

UCLA

UCLA Previously Published Works

Title

Catalyzing next-generation Artificial Intelligence through NeuroAI

Permalink

<https://escholarship.org/uc/item/1s66n7s9>

Journal

Nature Communications, 14(1)

ISSN

2041-1723

Authors

Zador, Anthony
Escola, Sean
Richards, Blake
et al.

Publication Date

2023-03-01

DOI

10.1038/s41467-023-37180-x

Peer reviewed

Catalyzing next-generation Artificial Intelligence through NeuroAI

Received: 11 September 2022

Accepted: 3 March 2023

Published online: 22 March 2023

 Check for updates

Anthony Zador^{1,29} ✉, Sean Escola^{2,29}, Blake Richards^{3,4,5,6,7}, Bence Ölveczky⁸, Yoshua Bengio³, Kwabena Boahen⁹, Matthew Botvinick¹⁰, Dmitri Chklovskii¹¹, Anne Churchland¹², Claudia Clopath¹³, James DiCarlo¹⁴, Surya Ganguli¹⁵, Jeff Hawkins¹⁶, Konrad Körding¹⁷, Alexei Koulikov¹, Yann LeCun^{18,19}, Timothy Lillicrap¹⁰, Adam Marblestone²⁰, Bruno Olshausen²¹, Alexandre Pouget²², Cristina Savin²³, Terrence Sejnowski²⁴, Eero Simoncelli²⁵, Sara Solla²⁶, David Sussillo^{18,27}, Andreas S. Tolias²⁸ & Doris Tsao²¹

Neuroscience has long been an essential driver of progress in artificial intelligence (AI). We propose that to accelerate progress in AI, we must invest in fundamental research in NeuroAI. A core component of this is the embodied Turing test, which challenges AI animal models to interact with the sensorimotor world at skill levels akin to their living counterparts. The embodied Turing test shifts the focus from those capabilities like game playing and language that are especially well-developed or uniquely human to those capabilities – inherited from over 500 million years of evolution – that are shared with all animals. Building models that can pass the embodied Turing test will provide a roadmap for the next generation of AI.

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 and catalyze economic growth, relieving workers from performing the most dangerous and menial jobs. However, to reach this potential, we still require advances that will make AI more

human-like in its capabilities. Historically, neuroscience has been a critical 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^{1,2}. It can still play this role. To accelerate progress in AI and realize its vast potential, we must invest in fundamental research in “NeuroAI.”

¹Cold Spring Harbor Laboratory, Cold Spring Harbor, NY 11724, USA. ²Department of Psychiatry, Columbia University, New York, NY 10027, USA. ³Mila, Montréal, QC H2S 3H1, Canada. ⁴School of Computer Science, McGill University, Montreal, Canada. ⁵Montreal Neurological Institute, McGill University, Montreal, Canada. ⁶Department of Neurology & Neurosurgery, McGill University, Montreal, Canada. ⁷Learning in Machines and Brains Program, CIFAR, Toronto, Canada. ⁸Department of Organismic and Evolutionary Biology, Harvard University, Cambridge, MA 02138, USA. ⁹Department of Bioengineering, Stanford University, Stanford, CA 94305, USA. ¹⁰Google Deepmind, London NIC 4AG, UK. ¹¹Flatiron Institute, Simons Foundation, New York, NY 10010, USA. ¹²Department of Neurobiology, University of California Los Angeles, Los Angeles, CA 90095, USA. ¹³Department of Bioengineering, Imperial College London, London SW7 2BW, UK. ¹⁴Department of Brain and Cognitive Sciences, MIT, Cambridge, MA 02139, USA. ¹⁵Department of Applied Physics, Stanford University, Stanford, CA 94305, USA. ¹⁶Numenta, Redwood City, CA 94063, USA. ¹⁷Department of Neuroscience, University of Pennsylvania, Philadelphia, PA 19104, USA. ¹⁸Meta, Menlo Park, CA 94025, USA. ¹⁹Department of Electrical and Computer Engineering, NYU, Brooklyn, NY 11201, USA. ²⁰Media Lab, MIT, Cambridge, MA 02140, USA. ²¹Helen Wills Neuroscience Institute, University of California Berkeley, Berkeley, CA 94720, USA. ²²Department of Basic Neurosciences, University of Geneva, Genève 1211, Switzerland. ²³Center for Neural Science, NYU, New York, NY 10003, USA. ²⁴Salk Institute for Biological Studies, La Jolla, CA 92037, USA. ²⁵Departments of Neural Science, Mathematics, and Psychology, NYU, New York, NY 10003, USA. ²⁶Department of Physiology, Northwestern University, Chicago, IL 60611, USA. ²⁷Department of Electrical Engineering, Stanford University, Stanford, CA 94305, USA. ²⁸Department of Neuroscience, Baylor College of Medicine, Houston, TX 77030, USA. ²⁹These authors contributed equally: Anthony M. Zador, Sean Escola. ✉e-mail: zador@cshl.edu

