

Toward Next-Generation Artificial Intelligence: Catalyzing the NeuroAI Revolution

By Anthony Zador, Blake Richards, Bence Ölveczky, Sean Escola, Yoshua Bengio, Kwabena Boahen, Matthew Botvinick, Dmitri Chklovskii, Anne Churchland, Claudia Clopath, James DiCarlo, Surya Ganguli, Jeff Hawkins, Konrad Koerding, Alexei Koulakov, Yann LeCun, Timothy Lillicrap, Adam Marblestone, Bruno Olshausen, Alexandre Pouget, Cristina Savin, Terrence Sejnowski, Eero Simoncelli, Sara Solla, David Sussillo, Andreas S. Tolias, Doris Tsao

Abstract: *Neuroscience has long been an important driver of progress in artificial intelligence (AI). We propose that to accelerate progress in AI, we must invest in fundamental research in NeuroAI.*

Over the coming decades, Artificial Intelligence (AI) will transform society and the world economy in ways that are as profound as the computer revolution of the last half century, and likely at an even faster pace. This AI revolution presents tremendous opportunities to unleash human creativity in the modern economy. New developments in AI systems have the potential to enable workers to attain greater productivity and relieve them from performing the most dangerous and menial jobs. But, to reach this potential, we still require advances that will make AI more human-like in its capabilities. Historically, neuroscience has been a key driver and source of inspiration for improvements in AI, particularly those that made AI more proficient in areas that humans and other animals excel at, such as vision, reward-based learning, interacting with the physical world, and language (Hassabis et al. 2017). It can still play this role. To accelerate progress in AI and realize its vast potential, we must invest in fundamental research in “NeuroAI”.

The seeds of the current AI revolution were planted decades ago, largely by researchers attempting to understand how brains compute (McCulloch and Pitts 1943). Indeed, the earliest efforts to build an “artificial brain” led to the invention of the modern “von Neumann computer architecture,” for which John von Neumann explicitly drew upon the very limited knowledge of the brain available to him in the 1940s (Von Neumann 2012). The deep convolutional networks that catalyzed the recent revolution in modern AI are built upon artificial neural networks (ANNs) directly inspired by the Nobel-prize winning work of David Hubel and Torsten Wiesel on visual processing circuits in the cat (Hubel and Wiesel 1962; LeCun and Bengio 1995). Similarly, the development of reinforcement learning (RL) drew a direct line of inspiration from insights into animal behavior and neural activity during learning (Thorndike and Bruce 2017; Rescorla 1972; Schultz, Dayan, and Montague 1997). Now, decades later, applications of ANNs and RL are coming so quickly that many observers assume that the long-elusive goal of human-level intelligence—sometimes referred to as “artificial general intelligence”—is within our grasp. However, in contrast to the optimism of those outside the field, many front-line AI researchers believe that major new breakthroughs are needed before we can build artificial systems capable of doing all that a human, or even a much simpler animal like a mouse, can do.

Although AI systems can easily defeat any human opponent in games such as chess (Campbell, Hoane, and Hsu 2002) and Go (Silver et al. 2016), they are not robust and often struggle when faced with novel

situations. Moreover, we have yet to build systems that can walk to the shelf, take down the chess set, set up the pieces, and move them around during a game. Similarly, no machine can build a nest, forage for berries, or care for young. Today's AI systems cannot compete with the sensorimotor capabilities of a four-year old child, or even simple animals. Many of the basic capacities required to navigate new situations—capacities that animals have or acquire effortlessly—turn out to be deceptively challenging for AI, in part because AI systems lack even the basic abilities to interact with an unpredictable world. A growing number of AI researchers doubt that merely scaling up current approaches will overcome these limitations. Given the need to achieve more natural intelligence in AI, it is quite likely that new inspiration from naturally intelligent systems is needed (Sinz et al. 2019).

