



Article development led by [acmqueue](https://queue.acm.org)
queue.acm.org

Is the ability to think scientifically the defining essence of intelligence?

BY EDLYN V. LEVINE

Cargo Cult AI

ASTRONOMER CARL SAGAN once wrote “science is more than a body of knowledge; it is a way of thinking.” This type of thinking requires skeptical rigor and brutal honesty to thoroughly investigate, reason, and seek to invalidate hypotheses before jumping to conclusions. But it is all too easy to jump to conclusions. Despite our self-proclaimed intelligence, humans are apt to believe remarkable fallacies based on a paucity of correlated information rather than rigorously seek to determine causal foundations.

This propensity of humans to believe wonderful, fanciful things so easily is what another physicist, Richard Feynman, called cargo cult science. Feynman named this phenomenon after a “cargo cult” of people in the Pacific Islands who believed that building replicas of landing strips and control towers would ensure supply planes continued to land after World War II.⁶ The planes never came. These people missed

the fact that it was the advent of war, not the presence of landing strips, that caused the planes to land there.

Today, some have speculated that large language models (LLMs) such as GPT-4 can be viewed as early versions of artificial general intelligence (AGI).³ In contrast to AI, which is often task-specific, AGI is assumed to be able to perform any general task that a human might be capable of doing. There is something unsettling about the opinion that LLMs are emergent AGI. LLMs exhibit many behaviors and precepts indicative of intelligence but are missing something essential: the stuffy rigor of scientific inquiry. Today’s AI models are missing the ability to reason abstractly, including asking and answering questions of “Why?” and “How?”

Is the ability to think scientifically the defining essence of intelligence? The truth is we don’t know. There is no comprehensive theory yet to explain what intelligence is or how it emerges from first principles. It would appear evident, however, that today’s LLMs are not able to reproduce scientific thinking that has enabled humans to combine Bacon’s empiricism and Descartes’s rationalism to expand the frontier of falsifiable knowledge in the form of scientific theories. Methods of scientific inquiry have enabled humans to establish aspects of universality, nondeterminism, and causality that ultimately enable manipulation of the natural world to advance human welfare.

Evidence abounds that the human brain does not innately think scientifically; however, it can be taught to do so. The same species that forms cargo cults around widespread and unfounded beliefs in UFOs, ESP, and anything read on social media also produces scientific luminaries such as Sagan and Feynman. Today’s cutting-edge LLMs are also not innately scientific. But unlike the human brain, there is good reason to believe they never will be unless new algorithmic paradigms are developed.



A Dominant Algorithmic Paradigm

Impressive progress in AI, including the recent sensation of ChatGPT, has been dominated by the success of a single, decades-old machine-learning approach called a multilayer (or deep) neural network. This approach was invented in the 1940s,¹⁷ and essentially all the foundational concepts of neural networks (nets)^{11,15} and associated methods—including convolutional neural networks⁷ and back-propagation¹⁹—were in place by the 1980s. However, it was not until the emergence of large digital datasets for training and sufficiently fast hardware in the form of graphics processor units (GPUs) that applications using neural nets have taken off.

The dominance of neural nets in today's AI is a tribute to their impressive emergent capabilities. A neural net is a mathematical function providing a representation of empirical information and computes an output for a

given input. The specific mathematical form of a neural net is that of a weighted, directed graph for which the vertices are called neurons and the edges are called connections. In the case of models such as GPT-3, which has 175 billion connections and thus 175 billion weights,² the function will have billions of terms.

The weights and biases of a neural net are determined through a process called deep learning that uses the back-propagation algorithm to progressively decrease the error between model prediction and training data.¹⁴ The resultant trained neural net model effectively transforms the training data into abstract representations that suppress trivial information and magnify or distort features critical for classification. These abstracted representations were originally used to enable classification of a plethora of diverse data inputs but can also be used in a generative capacity. Today, AI models generate anything

from chat prompts to images (for example, produced by generative adversarial networks). The transformer models behind these generative tasks, of which the LLM GPT-3 is an example, still use the foundational architecture of neural networks with the addition of attention to learn context by tracking relationships in sequential data.²⁴

Deep learning with neural nets has thus proven to be an extremely powerful and flexible computing framework. There are reasons to be concerned, however, that this approach will ultimately plateau if the goal is to achieve AGI capable of scientific reasoning. Neural nets may be fundamentally incapable of doing certain things such as establishing universality, nondeterminism, or causal inference. Even for what they can do, neural networks are incredibly resource intensive. How much more improvement can really be eked out of this approach on the pathway to AGI, and is it sustainable?