The seeds of the current AI revolution were planted decades ago, mainly by researchers attempting to understand how brains compute³. 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^{4,5}. Later, the Nobel-prize winning work of David Hubel and Torsten Wiesel on visual processing circuits in the cat neocortex inspired the deep convolutional networks that have catalyzed the recent revolution in modern AI^{6–8}. Similarly, the development of reinforcement learning was directly inspired by insights into animal behavior and neural activity during learning^{9–15}. 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 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 chess¹⁶ and Go¹⁷, they are not robust and often struggle when faced with novel situations. Moreover, we have yet to build effective systems that can walk to the shelf, take down the chess set, set up the pieces, and move them around during a game, although recent progress is encouraging¹⁸. 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 basic capacities required to navigate new situations—capacities that animals have or acquire effortlessly—turn out to be deceptively challenging for AI, partly 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 needed¹⁹.

Historically, many key AI advances, such as convolutional ANNs and reinforcement learning, were inspired by neuroscience. Neuroscience continues to provide guidance—e.g., attention-based neural networks were loosely inspired by attention mechanisms in the brain^{20–23}—but this is often based on findings that are decades old. The fact that such cross-pollination between AI and neuroscience is far less common than in the past 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. *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 fundamental ingredients of intelligence*

and catalyze the next revolution in AI. This will eventually lead to artificial agents with capabilities that match those of humans. The NeuroAI program we advocate is driven by the recognition that AI historically owes much to neuroscience and the promise that AI will continue to learn from it—but only if there is a large enough community of researchers fluent in both domains. We believe the time is right for a large-scale effort to identify and understand the principles of biological intelligence and abstract those for application in computer and robotic systems.

It is tempting to focus on the most characteristically human aspects of intelligent behavior, such as abstract thought and reasoning. However, 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 Moravec²⁴ 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 implies that the bulk of the work in developing general AI can be achieved by building systems that match the perceptual and motor abilities of animals and that the subsequent step to human-level intelligence would be considerably smaller. This is good news because progress on the first goal can rely on the favored subjects of neuroscience research—rats, mice, and non-human primates—for which extensive and rapidly expanding behavioral and neural datasets can guide the way. 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”²⁵ as a test of a machine’s ability to exhibit intelligent behavior indistinguishable from that of a human (Fig. 1, left). In that game, now known as the Turing test, a human judge evaluates natural language conversations between a real human and a machine trained to mimic human responses. By focusing on conversational abilities, Turing evaded asking whether a machine could “think,” a question he considered impossible to answer. The Turing test is based on the implicit belief that language represents the pinnacle of human intelligence and that a machine capable of conversation must surely be intelligent.

Until recently, no artificial system could come close to passing the Turing test. However, a class of modern AI systems called “large language models” can now engage in surprisingly cogent conversations²⁶. In part, their success reveals how easily we can be tricked into imputing intelligence, agency, and even consciousness to our interlocutor²⁷. Impressive though these systems are, because they are not grounded in real-world experiences, they nonetheless continue to struggle with many basic aspects of causal reasoning and physical common-sense. Thus, the Turing test does not probe our prodigious perceptual and motor abilities to interact with and reason about the physical world, abilities shared with animals and honed through countless generations of natural selection.