While many key AI advances, such as convolutional ANNs and RL were inspired by neuroscience, much of the current research in machine learning is following its own path by building on previously-developed approaches that were inspired by decades old findings in neuroscience, such as attention-based neural networks which were loosely inspired by attention mechanisms in the brain (Itti, Koch, and Niebur 1998; Larochelle and Hinton 2010; Xu et al. 2015). New influences from modern neuroscience exist, but they are spearheaded by a minority of researchers. This represents a missed opportunity. Over the last decades, through efforts such as the NIH BRAIN initiative and others, we have amassed an enormous amount of knowledge about the brain. This has allowed us to learn a great deal about the anatomical and functional structures that underpin natural intelligence. *The emerging field of NeuroAI, at the intersection of neuroscience and AI, is based on the premise that a better understanding of neural computation will reveal basic ingredients of intelligence and catalyze the next revolution in AI, eventually leading to artificial agents with capabilities that match and perhaps even surpass those of humans.* We believe it is the right time for a large-scale effort to identify and understand the principles of biological intelligence, and abstract those for application in computer and robotic systems.

Although it is tempting to focus on the most characteristically human aspects of intelligent behavior, such as abstract thought and reasoning, the basic ingredients of intelligence—adaptability, flexibility, and the ability to make general inferences from sparse observations—are already present in some form in basic sensorimotor circuits which have been evolving for hundreds of millions of years. As AI pioneer Hans Moravec (Moravec 1988) put it, abstract thought “is a new trick, perhaps less than 100 thousand years old...effective only because it is supported by this much older and much more powerful, though usually unconscious, sensorimotor knowledge.” This is good news, because it means that the favored subjects of neuroscience research—rats, mice and non-human primates—can serve as experimentally tractable models of natural intelligence. If AI could match their deceptively simple perceptual and motor abilities, the step to human-level intelligence would be considerably smaller. Thus, we believe that the NeuroAI path will lead to necessary advances if we figure out the core capabilities that all animals possess in embodied sensorimotor interaction with the world.

NeuroAI Grand Challenge: The Embodied Turing Test

In 1950, Alan Turing proposed the “imitation game” (Turing 1950) as a test of a machine's ability to exhibit intelligent behavior equivalent to, or indistinguishable from, that of a human. In that game, now

known as the Turing test, a human judge is asked to evaluate natural language conversations between a real human and a machine trained to mimic human responses. Turing proposed that, in lieu of being able to say concretely whether a machine could “think” (which he considered an impossible question to answer), we could instead determine whether a machine’s conversational abilities were indistinguishable from those of a human’s. Turing proposed that this was a reasonable replacement for the unanswerable question, “can machines think”. Implicit in the Turing test was the belief that language represents the pinnacle of human intelligence, and that a machine capable of conversation must surely be intelligent. To some degree, Turing was right, but in another way he was wrong. While no AI system has passed the Turing test, recent language systems trained purely on large text corpuses can engage in surprisingly cogent conversations. In part their success reveals how easily we can be tricked into imputing intelligence, agency and even consciousness to our interlocutor (Sejnowski 2022). Another is that these systems are still very poor at some reasoning tasks (Kosoy et al. 2022). Impressive though these recent successes are, their failures also serve to highlight that Turing was ignoring the fact that there is far more to intelligence than language ability. Many of the errors that current natural language processing systems make illustrate a fundamental lack of semantics, causal reasoning and common-sense. Words only have meaning for these models by virtue of their statistical co-occurrence, as opposed to their grounding in real-world experiences, so even the most advanced language models, despite their increasing power, continue to struggle with some basic aspects of physical common sense. Thus, the Turing test, as originally formulated, does not probe the ability, shared with animals, to make sense of the physical world in a flexible way. Moreover, this understanding is likely founded on our prodigious perceptual and motor abilities, honed through countless generations of natural selection.

We therefore propose an expanded Turing test, one that includes advanced sensorimotor abilities. The spirit of the original Turing test was to establish a simple qualitative standard against which our progress toward building artificially intelligent machines can be judged. An expanded “embodied Turing test” would benchmark and compare the interactions with the world of artificial systems versus humans and other animals. Because each animal has its own unique set of abilities, each animal defines its own embodied Turing test: An artificial beaver might be tested on its ability to build a dam, and an artificial squirrel on its ability to jump through trees. Nonetheless, many core sensorimotor capabilities are shared by almost all animals, and the ability of animals to rapidly evolve the sensorimotor skills needed to adapt to new environments suggests that these core skills provide a solid foundation. Below we highlight a few of these shared characteristics.

