



What the Rise of AI Means for Narrative Studies: A Response to “Why Computers Will Never Read (or Write) Literature” by Angus Fletcher

ABSTRACT: The role of AI in narrative studies is not a question of if but of when and of how we humans prepare for such a future. The if claim is addressed with a detailed rebuttal to Angus Fletcher’s “Why Computers Will Never Read (or Write) Literature.” A counter-argument based upon key AI concepts, the historical progress of AI, and landmark failures and breakthroughs brings readers up to date on the current state of AI as it relates to narrative studies. Numerous examples explain why the cycle of AI winters and springs is now broken, and there is a new global AI arms race. Scholars now have a windfall of increasingly sophisticated, multi-million dollar models that can analyze and generate narrative. Still, in light of the inherent complexity of natural language and the current limitations of even these state-of-the-art AI models, a human-in-the-loop is essential for the foreseeable future. We attempt to allay common yet misplaced concerns by reasserting the centrality of the human scholar to guide and interpret while using these tools. Leveraging these new AI models will yield new insights for narrative studies, and this important work will include shaping the language of fairness, equality, and ethics of models that increasingly impact the lives of billions. We invite narrative scholars to participate in this growing interdisciplinary movement that chooses active engagement over passive critique.

KEYWORDS: *Narrative Studies, Artificial Intelligence, Digital Humanities, Story Generation, Middle Reading*

Jon Chun was co-founder and CEO of the world's largest privacy/anonymity startup in Silicon Valley and Director of Development at a Fortune 500 Company before teaching at Kenyon College. Recent research includes his new tool, SentimentArcs, a self-supervised ensemble method for sentiment analysis of narrative. Chun and Elkins also publish on the new AI Transformer models for language generation.

Katherine Elkins is Associate Professor of Humanities and Comparative Literature at Kenyon College. She is founding Co-Director, alongside Jon Chun, of the KDH lab, which leverages AI advances for interdisciplinary research. Recent books include *Philosophical Approaches to Proust's In Search of Lost Time* (Oxford UP, 2022) and *The Shapes of Stories: Sentiment Analysis for Narrative* (Cambridge UP, February 2022).

WILL AI EVER BE ABLE TO read or write literature? This question was answered in the negative by Angus Fletcher in this journal in his article "Why Computers Will Never Read (or Write) Literature: A Logical Proof and a Narrative." We have been researching AI and narrative for several years now and strongly disagree with Fletcher's conclusion. His logical proof fundamentally misrepresents AI as being limited to a particular hardware, and that hardware as being limited to a particular gate-based logic. His narrative proof stops well short of the AI breakthroughs of the last decade that include transformative new language models. We start with a fuller account of the story of the rise of AI in order to appreciate just how well AI can both generate and analyze stories.¹ This lays the groundwork for a more informed critique of Fletcher's proofs followed by our own reservations about what the rise of AI means for narrative studies.

The Story of the Rise of AI

The history of AI has been marked by cycles of optimistic promises followed by disappointing failures in a pattern known as "AI springs and winters." The term 'Artificial Intelligence' was coined at a conference at Dartmouth in 1956 by an interdisciplinary team of ten leading researchers across fields like computer science, math, and cognitive science. The funding proposal for the two-month conference claimed the team would make significant advances in defining every aspect of human intelligence so that a machine could simulate it (McCarthy). This hubris is understandable given Moravec's Paradox (Sarkar). The reasoning was that if early digital computers could crunch complex scientific calculations with a speed and accuracy far beyond the abilities of the best human experts, they could easily master everyday cognitive tasks average humans find trivial, such as perception and language.

What makes us so sure that the current hype surrounding AI will not result in yet another AI winter? Unlike previous cycles, the commitment to AI development is no longer dependent upon the vagaries of government and academic funding. Current AI systems are increasingly commercially viable and continue to evolve to achieve superhuman performance in a growing number of task once thought the exclusive province of humans, with poker as just one example (Brown). Indeed, the rise of powerful AI systems has created something of an AI arms race in both the public and private sector (Ghi). This current environment brings unique opportunities for

researchers in narrative studies to benefit from the many AI advances through the rise of free and open source software.

Our work on narrative is most directly related to the AI subfields of Natural Language Processing (NLP) and Natural Language Generation (NLG). The initial dominant attitude towards language was that words were just symbols that could be logically manipulated according to rules of some unknown Universal Grammar (Cowie). For decades, leaders in computer science like Noam Chomsky sought to elucidate this universal grammar, believed to be evolutionarily written in the neural pathways of our brains. This approach proved unworkable, and it was soon eclipsed by more successful statistical methods in the 1980s which, in turn, have been eclipsed by the deep neural networks (DNN) of the last decade. Our work focuses on the most current large-scale DNN language models called Transformer Models (for example, BERT and GPT-2) (Vaswani).

This lengthy background is key to understanding why Fletcher's critique is misdirected at largely defunct older Chomskyan Symbolic approaches. He overlooks current AI approaches that rely on carefully-crafted white-box statistical models (termed GOFAI or Good Old-Fashioned AI) and massive black-box statistical universal approximation engines (DNN or Deep Neural Networks). These data-driven statistical methods arose because they achieved practical levels of success in NLP tasks (for example, Apple Siri, Google Translate) that decades of Symbolic research failed to achieve (Sejnowski). In fact, within the growing intellectual exchanges between neuroscientists, cognitive scientists and AI researchers, some have posited that at one level of abstraction the human brain performs complex statistical inference (Kriegeskorte).