From Times of Plenty to Times of Scarcity

The dramatic increases in computational power and memory capacity driven by Moore's Law have fueled an explosion in the data corpus and have enabled the use of resource-heavy approaches to deep learning. Training of Google's BERT (Bidirectional Encoder Representations from Transformers) required 3.3 billion tokens and more than 40 training epochs. Compare this with the average child, who may hear 45 million words by age five.²⁰ This is a factor of 3,000 fewer words than BERT and pales in comparison to the likely hundreds of billions of tokens used to train GPT-3.

Today's data and resource abundance stands in sharp contrast to foundational algorithmic work at the dawn of the computing era when innovations were based on scarcity. Computational memory and processing power were so limited and at such a premium that novel algorithmic approaches were needed to solve problems in scenarios where inefficient, brute-force methods were not possible.

Achieving AGI may require a return to this scarcity mindset in the design of new algorithmic approaches that could dramatically economize information processing and abstracted model generation. The skyrocketing costs and energy consumption associated with training neural networks of ever-larger sizes is unlikely to be sustainable²² and will require this shift. Today's large AI models can cost tens of millions of dollars to train²⁶ and they also consume terawatt-hours of energy annually.²¹ The energy consumed by the human brain is paltry by comparison.

The good news is the data representations in today's AI models are likely to be far from the algorithmically minimal representation required to achieve a certain capability,²⁵ so there is ample room for scarcity-driven algorithmic innovation.

Even if this resourcing problem is solved, there are still issues with fundamental limits of AI in its lack of ability to think scientifically. Current methods will not achieve AGI unless fundamental algorithmic innovations are introduced that enable AI to ask and answer questions of why.

Model Universality—or Lack Thereof

Neural nets are models. They provide a mathematical procedure for calculating a result rather than measuring a result directly. Humans have been developing models for centuries to aid with prediction and understanding, and ultimately to boost productivity. Rather than needing to make a measurement every time specific information is desired, such as the trajectory of a rocket or the energy stored in a capacitor, a mathematical procedure can often be determined that will enable accurate prediction of the result.

The development of models to make such predictions is foundational to theoretical science. The success of a mathematical model often depends on its predictive universality. Specifically, to what extent does the mathematical procedure developed to predict one phenomenon enable successful prediction of entirely different classes of phenomena?

Consider the development of a model to predict planetary motion, a problem addressed by astronomer Johannes Kepler in the 17th century. Kepler devised his famous three laws of planetary motion through careful study of data from fellow astronomer Tycho Brahe's detailed astronomical measurements. These three laws universally describe the orbital shape, speed, and period of planets in the solar system based on their distance from the sun. While these results can be generalized to other planetary systems or other orbital bodies (moons, artificial satellites, among others), they do not translate to non-orbital gravitational phenomena. It took Isaac Newton's breakthroughs in mechanical theory and the theory of gravity to develop a unified mathematical framework that could describe both the motion of the planets and the falling of an apple from a tree.

The Newtonian approach is therefore more universal than that of Kepler, but it is not the end of the story. There are physical scenarios for which the Newtonian model breaks down. Breakthroughs in the early 20th century, including Einstein's work in general relativity and the discovery of the theory of quantum mechanics, have provided more universal approaches to prediction of physical phenomena in different regimes. These mathematical models can then be used to accurately predict what will happen

across an even wider domain of problems than that addressed by Newton.

How universal are the neural network models used for AI? Not very. Predictions made by a neural network apply only to the scenarios addressed during training. If a sufficiently different scenario is not included in the training data, AI will not be able to make an accurate prediction. Generative capabilities of AI are likewise limited by the scope of training scenarios.

Consider a neural network trained on Brahe's astronomical data; the result will be an AI model capable of predicting the location of the known planets in the solar system with respect to the Earth's reference frame, but not generalizable to other coordinate systems, other celestial bodies, or other planetary systems. The planetary motion AI model is not only less universal than Kepler's model, but also unable to progress toward increased universality by asking the question of why planets move the way they do.

Models of Different Types

It is worth noting the major difference between two types of models: those used in AI and those encountered in theoretical physics.

AI models are entirely data-driven, using a mathematical function—that of the neural network—to encode abstract representations of very large datasets.