Interacting with the world. The defining feature of animals is their ability to move around and interact with their environment in purposeful ways. Despite recent advances in optimal control, reinforcement learning, and imitation learning, robotics is still far from achieving animal-level abilities in controlling their bodies and manipulating objects, even in simulation. Of course, neuroscience can provide guidance about the kinds of modular and hierarchical architectures that could be adapted to artificial systems to give them these capabilities (Merel, Botvinick, and Wayne 2019). It can also provide us with design principles like partial autonomy (how low-level modules in a hierarchy act semi-autonomously in the absence of input from high-level modules) and amortized control (how movements generated at first by a slow planning process are eventually transferred to a fast reflexive system). Understanding how specific neural circuits participate in different tasks – like locomotion; fine-grained control of limbs, hands and

fingers; perception; and action selection – may provide a path for how such systems could be implemented in robots, and could also inspire solutions for other forms of ‘intelligence’, including in more cognitive realms. For example, we speculate that incorporating principles of circuitry for low-level motor control could help provide a better basis for higher-level motor planning in AI systems.

Flexibility of animal behavior. Another goal is to develop AI systems that can engage a large repertoire of flexible and diverse tasks in a manner that echoes the incredible range of behaviors that individual animals can generate. Modern AI can easily learn to outperform humans at video games like Breakout using nothing more than pixels on a screen and game scores (Mnih et al. 2015). However, these systems, unlike human players, are brittle and highly sensitive to small perturbations: changing the rules of the game slightly, or even a few pixels on the input, can lead to catastrophically poor performance (Huang et al. 2017). This is because these systems learn a mapping from pixels to actions that need not involve an understanding of the agents and objects in the game and the physics that governs them. Similarly, a self-driving car does not inherently know about the danger of a crate falling off a truck in front of it, unless it has literally seen examples of crates falling off trucks leading to bad outcomes. And even if it has been trained on the dangers of falling crates, the system might consider an empty plastic bag being blown out of the car in front of it as an obstacle to avoid at all cost rather than an irritant, again, because it doesn’t actually understand what a plastic bag is or how unthreatening it is physically. This inability to handle scenarios that have not appeared in the training data is a significant challenge to widespread reliance on AI systems.

To be successful in an unpredictable and changing world, an agent must be flexible and master novel situations by using its general knowledge about how such situations are likely to unfold. This is arguably what animals do. Animals are born with most of the skills needed to thrive, or can rapidly acquire them from limited experience, thanks to their strong foundation in real-world interaction, courtesy of evolution and development (Zador 2019). Thus, it is clear that training from scratch for a specific task is not how animals obtain their impressive skills; animals do not arrive into the world *tabula rasa* and then rely on large labeled training sets to learn. Although machine learning has been pursuing approaches for sidestepping this *tabula rasa* limitation, including self-supervised learning, transfer learning, continual learning, meta learning, one-shot learning and imitation learning (Bommasani et al. 2021), none of these approaches comes close to achieving the flexibility found in most animals. Thus, we argue that understanding the neural circuit-level principles that provide the foundation for behavioral flexibility in the real-world, even in simple animals, has the potential to greatly increase the flexibility and utility of AI systems. Put another way, we can greatly accelerate our search for general-purpose circuits for real-world interaction by taking advantage of the optimization process that evolution has already engaged in (Gupta et al. 2021; Stöckl, Lang, and Maass 2022; Koulakov et al. 2022; Stanley et al. 2019; Pehlevan and Chklovskii 2019).