We therefore propose an expanded “embodied Turing test,” one that includes advanced sensorimotor abilities (Fig. 1, right). 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. This embodied Turing test would benchmark and compare the interactions with the world of artificial systems versus humans and other animals. Similar ideas have been proposed previously^{28–32}. However, in light of recent advances enabling large-scale behavioral and neural measurements, as well as large-scale simulations of embodied agents in silico, we believe the time is ripe to instantiate a major research effort in this direction. As each animal has its own unique set of abilities, each animal defines its own embodied




Fig. 1 | Turing tests: comparisons between the performance of AI systems and their living counterparts. *Left:* The original Turing test as proposed by Alan Turing²⁵. If a human tester cannot determine whether their interlocutor is an AI system or another human, the AI passes the test. Modern large language models have made substantial progress towards passing this test²⁶. *Right:* The embodied Turing test. An AI animal model—whether robotic or in simulation—passes the test if its behavior is indistinguishable from that of its living counterpart. No AI systems are close to passing this test. Here, an artificial beaver is tested on the species-specific behavior of dam construction.

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. This implies that after developing an AI system to faithfully reproduce the behavior of one species, the adaptation of this system to other species—and even to humans—may be straightforward. Below we highlight a few of the characteristics that are shared across species.

Animals engage their environments

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 capabilities³³. 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). These principles could guide the design of systems for perception, action selection, locomotion, and fine-grained control of limbs, hands, and fingers. Understanding how specific neural circuits participate in different tasks 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.

Animals behave flexibly

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 scores³⁴. 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 performance³⁵. 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 costs 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 development³⁶. 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 learning³⁷, 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^{38–45}.

Animals compute efficiently

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 megawatts-hours, enough to power a small town for a day⁴⁶. Biological systems are, by contrast, much more energy efficient: The human brain uses about 20 watts⁴⁷. 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 models²⁶, rely on very large feedforward architectures with self-attention to process sequences over time²³, 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 that can solve the temporal credit assignment problem with great efficiency. Uncovering the mechanisms by which this happens could potentially enable us to increase the energy efficiency of artificial systems. Alternatively, it has been proposed that the synaptic dynamics within adjacent dendritic spines could serve as a mechanism for learning sequential structure, a scheme that could potentially be efficiently implemented in hardware⁴⁸. Second, at an implementation level, neural circuits differ from digital computers. Neural circuits compute effectively despite the presence of unreliable or “noisy” components. For example, synaptic release, the primary means of communication between neurons, can be so unreliable that only one in every ten messages is transmitted⁴⁹. Furthermore, 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”⁵⁰. As biological circuits operate in a regime where spikes are sparse—even very active neurons rarely fire at more than 100 spikes per second and typical cortical firing rates may be less than 1 spike/second—they are much more energy efficient⁵¹. Spike-based computation has also been shown to be orders of magnitude faster and more energy efficient in recent hardware implementation⁵².

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 others (e.g., threats). Layered on top of these foundational abilities are more sophisticated skills, such as 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 passing the embodied Turing test by breaking it down into a series of incrementally challenging ones that build on each other, and iteratively optimizing on this series⁵³. Specifically, the embodied Turing test comprises challenges that include a wide range of organisms used in neuroscience research, including worms, flies, fish, rodents and primates. This would enable us to deploy the vast amount of knowledge we have accumulated about the behavior, biomechanics, and neural mechanisms of these model organisms to both precisely define each species-specific embodied Turing test and serve as strong inductive biases to guide the development of robust AI controllers that can pass it.

The performance of these artificial agents could be compared with that of animals. Rich behavioral datasets representing a large swath of a species' ethogram have now been collected and can be deployed to benchmark performance on species-specific embodied Turing tests. Furthermore, these datasets are being rapidly expanded given new tools in 3D videography^{54–57}. Additionally, detailed biomechanical measurements support highly realistic animal body models, complete with skeletal constraints, muscles, tendons, and paw features⁵⁸. Combined with the open-sourcing of powerful, fast physics simulators and virtual environments^{59,60}, these models will afford the opportunity for embodied Turing test research to be performed in silico at scale³³. Finally, existing extensive neural datasets with simultaneous neural recordings across multiple brain regions during behavior, combined with increasingly detailed neural anatomy and connectomics, provide a powerful roadmap for the design of AI systems that can control virtual animals to recapitulate the behaviors of their in vivo counterparts and thus pass the embodied Turing test.