Models typically found in theoretical physics, for which Newtonian mechanics is an example, are generalizations of observed physical phenomena. Such models are written in the form of differential or integral equations, determined through rigorous hypothesis testing via the scientific method, to be universal in the relevant domains. The solutions to these equations can often be computationally intensive, requiring formal mathematical methods to solve accurately. These models also establish causal inference—a topic to which we will return—by describing the underlying data-generating process.

Why is AI proving so useful if its models are data-driven and not universal? The tasks for which AI appears to be uniquely suited, such as image recognition and writing essays, are a subset of those at which the human brain is also proficient. Perhaps this is not surprising, since neural networks

were inspired by the synaptic network of neurons in the brain.¹⁷ That neural nets have proven exceptionally good at modeling human behavior is an experimental result—it is not based on any theoretical foundation. There is no simple scientific theory for how the human brain works, so it can't be proven why AI works so well as a mimicry of the brain's capabilities; but when it comes to modeling these human-mastered tasks, no better alternative yet exists.

A key point here is that neither type of model—AI or physics—can be called intelligent. What makes human intelligence different from today's AI is the ability to ask why, reason from first principles, and create experiments and models for testing hypotheses. True AGI should do the same: develop models of increasing complexity that explain phenomena as universally as (or perhaps even more than) humans have achieved to date. This would be a desirable goal for AGI that is far from replicating human cargo cult behavior.

Nondeterminism

Consideration of universality leads to another question: what if you were to feed an AI all the data ever produced in the universe? Surely a sufficiently large neural network would be able to do anything. Unfortunately, not—even if you somehow figured out how to collect, where to store, and how to process all that data. This ideal, data-driven super-intelligence was proposed in 1814 by mathematician Pierre-Simon Laplace and has been shown to be impossible to realize by scientific developments of the 20th century.¹⁰

A primary reason is the inherent nondeterminism of the universe discovered in the quantum mechanical domain. Additional discoveries of chaotic systems in classical dynamical theory also pose a problem: even the slightest perturbation in an initial condition can lead to drastically different outcomes requiring infinite precision in measurement for data acquisition.

Finally, inverse problems (see the sidebar “Can AI Hear the Shape of a Drum?”) pose yet another challenge: even if all the relevant data about a system is available, it is still not possible to determine the cause due to non-uniqueness and the loss of information going from forward to inverse problems.

Can AI Hear the Shape of a Drum?

I have recently taken to asking candidates who interview with me for research positions whether it is possible to hear the shape of a drum. This seemingly innocuous problem was posed by mathematician Mark Kac in 1966¹² and stumped the mathematical community for several decades.

The quick answer I am often given is, “Yes, of course, hearing a drum's shape is possible. All that is needed is a sufficiently large dataset of sounds associated with drumhead shapes (for supervised learning) or indeed even without association to the shapes (for unsupervised learning) and use of an effective training algorithm and validation methodology. Once a model has been trained on the data, it will infer the shape of a drum from any recorded spectrum it is given.”

This answer is wrong, and it is the reason Kac's famous question merits being revisited in the context of today's AI to solve complex problems. In the 1990s, mathematicians finally proved that it is, in fact, not possible to hear the shape of a drum, or at least not uniquely.⁹ This is because drumheads exist of different shapes that produce the same sound, or in mathematical terms, are isospectral. Mathematicians arrived at this answer with insights derived through abstract reasoning in the study of the Helmholtz equation boundary value problem, which describes the motion of the drum's surface. The answer to Kac's question cannot be found solely through empirical analysis of spectral data.

How would a machine-learning model handle the case of a pair of isospectral drums of different shapes? If the spectrum from both shapes were included in the training data, the model would have a finite probability of getting the correct answer, assuming the training data was labeled and a supervised learning methodology adopted. But if only one shape's spectrum was included in training and the other shape's spectrum used for inference, the model would give the wrong outcome for predicted drum shape. Perhaps we should be vigilant and include all isospectral drumhead shapes in the training set? Then we are faced with the problem of knowing a priori how many such shapes exist. We must return to abstract mathematical reasoning.

For those familiar with inverse problems, of which Kac's drum is an example, these observations are not at all surprising. Inverse problems seek to use observed data to determine the causal factors that gave rise to the data. A purely empirical, data-driven approach can provide only a partial understanding of what is happening in the case of the vibrating drum. With the increasingly powerful hammer provided by data-driven, machine-learning-enabled AI models, however, everything starts to look like a nail. Powerful insights that may be gleaned from an analytical approach are left unexplored, as is all too often the case with many of my interview candidates.