Energy efficiency. One important challenge for modern AI – that our brains have overcome – is energy efficiency. Training a neural network requires enormous amounts of energy. For example, training a large language model such as GPT-3 requires over 1000 megawatt-hours, enough to power a small town for a day (Patterson et al. 2021). The total amount of energy being used to train AI systems is large and growing rapidly. Biological systems are, by contrast, much more energy efficient: The human brain uses

about 20 watts (Sokoloff 1960). The difference in energy requirement between brains and computers derives from differences in information processing. First, at an algorithmic level, modern large-scale ANNs, such as large language models (Brown et al. 2020), rely on very large feedforward architectures with self-attention to process sequences over time (Vaswani et al. 2017), ignoring the potential power of recurrence for processing sequential information. One reason for this is that currently we do not have efficient mechanisms for credit assignment calculations in recurrent networks. In contrast, brains utilize flexible recurrent architectures for dealing with sequences over time and apparently can solve the temporal credit assignment problem with great efficiency—even more efficiently than the feedforward credit assignment mechanisms used in current ANNs. If we could use the brain for guidance on how to craft efficient training mechanisms for recurrent circuits then we could potentially increase our abilities in processing sequential data while further increasing the energy efficiency of our systems. Second, at an implementation level, biological neurons interact mainly by transmitting action potentials (spikes), an asynchronous communication protocol. Like the interactions between conventional digital elements, the output of a neuron can be viewed as a string of 0s and 1s; but unlike a digital computer, the energy cost of a “1” (i.e. of a spike) is several orders of magnitude higher than that of a “0” (Attwell and Laughlin 2001). Because biological circuits operate in a regime where spikes are sparse—even very active neurons rarely exceed a duty cycle¹ of 10%, and most operate at much lower rates—they are much more energy efficient (Lennie 2003).

In addition, other factors may contribute to the energy efficiency of biological networks. For example, biological networks compute effectively even though some of the components are highly unreliable, or “noisy”. Synaptic release—the means by which neurons communicate—can be so unreliable that only 1 out of every 10 messages is transmitted (Dobrunz and Stevens 1997). Circuits are organized so that spike trains are also highly variable, a feature that may allow neural circuits to perform probabilistic inference, a robust form of computation in the presence of uncertainty (Ma et al. 2006). Although there are on-going efforts to exploit the potential of spiking networks (Davies et al. 2018; DeBole et al. 2019), to date no “killer application” has emerged that these networks are able to execute with the energy efficiency of biological circuits. Arguably, the major problem has been that current “neuromorphic chips” neither replicate innate neural circuit functions, nor are they easy to train (Roy, Jaiswal, and Panda 2019). As such, though they are more energy efficient, they are far less useful than their energy hungry digital counterparts. Thus, we believe that obtaining greater energy efficiency in AI could come not only from borrowing the idea of sparse spiking networks, but also by providing neuromorphic chips with innate neural circuit functions and learning rules.

A roadmap for solving the embodied Turing test

How might artificial systems that pass the embodied Turing test be developed? One natural approach would be to do so incrementally, guided by our evolutionary history. For example, almost all animals engage in goal-directed locomotion; they move toward some stimuli (e.g. food sources) and away from

¹ Assuming a temporal discretization of e.g. 1 msec, the duty cycle can be defined as the average number of spikes per second divided by 1000, i.e. the number of “1”s out of 1000 possible positions. Thus, a neuron with an average firing rate of 100 spikes/second has a duty cycle of 10%, and a neuron with an average firing rate of 1 spike/second has a duty cycle of 0.1%

others (e.g. threats). Layered on top of these foundational abilities are more sophisticated skills, such the ability to combine different streams of sensory information (e.g. visual and olfactory), to use this sensory information to distinguish food sources and threats, to navigate to previous locations, to weigh possible rewards and threats to achieve goals, and to interact with the world in precise ways in service of these goals. Most of these—and many other—sophisticated abilities are found to some extent in even very simple organisms, such as worms. In more complex animals, such as fish and mammals, these abilities are elaborated and combined with new strategies to enable more powerful behavioral strategies.

This evolutionary perspective suggests a strategy for solving the embodied Turing test by breaking it down into a series of incremental challenging ones that build on each other, and iteratively optimizing on this series (Cisek and Hayden 2022). Moreover, organisms representing solutions to the lower and intermediate challenges could include worms, flies, fish, rodents and primates, widely used systems in neuroscience research. This would enable us to build on the vast amount of knowledge we have accumulated about the circuitry and mechanisms underlying the behaviors of these model organisms. Much of this research could be performed *in silico*, using virtual environments and virtual animals (Merel, Botvinick, and Wayne 2019; Merel et al. 2019). To achieve the required level of behavioral flexibility, the artificial systems that pass the embodied Turing test would be challenged with a constellation of species-specific tests probing self-supervised learning, continual learning, transfer learning, meta learning, and life-long memory. These challenges can be standardized to permit the quantification of progress. Ultimately, successful virtual organisms could be adapted to the physical world with additional efforts in robotics and deployed to solve real-world problems.