Importantly, the specifics of the embodied Turing test for each species can be tuned to the needs of different groups of researchers. We can test the capacity of AI systems in terms of sensorimotor control, self-supervised and continual learning, generalization, memory-guided behavior on both short and life-long timescales, and social interactions. Despite these potentially different areas of interest, the challenges that compose the embodied Turing test can be standardized to permit the quantification of progress and comparison between research efforts. Standardization can be fostered by stakeholders including government and private funders, large research organizations such as the Allen Institute, and major collaborations like the International Brain Lab, with an eye toward the development of common APIs and support for competitions as has been an important impetus for much progress in machine learning and robotics^{61,62}. Ultimately, virtual organisms that demonstrate successful recapitulation of behaviors of interest can 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 deployed in three main areas. *First, we must 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. There is thus a critical need for researchers with dual training in AI and neuroscience to apply insights from neuroscience to advance AI and to help design experiments that generate new insights relevant to AI. Although there is already some research of this type, it exists largely at the margins of

mainstream neuroscience; training in neuroscience has thus far been motivated and funded mainly by the goal of improving human health and of understanding the brain as such. This lack of alignment between fields might explain, e.g., the multi-decade gap between Hubel and Wiesel's discovery of the structure of the visual system⁶ and the development and application of convolutional neural networks in modern machine learning⁸. Thus, the success of a NeuroAI research program depends on the formation of a community of researchers for whom the *raison d'être* of their training is to exploit synergies between neuroscience and AI. Explicit design of new training programs can ensure that the NeuroAI research community reflects the demographics of society as a whole and is equipped with the ethical tools needed to ensure that the development of AI benefits society⁶³.

Second, we must 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 pass 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 hardware⁶⁴. 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 tests, 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 must 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. 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 broad political agreement that investments in AI research are essential to humanity's technological future. Indeed, IARPA (Intelligence Advanced Research Projects Activity) was a pioneer in this field, launching the Machine Intelligence from Cortical Networks (MICrONS) project. This project spearheaded the collection of an unprecedented data set consisting of a portion of a mouse connectome and associated functional responses with the specific goal of catalyzing the development of next-generation AI algorithms⁶⁵. Nonetheless, community-wide efforts to bridge the fields of neuroscience and AI will require robust investments from government resources, as well as oversight of project milestones, commercialization support, ethics, and big bets on innovative ideas. In the U.S., 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 in the U.S. are predominantly ancillary programs through the NIH, NSF, DoD, and even the EPA—each of which have their own directives and goals—and this pattern is shared by funding agencies globally. 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 and consequently unlock the potential for AI to benefit humanity.

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. However, the analogy fails, in part because pioneers of aviation did indeed study birds^{66,67}, and some still do^{68,69}. 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. However, 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

1. Hassabis, D., Kumaran, D., Summerfield, C. & Botvinick, M. Neuroscience-inspired artificial intelligence. *Neuron* **95**, 245–258 (2017).
2. Macpherson, T. et al. Natural and artificial intelligence: a brief introduction to the interplay between AI and neuroscience research. *Neural Netw.* **144**, 603–613 (2021).
3. McCulloch, W. S. & Pitts, W. A logical calculus of the ideas immanent in nervous activity. *Bull. Math. Biophys.* **5**, 115–133 (1943).
4. von Neumann, J. *First Draft of a Report on the EDVAC*. <https://doi.org/10.5479/sil.538961.39088011475779> (1945).
5. von Neumann, J. *The Computer and the Brain* (Yale University Press, 2012).
6. Hubel, D. H. & Wiesel, T. N. Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *J. Physiol.* **160**, 106–154 <https://doi.org/10.1113/jphysiol.1962.sp006837> (1962).
7. Fukushima, K. Neocognitron: a self organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biol. Cybern.* **36**, 193–202 (1980).
8. LeCun, Y. & Bengio, Y. Convolutional networks for images, speech, and time series. In: *The Handbook of Brain Theory and Neural*. 255–258 (ACM, 1995).
9. Thorndike, E. L. Animal intelligence: an experimental study of the associative processes in animals. <https://doi.org/10.1037/10780-000> (1898).
10. Thorndike, E. L. The law of effect. *The Am. J. Psychol.* **39**, 212 <https://doi.org/10.2307/1415413> (1927).