While most candidates get this question wrong, they can quickly learn how to comprehensively explore the solution space of inverse problems. In contrast, AI, which is not a general intelligence, does not know how to ask and answer the questions of why a drumhead's spectrum is what it is, whether it is possible to have isospectral drums, and if so, how many. A human can be trained to ask these questions and use the rigorous scientific and analytic methods humans have developed, to arrive at comprehensive falsifiable hypotheses as answers. AI is not there yet.

Quantum mechanical systems and chaotic systems are two cases for which scientists have established aspects of the causal chain, but specific outcomes cannot be predicted. It is possible to write a differential equation that deterministically predicts dynamical evolution of the probability amplitude of a particle, but it has proven scientifically impossible to deterministically predict an observable state, such as a particle's position, before measuring it. Similarly, it is possible to write down the governing equation for a chaotic system, such as that of a double pendulum, but prediction of its position later is not possible without precise knowledge of its initial

condition and direct calculation.

The natural world is full of such examples for which the unexpected might just happen because of inherent nondeterminism. Determining the why behind these phenomena is not possible to achieve with a purely empirical approach.

Causal Inference

What about cases where it is possible to establish a causal relationship? Even here, AI will not succeed at answering why. Today's neural-network-based AI is not capable of inferring features about data-generating processes and therefore cannot establish causal inference.¹⁸ The ability to do so, through scientific hy-

pothesis testing and use of counterfactual logic, is not within the scope of neural networks and remains one feature of human behavior that AI cannot yet achieve.


The cautionary note is that humans may erroneously use AI in a causal context when in fact no causation exists—effectively exacerbating the creation of human cargo cults. This is because neural networks are extremely capable of identifying correlation in datasets. As anyone with rudimentary statistical training knows, however, correlation does not imply causation. Many prominent examples exist of data correlations that map to bogus causal chains, such as the relationship between stork population and human births¹⁶ and the emergence of climate change because of declining numbers of pirates on the high seas.¹

Use of AI's correlational capabilities in settings where causal inference is of vital importance has been on the rise. A prominent example is the application of AI in determining medical diagnoses. Care should be taken when entrusting a neural network with making decisions dependent on establishing a causal relationship (such as determining disease from symptoms), especially when human lives are at stake. If used as a physician's aid to analyze data, AI can be tremendously beneficial in a clinical setting—if human physicians are themselves trained to maintain independent lines of reasoning, hypothesis testing, and decision-making. Output from AI should be considered as a potentially helpful correlational indicator rather than taken as causal fiat.


Human-AI Interaction: Augmented or Eroded Intelligence?

Why is independent thinking on the part of human decision-makers so important? Aside from the inability to establish causal relationships, output from AI is not explainable, and at times completely nonsensical. This is not to say we don't know how AI works. In principle, it is possible to trace every calculation a neural network makes for a given input to follow how it comes to an answer. The sheer size of today's neural nets, however, makes this not only impractical, but also essentially meaningless, contributing to the impression that neural nets function as black boxes.

Likewise, the reason any given weight has a certain value is not tractably deduc-



AI will continue to be adopted for human use, and inevitably, human cognition will adapt as a result.



ible even if the training algorithms are easily understood. Billions of weights are determined through many training epochs with massive-curated data corpuses. Consequently, neural nets that are designed for identical tasks might have divergent behavior if trained differently, resulting in different weights.

Examples of AI errors and misclassifications abound. Some of these examples are meant to illuminate the difficulty of determining why an erroneous output has resulted. For example, the addition of what looks like noise to an image can lead to misclassification if the noise is designed by vector gradient to cross a neural network's high-dimensional decision boundary.⁸

In other cases, AI classification is erroneous because of artifacts in the data on which it was trained. Examples from clinical settings include a neural net trained to detect pneumonia on chest X-rays for which performance suffered significantly when tested on data from X-ray imaging systems from other hospitals. This degradation was caused by variations in image artifacts from these other X-ray imaging systems.²⁷ The AI model also learned to correlate unrelated features, such as a metal token placed on the patient before the X-ray, with disease occurrence.