What we need

Achieving these goals will require significant resources and also contribution across many disciplines beyond traditional AI and neuroscience, including psychology, engineering, linguistics, etc. More than simply harnessing existing expertise in these fields, an imperative ***will be to train a new generation of AI researchers who are equally at home in engineering/computational science and neuroscience***. These researchers will chart fundamentally new directions in AI research by drawing on decades of progress in neuroscience. The greatest challenge will be in determining how to exploit the synergies and overlaps in neuroscience, computational science and other relevant fields to advance our quest: identifying what details of the brain's circuitry, biophysics, and chemistry are important and what can be disregarded in the application to AI. Hence, there is a critical need for researchers with suitable training across the different fields to abstract neuroscience knowledge in a way that makes it applicable to computers and help design experiments to generate new neurobiological insights with relevance to AI. The success of this research program depends on the formation of a community of researchers with expertise in neuroscience and AI. Moreover, explicit design of new training programs can ensure that the NeuroAI research community reflects the demographics of society as a whole.

Second, we will need to create a shared platform capable of developing and testing these virtual agents.

One of the greatest technical challenges that we will face in creating an iterative, embodied Turing test and evolving artificial organisms to solve it is the amount of computational power required. Currently, training just one large neural network model on a single embodied task (e.g. control of a body in

3-dimensional space) can take days on specialized distributed hardware (Liu et al. 2021). For multiple research groups to iteratively work together to optimize and evaluate a large number of agents over multiple generations on increasingly complex embodied Turing tasks, a large investment in a shared computational platform will be required. Much like a particle accelerator in physics or large telescope in astronomy, this sort of large-scale shared resource will be essential for moving the brain-inspired AI research agenda forward. It will require a major organizational effort, with government and ideally also industry support, that has as its central goal scientific progress on animal and human-like intelligence.

Third, *we will need to support fundamental theoretical and experimental research on neural computation*. We have learned a tremendous amount about the brain over the last decades, through the efforts of the NIH, in no small measure due to the BRAIN Initiative, and other major funders, and we are now reaching an understanding of the vast diversity of the brain's individual cellular elements, neurons, and how they function as parts of simple circuits. With these building blocks in place, we are poised to shift our focus toward understanding how the brain functions as an integrated intelligent system. This will require insight into how a hundred billion neurons of a thousand different types, each one communicating with thousands of other neurons, with variable, adaptable connections, are wired together, and the computational capabilities – the intelligence – that emerges. We must *reverse engineer the brain* to abstract the underlying principles. Note that the development of virtual agents will itself greatly accelerate this effort by allowing for direct comparisons between experiments in real and 'in-silico' animals, efforts that will provide insights into the neural circuit-level attributes and mechanisms essential for robust control, flexible behavior, energy efficiency, and intelligent behavior. Taking advantage of the powerful synergies between neuroscience and AI will require program and infrastructure support to organize and enable research across the disciplines at a large scale.

Fortunately, there is now bipartisan agreement in Washington, D.C., that investments into AI research are essential to the technological future of the U.S. Community-wide efforts to bridge the fields of neuroscience and AI will require robust investments from federal resources, as well as oversight of project milestones, commercialization support, ethics and big bets on innovative ideas. There are currently some lines of federal resourcing such as the NSF's National Artificial Intelligence Research Institutes explicitly dedicated to driving innovation and discovery in AI from neuroscience research, but these are largely designed to support a traditional academic model with different groups investigating different questions, rather than the creation of a centralized effort that could create something like the embodied Turing test. Likewise, AI support grants are predominantly ancillary programs through the NIH, NSF, DoD, and even the EPA – each of which have their own directives and goals. This leaves a significant funding gap for technology development as an end in itself. The creation of overarching directives either through existing entities, or as a stand-alone agency, to support NeuroAI and AI research would drive this mission, solidifying the U.S. government as an international leader in AI R&D.