11. Thorndike, E. L. The fundamentals of learning. <https://doi.org/10.1037/10976-000> (1932).
12. Crow, T. J. Cortical synapses and reinforcement: a hypothesis. *Nature* **219**, 736–737 (1968).
13. Rescorla, R. A. A theory of pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement. In: Black, A. H. & Prokasy, W. F. (eds.) *Classical Conditioning II: Current Research and Theory*. 64–99 (Century-Crofts, 1972).
14. Klopff, A. H. *Brain Function and Adaptive Systems: A Heterostatic Theory* (AIR FORCE CAMBRIDGE RESEARCH LABS HANSCOM AFB MA, 1972).
15. Schultz, W., Dayan, P. & Montague, P. R. A neural substrate of prediction and reward. *Science* **275**, 1593–1599 (1997).
16. Campbell, M., Hoane, A. J. & Hsu, F.-H. Deep blue. *Artif. Intell.* **134**, 57–83 (2002).
17. Silver, D. et al. Mastering the game of Go with deep neural networks and tree search. *Nature* **529**, 484–489 (2016).
18. Reed, S. et al. A generalist agent. <https://arxiv.org/abs/2205.06175> (2022).
19. Sinz, F. H., Pitkow, X., Reimer, J., Bethge, M. & Tolias, A. S. Engineering a less artificial intelligence. *Neuron* **103**, 967–979 (2019).
20. Itti, L., Koch, C. & Niebur, E. A model of saliency-based visual attention for rapid scene analysis. *IEEE Trans. Pattern Anal. Mach. Intell.* **20**, 1254–1259 (1998).
21. Larochelle, H. & Hinton, G. Learning to combine foveal glimpses with a third-order Boltzmann machine. *Adv. Neural Inform. Process. Syst.* **23**, 1243–1251 (2010).
22. Xu, K. et al. Show, attend and tell: neural image caption generation with visual attention. In: *Proceedings of the 32nd International Conference on Machine Learning* (eds. Bach, F. & Blei, D.) vol. 37, 2048–2057 (PMLR, 2015).
23. Vaswani, A. et al. Attention is all you need. *Adv. Neural Inf. Process. Syst.* **30**, 6000–6010 (2017).
24. Moravec, H. *Mind Children: The Future of Robot and Human Intelligence* (Harvard University Press, 1988).
25. Turing, A. M. I.—Computing machinery and intelligence. *Mind* **LIX**, 433–460 (1950).
26. Brown, T. et al. Language models are few-shot learners. *Adv. Neural Inf. Process. Syst.* **33**, 1877–1901 (2020).
27. Sejnowski, T. Large language models and the reverse turing test. <https://arxiv.org/abs/2207.14382> (2022).
28. Brooks, R. A. Intelligence without representation. *Artificial Intelligence*, **47**, 139–159 [https://doi.org/10.1016/0004-3702\(91\)90053-m](https://doi.org/10.1016/0004-3702(91)90053-m) (1991).
29. Meyer, J.-A. & Wilson, S. W. *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior* (Bradford Books, 1991).
30. Pfeifer, R. & Scheier, C. Understanding intelligence. <https://doi.org/10.7551/mitpress/6979.001.0001> (2001).
31. Pfeifer, R. & Bongard, J. *How the Body Shapes the Way We Think: A New View of Intelligence* (MIT Press, 2006).
32. Ortiz, C. L. Why we need a physically embodied turing test and what it might look like. *AI Magazine*. vol. 37, 55–62 <https://doi.org/10.1609/aimag.v37i1.2645> (2016).
33. Merel, J., Botvinick, M. & Wayne, G. Hierarchical motor control in mammals and machines. *Nat. Commun.* **10**, 5489 (2019).
34. Mnih, V. et al. Human-level control through deep reinforcement learning. *Nature* **518**, 529–533 (2015).
35. Huang, S., Papernot, N., Goodfellow, I., Duan, Y. & Abbeel, P. Adversarial attacks on neural network policies. <https://arxiv.org/abs/1702.02284> (2017).
36. Zador, A. M. A critique of pure learning and what artificial neural networks can learn from animal brains. *Nat. Commun.* **10**, 3770 (2019).