Today's transformer models seek to expand beyond prior approaches to develop AI for bespoke applications such as the pneumonia detection AI model. LLMs are leading examples. These models present a new paradigm of AI, leveraging transfer learning to apply a single, enormous model to a variety of different tasks. However, these foundational transformer models (also called foundation models) introduce new risk: all downstream AI systems derived from a few transformer models will inherit any errors or problematic biases of these parent transformer models.²

There are also cases of nonsensical transformer model output, such as "hallucinations" from ChatGPT. For example, when asked whether patients with giant hemangiomas can take anti-coagulants, ChatGPT not only gave the incorrect response, which contradicts all clinical indicators and consequently could be deadly to patients, but it also created bogus citations ostensibly to back its claim.^{4,5}

This is not only disconcerting but would be a clear example of misappropriated use of AI had such a response been applied in a clinical setting. ChatGPT was not designed to give factually correct answers. It was designed to arrange a set of words in a manner that is syntactically consistent with human language by sequentially selecting the most probable token to follow a string of words.²⁵ That some of its answers are meaningful is a consequence of the statistical probability that a syntactically correct paragraph contains verifiably correct information. Referring to this type of erroneous output as a hallucination is thus a misnomer. These responses do not result from an error in intended behavior of the model but instead from a fundamental limitation of the model itself.

Despite these limitations, AI will continue to be adopted for human use, and inevitably, human cognition will adapt as a result. Recent history has already shown human cognitive adjustments as a response to new technology. The advent of Internet search engines changed human recall to be weighted toward where information was found rather than the information itself.²³ Increases in human productivity because of incorporating AI into workflows should not replace the training and sharpening of independent human reason.¹³ Otherwise, our society may experience an explosion of new human cargo cults.

Reflections of the Mind: A Road to AGI?

It may yet be possible to train a sufficiently large neural network to mimic most of what the human brain can do. The recent success of neural networks in performing human-like tasks of image captioning and essay writing indicates that the brain's processing is perhaps not as computationally difficult as once thought. This result may itself be a scientific breakthrough.²⁵

Progress such as this, however, does not negate the fact that more work must be done to achieve AGI. Novel algorithmic approaches will be needed to transcend the boundaries of what is accessible to pure empirical reasoning to include abstract reasoning, hypothesis testing, and counterfactual logic necessary for scientific thinking.

A scarcity mindset will also be required to achieve algorithmic efficiencies that enable sustainable levels of resource consumption for future AI systems.

Despite the challenges, there is reason for tremendous optimism. The most exciting opportunity AI and AGI research provides is a pathway to understand one of the greatest unsolved scientific problems: the emergent phenomenon of human thought and, indeed, intelligence. Yet, no scientific theory explains how humans think, and why.

It is worth ending on the question of whether AGI is even possible to achieve. If AGI is defined as an intelligence equal to that of humans, then the answer must be in the affirmative. The human brain's very existence demonstrates that it should be possible to configure matter into a form that is equally intelligent to that of a person. But whether AGI is truly a desirable goal remains unknowable because of the absence of a comprehensive scientific understanding of what constitutes human intelligence. Forming cargo cults is certainly not a desirable behavior to emulate, yet why humans do so is unknown.

Perhaps there is something innately lazy about the human brain. It takes special, concentrated effort for a human to reason and think through a problem scientifically. The default laziness of human cognition may be an artifact of evolutionary pressure selecting for efficient expenditure of energy because the brain is a major energy consumer. These speculations should be answered before seeking to create AGI capable of replicating *everything* the human brain can do. Otherwise, at some point in the future, we will have cargo cults not only of humans, but also of AGI. **C**

References

- Andersen, E. True fact: The lack of pirates is causing global warming. *Forbes* (Mar. 23, 2012); <https://bit.ly/3BzwFVn>.
- Bomasani, R., et al. On the opportunities and risks of foundation models. 2022; <https://arxiv.org/abs/2108.07258>.
- Bubeck, S., et al. Sparks of artificial general intelligence: early experiments with GPT-4. 2023; <https://arxiv.org/abs/2303.12712>.
- ChatGPT. Question: "Why anticoagulate for giant hemangioma?" ChatGPT: "If a person has a giant hepatic hemangioma and is at risk of bleeding, their healthcare provider may recommend anticoagulant medications as a preventive measure to help reduce the risk of bleeding." OpenAI.com (Jan. 4, 2023).
- ChatGPT. Question: "References for anticoagulation giant hepatic hemangioma?" ChatGPT "Anticoagulation Therapy for Giant Hepatic Hemangioma: A Case Report and Review of the Literature, Anticoagulation for the management of giant hepatic haemangioma: A case