Conclusions

Despite the long history of neuroscience driving advances in AI and the tremendous potential for future advances, most engineers and computational scientists in the field are unaware of the history and

opportunities. The influence of neuroscience on shaping the thinking of von Neumann, Turing and other giants of computational theory are rarely mentioned in a typical computer science curriculum. Leading AI conferences such as NeurIPS, which once served to showcase the latest advances in both computational neuroscience and machine learning, now focus almost exclusively on the latter. Even some researchers aware of the historical importance of neuroscience in shaping the field often argue that it has lost its relevance. “Engineers don’t study birds to build better planes” is the usual refrain. But the analogy fails, in part because pioneers of aviation did indeed study birds (Lilienthal 1911; Culick 2001), and some still do (Shyy et al. 2008; Akos et al. 2010). Moreover, the analogy fails also at a more fundamental level: The goal of modern aeronautical engineering is not to achieve “bird-level” flight, whereas a major goal of AI is indeed to achieve (or exceed) “human-level” intelligence. Just as computers exceed humans in many respects, such as the ability to compute prime numbers, so too do planes exceed birds in characteristics such as speed, range and cargo capacity. But if the goal of aeronautical engineers were indeed to build a machine with the “bird-level” ability to fly through dense forest foliage and alight gently on a branch, they would be well-advised to pay very close attention to how birds do it. Similarly, if AI aims to achieve animal-level common-sense sensorimotor intelligence, researchers would be well-advised to learn from animals and the solutions they evolved to behave in an unpredictable world.

References

- Akos, Zsuzsa, Máté Nagy, Severin Leven, and Tamás Vicsek. 2010. “Thermal Soaring Flight of Birds and Unmanned Aerial Vehicles.” *Bioinspiration & Biomimetics* 5 (4): 045003.
- Attwell, D., and S. B. Laughlin. 2001. “An Energy Budget for Signaling in the Grey Matter of the Brain.” *Journal of Cerebral Blood Flow and Metabolism: Official Journal of the International Society of Cerebral Blood Flow and Metabolism* 21 (10): 1133–45.
- Bommasani, Rishi, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, et al. 2021. “On the Opportunities and Risks of Foundation Models.” *arXiv [cs.LG]*. arXiv. <http://arxiv.org/abs/2108.07258>.
- Brown, Tom, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D. Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, and Amanda Askell. 2020. “Language Models Are Few-Shot Learners.” *Advances in Neural Information Processing Systems* 33: 1877–1901.
- Campbell, Murray, A. Joseph Hoane, and Feng-Hsiung Hsu. 2002. “Deep Blue.” *Artificial Intelligence* 134 (1): 57–83.
- Cisek, Paul, and Benjamin Y. Hayden. 2022. “Neuroscience Needs Evolution.” *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences* 377 (1844): 20200518.
- Culick, Fred. 2001. “What the Wright Brothers Did and Did Not Understand about Flight Mechanics - In Modern Terms.” In *37th Joint Propulsion Conference and Exhibit*. Joint Propulsion Conferences. American Institute of Aeronautics and Astronautics.
- Davies, Mike, Narayan Srinivasa, Tsung-Han Lin, Gautham China, Yongqiang Cao, Sri Harsha Choday, Georgios Dimou, et al. 2018. “Loihi: A Neuromorphic Manycore Processor with On-Chip Learning.” *IEEE Micro* 38 (1): 82–99.
- DeBole, Michael V., Brian Taba, Arnon Amir, Filipp Akopyan, Alexander Andreopoulos, William P. Risk, Jeff Kusnitz, et al. 2019. “TrueNorth: Accelerating From Zero to 64 Million Neurons in 10 Years.” *Computer* 52 (5): 20–29.
- Dobrunz, L. E., and C. F. Stevens. 1997. “Heterogeneity of Release Probability, Facilitation, and Depletion at Central Synapses.” *Neuron* 18 (6): 995–1008.
- Gupta, Agrim, Silvio Savarese, Surya Ganguli, and Li Fei-Fei. 2021. “Embodied Intelligence via Learning and Evolution.” *Nature Communications* 12 (1): 5721.
- Hassabis, Demis, Dharshan Kumaran, Christopher Summerfield, and Matthew Botvinick. 2017. “Neuroscience-Inspired Artificial Intelligence.” *Neuron* 95 (2): 245–58.
- Huang, Sandy, Nicolas Papernot, Ian Goodfellow, Yan Duan, and Pieter Abbeel. 2017. “Adversarial Attacks on Neural Network Policies.” *arXiv [cs.LG]*. arXiv. <http://arxiv.org/abs/1702.02284>.
- Hubel, D. H., and T. N. Wiesel. 1962. “Receptive Fields, Binocular Interaction and Functional Architecture in the Cat’s Visual Cortex.” *The Journal of Physiology*. <https://doi.org/10.1113/jphysiol.1962.sp006837>.
- Itti, L., C. Koch, and E. Niebur. 1998. “A Model of Saliency-Based Visual Attention for Rapid Scene Analysis.” *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20 (11): 1254–59.
- Kosoy, Eliza, Adrian Liu, Jasmine L. Collins, David Chan, Jessica B. Hamrick, Nan Rosemary Ke, Sandy Huang, Bryanna Kaufmann, John Canny, and Alison Gopnik. 11--13 Apr 2022. “Learning Causal Overhypotheses through Exploration in Children and Computational Models.” In *Proceedings of the First Conference on Causal Learning and Reasoning*, edited by Bernhard Schölkopf, Caroline Uhler, and Kun Zhang, 177:390–406. Proceedings of Machine Learning Research. PMLR.
- Koulakov, Alexei, Sergey Shuvaev, Divyansha Lachi, and Anthony Zador. 2022. “Encoding Innate Ability through a Genomic Bottleneck.” *bioRxiv*. <https://doi.org/10.1101/2021.03.16.435261>.
- Larochelle, Hugo, and Geoffrey Hinton. 2010. “Learning to Combine Foveal Glimpses with a