37. Bommasani, R. et al. On the opportunities and risks of foundation models. <https://arxiv.org/abs/2108.07258> (2021).
38. Elman, J. L. Learning and development in neural networks: the importance of starting small. *Cognition* **48**, 71–99 (1993).
39. Lake, B. M., Ullman, T. D., Tenenbaum, J. B. & Gershman, S. J. Building machines that learn and think like people. *Behav. Brain Sci.* **40**, e253 (2017).
40. Doya, K. & Taniguchi, T. Toward evolutionary and developmental intelligence. *Curr. Opin. Behav. Sci.* **29**, 91–96 <https://doi.org/10.1016/j.cobeha.2019.04.006> (2019).
41. Pehlevan, C. & Chklovskii, D. B. Neuroscience-inspired online unsupervised learning algorithms: artificial neural networks. *IEEE Signal Process. Mag.* **36**, 88–96 (2019).
42. Stanley, K. O., Clune, J., Lehman, J. & Miikkulainen, R. Designing neural networks through neuroevolution. *Nat. Mach. Intell.* **1**, 24–35 (2019).
43. Gupta, A., Savarese, S., Ganguli, S. & Fei-Fei, L. Embodied intelligence via learning and evolution. *Nat. Commun.* **12**, 5721 (2021).
44. Stöckl, C., Lang, D. & Maass, W. Structure induces computational function in networks with diverse types of spiking neurons. *bioRxiv*. <https://doi.org/10.1101/2021.05.18.444689> (2022).
45. Koulakov, A., Shuvaev, S., Lachi, D. & Zador, A. Encoding innate ability through a genomic bottleneck. *bioRxiv*. <https://doi.org/10.1101/2021.03.16.435261> (2022).
46. Patterson, D. et al. Carbon emissions and large neural network training. <https://arxiv.org/abs/2104.10350> (2021).
47. Sokoloff, L. The metabolism of the central nervous system in vivo. *Handb. Physiol. Sect. I Neurophysiol.* **3**, 1843–1864 (1960).
48. Boahen, K. Dendrocentric learning for synthetic intelligence. *Nature* **612**, 43–50 (2022).
49. Dobrunz, L. E. & Stevens, C. F. Heterogeneity of release probability, facilitation, and depletion at central synapses. *Neuron* **18**, 995–1008 (1997).
50. Attwell, D. & Laughlin, S. B. An energy budget for signaling in the grey matter of the brain. *J. Cereb. Blood Flow. Metab.* **21**, 1133–1145 (2001).
51. Lennie, P. The cost of cortical computation. *Curr. Biol.* **13**, 493–497 (2003).
52. Davies, M. et al. Advancing neuromorphic computing with loihi: a survey of results and outlook. *Proc. IEEE Inst. Electr. Electron. Eng.* **109**, 911–934 (2021).
53. Cisek, P. & Hayden, B. Y. Neuroscience needs evolution. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **377**, 20200518 (2022).
54. Mathis, A. et al. DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nat. Neurosci.* **21**, 1281–1289 (2018).
55. Wu, A. et al. Deep Graph Pose: a semi-supervised deep graphical model for improved animal pose tracking. *Adv. Neural Inf. Process. Syst.* **33**, 6040–6052 (2020).
56. Marshall, J. D. et al. Continuous whole-body 3D kinematic recordings across the rodent behavioral repertoire. *Neuron* **109**, 420–437.e8 (2021).
57. Pereira, T. D. et al. Publisher Correction: SLEAP: A deep learning system for multi-animal pose tracking. *Nat. Methods* **19**, 628 (2022).
58. Merel, J. et al. Deep neuroethology of a virtual rodent. in *International Conference on Learning Representations* (Association for Computing Machinery, 2020).
59. Todorov, E., Erez, T. & Tassa, Y. MuJoCo: A physics engine for model-based control. in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems* (IEEE, 2012).
60. Brockman, G. et al. OpenAI Gym. (2016) <https://doi.org/10.48550/arXiv.1606.01540>.
61. Kitano, H., Asada, M., Kuniyoshi, Y., Noda, I. & Osawa, E. RoboCup: The Robot World Cup Initiative. in: *Proceedings of the first international conference on Autonomous Agents*. 340–347 (Association for Computing Machinery, 1997).
62. Bell, R. M. & Koren, Y. Lessons from the Netflix prize challenge. *ACM SIGKDD Explorations Newsletter*. vol. 9, 75–79 <https://doi.org/10.1145/1345448.1345465> (2007).
63. Doya, K., Ema, A., Kitano, H., Sakagami, M. & Russell, S. Social impact and governance of AI and neurotechnologies. *Neural Netw.* **152**, 542–554 (2022).
64. Liu, S. et al. From motor control to team play in simulated humanoid football. <https://arxiv.org/abs/2105.12196> (2021).
65. MICrONS Consortium et al. Functional connectomics spanning multiple areas of mouse visual cortex. *bioRxiv* <https://doi.org/10.1101/2021.07.28.454025> (2021).
66. Lilienthal, O. *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, 1911).
67. Culick, F. What the Wright Brothers did and did not understand about flight mechanics—in modern terms. in *37th Joint Propulsion Conference and Exhibit* (American Institute of Aeronautics and Astronautics, 2001).
68. Shyy, W., Lian, Y., Tang, J., Viieru, D. & Liu, H. *Aerodynamics of Low Reynolds Number Flyers*. (Cambridge University Press, 2008).
69. Akos, Z., Nagy, M., Leven, S. & Vicsek, T. Thermal soaring flight of birds and unmanned aerial vehicles. *Bioinspir. Biomim.* **5**, 045003 (2010).

