihelka I, Grabska-Barwińska A, et al. Hybrid computing using a neural network with dynamic external memory. *Nature* 2016;538(7626):471–6.
- [27] Santoro A, Bartunov S, Botvinick M, Wierstra D, Lillicrap T. One-shot learning with memory-augmented neural networks. 2016. arXiv:1605.06065.
- [28] Denk W, Strickler JH, Webb WW. Two-photon laser scanning fluorescence microscopy. *Science* 1990;248(4951):73–6.
- [29] Nishiyama J, Yasuda R. Biochemical computation for spine structural plasticity. *Neuron* 2015;87(1):63–75.
- [30] Cichon J, Gan WB. Branch-specific dendritic Ca^{2+} spikes cause persistent synaptic plasticity. *Nature* 2015;520(7546):180–5.
- [31] Kirkpatrick J, Pascanu R, Rabinowitz N, Veness J, Desjardins G, Rusu AA, et al. Overcoming catastrophic forgetting in neural networks. *Proc Natl Acad Sci USA* 2017;114(13):3521–6.
- [32] Sutton R, Barto A. Introduction to reinforcement learning. Cambridge, MA: MIT Press; 1998.
- [33] Sutton RS, Barto AG. Toward a modern theory of adaptive networks: expectation and prediction. *Psychol Rev* 1981;88(2):135–70.
- [34] Insel TR, Landis SC, Collins FS. The NIH BRAIN Initiative. *Science* 2013;340(6133):687–8.
- [35] Jeong S, Lee Y, Jun B, Ryu Y, Sohn J, Kim S, et al. Korea Brain Initiative: emerging issues and institutionalization of neuroethics. *Neuron* 2019;101(3):390–3.
- [36] Amunts K, Ebell C, Müller J, Telefont M, Knoll A, Lippert T. The human brain project: creating a European research infrastructure to decode the human brain. *Neuron* 2016;92(3):574–81.
- [37] Okano H, Sasaki E, Yamamori T, Iriki A, Shimogori T, Yamaguchi Y, et al. Brain/MINDS: a Japanese national brain project for marmoset neuroscience. *Neuron* 2016;92(3):582–90.
- [38] Jabalpurwala I. Brain Canada: one brain one community. *Neuron* 2016;92(3):601–6.
- [39] Alliance A. A neuroethics framework for the Australian Brain Initiative. *Neuron* 2019;101(3):365–9.
- [40] Deisseroth K. Optogenetics. *Nat Methods* 2011;8(1):26–9.
- [41] Pégard NC, Mardinly AR, Oldenburg IA, Sridharan S, Waller L, Adesnik H. Three-dimensional scanless holographic optogenetics with temporal focusing (3D-SHOT). *Nat Commun* 2017;8(1):1228.
- [42] Hochbaum DR, Zhao Y, Farhi SL, Klapoetke N, Werley CA, Kapoor V, et al. All-optical electrophysiology in mammalian neurons using engineered microbial rhodopsins. *Nat Methods* 2014;11(8):825–33.
- [43] Ji N, Freeman J, Smith SL. Technologies for imaging neural activity in large volumes. *Nat Neurosci* 2016;19(9):1154–64.
- [44] Weisenburger S, Vaziri A. A guide to emerging technologies for large-scale and whole-brain optical imaging of neuronal activity. *Annu Rev Neurosci* 2018;41(1):431–52.
- [45] Ahrens MB, Orger MB, Robson DN, Li JM, Keller PJ. Whole-brain functional imaging at cellular resolution using light-sheet microscopy. *Nat Methods* 2013;10(5):413–20.
- [46] Kim TH, Zhang Y, Lecoq J, Jung JC, Li J, Zeng H, et al. Long-term optical access to an estimated one million neurons in the live mouse cortex. *Cell Rep* 2016;17(12):3385–94.
- [47] McConnell G, Trägårdh J, Amor R, Dempster J, Reid E, Amos WB. A novel optical microscope for imaging large embryos and tissue volumes with sub-cellular resolution throughout. *Elife* 2016;5:e18659.
- [48] Stirman JN, Smith IT, Kudenov MW, Smith SL. Wide field-of-view, multi-region, two-photon imaging of neuronal activity in the mammalian brain. *Nat Biotechnol* 2016;34(8):857–62.
- [49] Chen JL, Carta S, Soldado-Magraner J, Schneider BL, Helmchen F. Behaviour-dependent recruitment of long-range projection neurons in somatosensory cortex. *Nature* 2013;499:336–40.
- [50] Sofroniew NJ, Flickinger D, King J, Svoboda K. A large field of view two-photon mesoscope with subcellular resolution for in vivo imaging. *Elife* 2016;5:e14472.
- [51] Joesch M, Mankus D, Yamagata M, Shahbazi A, Schalek R, Suissa-Peleg A, et al. Reconstruction of genetically identified neurons imaged by serial-section electron microscopy. *Elife* 2016;5:e15015.
- [52] Friedrich J, Yang W, Soudry D, Mu Y, Ahrens MB, Yuste R, et al. Multi-scale approaches for high-speed imaging and analysis of large neural populations. *PLoS Comput Biol* 2017;13(8):e1005685.
- [53] Berens P, Freeman J, Deneux T, Chenkov N, McColgan T, Speiser A, et al. Community-based benchmarking improves spike rate inference from two-photon calcium imaging data. *PLoS Comput Biol* 2018;14(5):e1006157.
- [54] Paninski L, Cunningham JP. Neural data science: accelerating the experiment-analysis-theory cycle in large-scale neuroscience. *Curr Opin Neurobiol* 2018;50:232–41.
- [55] Hoffer E, Hubara I, Soudry D. Train longer, generalize better: closing the generalization gap in large batch training of neural networks. In: Guyon I, Luxburg UV, Bengio S, Wallach H, Fergus R, Vishwanathan S, Garnett R, editors. Proceedings of the Neural Information Processing Systems 2017; 2017 Dec 4–9; Long Beach, CA, USA. p. 1731–41.
- [56] Kadmon J, Sompolinsky H. Optimal architectures in a solvable model of deep networks. In: Lee DD, Sugiyama M, Luxburg UV, Guyon I, Garnett R, editors. Proceedings of the Neural Information Processing Systems 2016; 2016 Dec 5–10; Barcelona, Spain. p. 4788–96.