- Third-Order Boltzmann Machine.” *Advances in Neural Information Processing Systems* 23.
<https://proceedings.neurips.cc/paper/2010/hash/677e09724f0e2df9b6c000b75b5da10d-Abstract.html>
- LeCun, Yann, and Yoshua Bengio. 1995. “Convolutional Networks for Images, Speech, and Time Series.” *The Handbook of Brain Theory and Neural* 3361 (10).
<http://www.iro.umontreal.ca/~lisa/pointeurs/handbook-convo.pdf>.
- Lennie, Peter. 2003. “The Cost of Cortical Computation.” *Current Biology: CB* 13 (6): 493–97.
- Lilienthal, Otto. 1911. *Birdflight as the Basis of Aviation: A Contribution Towards a System of Aviation, Compiled from the Results of Numerous Experiments Made by O. and G. Lilienthal*. Longmans, Green.
- Liu, Siqu, Guy Lever, Zhe Wang, Josh Merel, S. M. Ali Eslami, Daniel Hennes, Wojciech M. Czarnecki, et al. 2021. “From Motor Control to Team Play in Simulated Humanoid Football.” *arXiv [cs.AI]*. arXiv. <http://arxiv.org/abs/2105.12196>.
- Ma, Wei Ji, Jeffrey M. Beck, Peter E. Latham, and Alexandre Pouget. 2006. “Bayesian Inference with Probabilistic Population Codes.” *Nature Neuroscience* 9 (11): 1432–38.
- McCulloch, Warren S., and Walter Pitts. 1943. “A Logical Calculus of the Ideas Immanent in Nervous Activity.” *The Bulletin of Mathematical Biophysics* 5 (4): 115–33.
- Merel, Josh, Diego Aldarondo, Jesse Marshall, Yuval Tassa, Greg Wayne, and Bence Ölveczky. 2019. “Deep Neuroethology of a Virtual Rodent.” *arXiv [q-bio.NC]*. arXiv. <http://arxiv.org/abs/1911.09451>.
- Merel, Josh, Matthew Botvinick, and Greg Wayne. 2019. “Hierarchical Motor Control in Mammals and Machines.” *Nature Communications* 10 (1): 5489.
- Mnih, Volodymyr, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, et al. 2015. “Human-Level Control through Deep Reinforcement Learning.” *Nature* 518 (7540): 529–33.
- Moravec, Hans. 1988. *Mind Children: The Future of Robot and Human Intelligence*. Harvard University Press.
- Patterson, David, Joseph Gonzalez, Quoc Le, Chen Liang, Lluís-Miquel Munguia, Daniel Rothchild, David So, Maud Texier, and Jeff Dean. 2021. “Carbon Emissions and Large Neural Network Training.” *arXiv [cs.LG]*. arXiv. <http://arxiv.org/abs/2104.10350>.
- Pehlevan, Cengiz, and Dmitri B. Chklovskii. 2019. “Neuroscience-Inspired Online Unsupervised Learning Algorithms: Artificial Neural Networks.” *IEEE Signal Processing Magazine* 36 (6): 88–96.
- Rescorla, R. A. 1972. “A Theory of Pavlovian Conditioning : Variations in the Effectiveness of Reinforcement and Nonreinforcement.” *Current Research and Theory*, 64–99.
- Roy, Kaushik, Akhilesh Jaiswal, and Priyadarshini Panda. 2019. “Towards Spike-Based Machine Intelligence with Neuromorphic Computing.” *Nature* 575 (7784): 607–17.
- Schultz, W., P. Dayan, and P. R. Montague. 1997. “A Neural Substrate of Prediction and Reward.” *Science* 275 (5306): 1593–99.
- Sejnowski, Terrence. 2022. “Large Language Models and the Reverse Turing Test.” *arXiv [cs.CL]*. arXiv. <http://arxiv.org/abs/2207.14382>.
- Shyy, Wei, Yongsheng Lian, Jian Tang, Dragos Viieru, and Hao Liu. 2008. *Aerodynamics of Low Reynolds Number Flyers*.
- Silver, David, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, et al. 2016. “Mastering the Game of Go with Deep Neural Networks and Tree Search.” *Nature* 529 (7587): 484–89.
- Sinz, Fabian H., Xaq Pitkow, Jacob Reimer, Matthias Bethge, and Andreas S. Tolias. 2019. “Engineering a Less Artificial Intelligence.” *Neuron* 103 (6): 967–79.
- Sokoloff, Louis. 1960. “The Metabolism of the Central Nervous System in Vivo.” *Handbook of*