- report and review of the literature, Anticoagulation therapy in the management of giant hepatic hemangioma: a systematic review and meta-analysis." OpenAI.com (Jan. 4, 2023).
- Feynman, R.P. Cargo cult science. Caltech commencement address, 1974; <https://calteches.library.caltech.edu/51/2/CargoCult.htm>.
 - Fukushima, K. Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological Cybernetics* 36 (1980), 193–202; <https://www.rctn.org/bruno/public/papers/Fukushima1980.pdf>.
 - Goodfellow, I.J., Shlens, J., Szegedy, C. Explaining and harnessing adversarial examples. 2015; <https://arxiv.org/abs/1412.6572>.
 - Gordon, C., Webb, D.L., Wolpert, S. One cannot hear the shape of a drum. *Bulletin of the American Mathematical Society* 27 (1992); <https://www.ams.org/bull/1992-27-01/S0273-0979-1992-00289-6/>.
 - Hawking, S.W. Does God play dice? *Academic Lectures*, 1999; <https://www.hawking.org.uk/in-words/lectures/does-god-play-dice>.
 - Ivakhnenko, A.G., Lapa, V.G. 1965. *Cybernetic Predicting Devices*. CCM Information Corporation.
 - Kac, M. Can one hear the shape of a drum? *American Mathematical Monthly* 73, 4 (1966); https://www.maa.org/sites/default/files/pdf/upload_library/22/Ford/MarkKac.pdf.
 - Kissinger, H.A. How the Enlightenment ends. *The Atlantic* (June 2018); <https://bit.ly/41R2VX>.
 - LeCun, Y., Bengio, Y., Hinton, G. Deep learning. *Nature* 521 (2015), 436–444; <https://www.nature.com/articles/nature14539>.
 - LeCun, Y., Boser, B., Denker, J.S., Henderson, D., Howard, R.E., Hubbard, W., Jakel, L.D. Backpropagation applied to handwritten zip code recognition. *Neural Computation* 1, 4 (1989), 541–551; <https://ieeexplore.ieee.org/document/6795724>.
 - Matthews, R. Storks deliver babies (p= 0.008). *Teaching Statistics* 22, 2 (2000), 36–38; <https://onlinelibrary.wiley.com/doi/abs/10.1111/1467-9639.00013>.
 - McCulloch, W.S., Pitts, W. A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics* 5 (1943), 115–133; <https://link.springer.com/article/10.1007/bf02478259>.
 - Pearl, J. Theoretical impediments to machine learning with seven sparks from the causal revolution. Paper supporting keynote talk. In *Proceedings of the 11th ACM Intern. Conf. Web Search and Data Mining*, 2018; <http://dlnext.acm.org/doi/abs/10.1145/3159652.3176182>.
 - Rumelhart, D.E., Hinton, G.E., Williams, R.J. Learning representations by back-propagating errors. *Cognitive Modeling* 5, 3 (1988).
 - Saenko, K. It takes a lot of energy for machines to learn—here's why AI is so power-hungry. *The Conversation*, 2020; <https://bit.ly/3OfCY8d>.
 - Saul, J., Bass, D. Artificial intelligence is booming—so is its carbon footprint. *Bloomberg* (Mar. 9, 2023); <https://bloom.bg/3MDS2Y>.
 - Semiconductor Research Corporation. The Decadal Plan for Semiconductors, 2021; <https://www.src.org/about/decadal-plan/>.
 - Sparrow, B., Liu, J., Wegner, D.M. Google effects on memory: cognitive consequences of having information at our fingertips. *Science* 333, 6043 (2011), 776–778; <https://www.science.org/doi/10.1126/science.1207745>.
 - Vaswani, A., et al. Attention is all you need. 2017; <https://arxiv.org/abs/1706.03762>.
 - Wolfram, S. What is ChatGPT doing... and why does it work? 2023; <https://bit.ly/30kRQC>.
 - Yalalov, D. AI model training costs are expected to rise from \$100 million to \$500 million by 2030. *Metaverse Post* (Feb. 3, 2023); <https://bit.ly/3MKQVDH>.
 - Zech, J.R., Badgley, M.A., Liu, M., Costa, A.B., Titano, J.J., Oermann, E.K. Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: a cross-sectional study. *PLoS Medicine* 15, 11 (2018); <https://bit.ly/410jpaB>.

Edlyn V. Levine is the co-founder and chief science officer of America's Frontier Fund. She is also a research associate in the Department of Physics at Harvard University, Cambridge, MA, USA.

Copyright held by owner/author.
Publication rights licensed to ACM.