Acknowledgements

This Perspective grew out of discussions following the first conference From Neuroscience to Artificially Intelligent Systems (NAISys), held in CSHL in 2020, and includes many of the invited speakers from that meeting. A.Z. would like to thank Cat Donaldson for convincing him that this manuscript would be worth writing, and for contributing to early versions. B.Ö. would like to acknowledge valuable contributions on the manuscript from Diego Aldorando. The authors would like to acknowledge the following funding sources—A.Z.: Schmidt Foundation, Eleanor Schwartz Foundation, Robert Lourie Foundation; S.E.: NIH 5U19NS104649, 5R01NS105349; R.B.: CIFAR, NSERC RGPIN-2020-05105, RGPAS-2020-00031; B.Ö.: NIH R01NS099323, R01GM136972; Y.B.: CIFAR; K.B.: Stanford Institute for Human-Centered Artificial Intelligence, NSF 2223827; A.C.: NIH U19NS123716; J.D.: Semiconductor Research Corporation, DARPA, Office of Naval Research MURI-114407, MURI-N00014-21-1-2801, N00014-20-1-2589, NSF CCF-1231216, NSF 2124136, Simons SCGB 542965; S.G.: Simons Foundation, James S. McDonnell Foundation, NSF CAREER Award; A.M.: Federation of American Scientists; B.O.: NSF IIS-1718991; C.S.: NIH 1R01MH125571-01, R01NS127122, NSF 1922658, Google faculty award; E.S.: Simons Foundation, NIH EY022428; S.S.: NIH NINDS NS053603; D.S.: Simons 543049; A.T.: IARPA D16PC00003, DARPA HRO011-18-2-0025, NIH R01EY026927, R01MH109556, NSF NeuroNex DBI-1707400; D.T.: HHMI.

Author contributions

A.Z. conceived the project and wrote the first draft. A.Z., S.E., B.R., B.Ö. provided extensive editing to successive drafts. Y.B., K.B., M.B., D.C., A.C., C.C., J.D., S.G., J.H., K.K., A.K., Y.L., T.L., A.M., B.O., A.P., C.S., T.S., E.S., S.S., D.S., A.T., and D.T. provided additional guidance and edits at various stages.

Competing interests

The authors declare their engagements with the following relevant for-profit entities—A.Z.: Cajal Neuroscience; S.E.: Herophilus; R.B.: Google DeepMind; B.Ö.: Blackbird Neuroscience; K.B.: Femtosense Inc., Radical

Semiconductor, Neurovigil; C.C.: Google DeepMind; S.G.: Meta; J.H.: Numenta; K.K.: Paradromics; T.L.: Google DeepMind; A.M.: Google DeepMind, Kernel; D.S.: Meta; A.T.: Vathes Inc., Upload AI LLC, BioAvatar LLC. The remaining authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to Anthony Zador.

Peer review information *Nature Communications* thanks Wolfgang Maass and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2023