- Physiology, Section I, Neurophysiology* 3: 1843–64.
- Stanley, Kenneth O., Jeff Clune, Joel Lehman, and Risto Miikkulainen. 2019. “Designing Neural Networks through Neuroevolution.” *Nature Machine Intelligence* 1 (1): 24–35.
- Stöckl, Christoph, Dominik Lang, and Wolfgang Maass. 2022. “Structure Induces Computational Function in Networks with Diverse Types of Spiking Neurons.” *bioRxiv*.
<https://doi.org/10.1101/2021.05.18.444689>.
- Thorndike, L., and Darryl Bruce. 2017. *Animal Intelligence: Experimental Studies*. Routledge.
- Turing, A. M. 1950. “I.—COMPUTING MACHINERY AND INTELLIGENCE.” *Mind; a Quarterly Review of Psychology and Philosophy* LIX (236): 433–60.
- Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. “Attention Is All You Need.” *Advances in Neural Information Processing Systems* 30. <https://proceedings.neurips.cc/paper/7181-attention-is-all-you-need>.
- Von Neumann, John. 2012. *The Computer and the Brain*. 3rd ed. Yale University Press.
<https://doi.org/10.12987/9780300188080>.
- Xu, K., J. Ba, R. Kiros, K. Cho, and A. Courville. 2015. “Show, Attend and Tell: Neural Image Caption Generation with Visual Attention.” *International*. <https://proceedings.mlr.press/v37/xuc15.html>.
- Zador, Anthony M. 2019. “A Critique of Pure Learning and What Artificial Neural Networks Can Learn from Animal Brains.” *Nature Communications* 10 (1): 3770.