How Well Can AI Tell Stories

The question of whether AI can write narrative has already been answered to some degree depending on the definition of "story." For over a decade, companies like Narrative Science, Automated Insights, and Yseop have used traditional template and rules-based Natural Language Generation (NLG) to create newspaper articles and financial reports for companies like the Associated Press and Société Générale (Peiser). Last year, GPT-3 generated text for a top article in a leading tech blog (Hao) and mainstream media newspaper (GPT3). In general, humans cannot distinguish between AI or human-generated content in these best-case scenarios. Children can hear interactive stories on Amazon Alexa from companies like Disney, and teens can play immersive video games with emerging narrative elements.

Before recent developments, the stories generated with an older template-rule AI came with key limitations: very narrow focus, inability to generalize creatively, and labor-intensive human engineering of canned keyword prompts and scripts. This is why the public release of new large-scale Transformer language models like GPT and BERT has been greeted with both excitement and trepidation in leading cultural magazines like *The New Yorker* (Marche) and *The Atlantic* (DiResta). For the first time, NLG has demonstrated the ability to overcome critical limitations of the

traditional template-rule approach and generate text that is at times indistinguishable from the best human writers (Elkins and Chun, “Turing Test?”).

Transformer models like GPT and BERT learn language and storytelling much like humans learn the craft of good writing, through repeated exposure to good examples and practice. Current state-of-the-art Transformers are trained on billions of words that can include Gutenberg novels (narrative), Wikipedia (common sense), and upvoted Reddit threads (reasoning and debate). Through massive complex statistical analysis, Transformers learn to indirectly model the hidden latent patterns universal to human language including not only syntax, but also semantics and pragmatics. Models can be further fine-tuned to learn specific genres, writing styles, and formal aspects of storytelling (Gwern).

One of the most impressive aspects of Transformer models is that one single, universal model has surpassed generations of many disparate older NLP models that were specifically hand-crafted to solve a single task like foreign language translation or sentiment analysis. Instead, a Transformer model is fed well-structured prompts to elicit similarly structured yet open-ended responses. For example, if you input several lines of dialogue formatted like a play or TV script, you will be given back a continuation of dialogue not only in form but replicating content and style. If fed the first several lines of a creative story, Transformers can generate variations that continue the narrative in both familiar and unexpectedly creative directions. As of March 2021, OpenAI reported tens of thousands of developers working on over 300 apps that are generating 4.5 billions words each day using GPT-3 (OpenAI).

Of course, Transformers come with their own limitations and precautions that require careful training and close human editorial supervision. As we discuss elsewhere, even the older GPT-2 model performs fairly well when tasked with writing poetry, short stories, and scripts. Long-form narrative remains a challenge. Even with its successes, there are a variety of problems that point to specific limitations. Common failure modes include repetition, grammatical errors, leaps or omission in turn-based dialogue, invented facts, illogical reasoning, bias and racism, and factual errors stemming from a lack of real-world knowledge like cooking or physics. These limitations arise because Transformers are a probabilistic model that (a) learn from massive transcripts of human generated language and (b) create language by trial and error rather than relying on pre-programmed knowledge and fixed rules. For example, although not explicitly taught how to add numbers or solve SAT-like analogies, Transformers indirectly acquire the rudimentary ability to do so while also exhibiting errors similar to that of a young child (Brown).

For the past two years, we have fine-tuned many GPT-2 models to generate novels, screen plays, improv dialogue (Dennen), poems, song lyrics, political tracts, news articles, and much more (Elkins and Chun, “Turing Test?”). Our analysis over this time has led us to a very different outlook on the current and future role of AI in literary studies. Our first impression upon generating stories with GPT-2 is how much better it is in every metric compared to the “word salads” generated by previous generations of NLG (Sharp). GPT-2 learns English by training a language model with up to 1.5 billion parameters on a corpus of 8 million web pages filtered for quality (Radford). By extracting latent statistical correlation within the training corpus,

GPT-2 models can learn not only linguistic semiotics but also even longer-range abstract elements of narrative like plot and theme. Performance continues to improve as these Transformer models grow: the latest GPT-3 model that we're exploring has 175 billion parameters trained on 499 billion words at an estimated cost of \$4.6 million (Brown). Advances continue with the Beijing Academy of AI releasing a 1.75 trillion parameter Transformer model WuDao in June 2020.²

AI Narrative Analysis or Middle Reading

Over the past several years, we have also been experimenting with AI approaches to narrative analysis using a variety of AI techniques including symbolic, statistical machine learning, deep neural nets, probabilistic graphical models, and others. Unlike the more well-known “distant reading” made famous by Franco Moretti that “reads” over large corpora (Moretti), we focus on what we call “Middle Reading” on the scale of a single novel or author. In our lab, we study how AI NLP techniques can be used to surface formal elements of narrative, correlates of plot, thematic interweaving, and character development.³

While one can easily mistake many (though certainly not all) AI-written stories for those written by a human, a definition of machine “reading” diverges more fundamentally from our human experience. This is especially true if one understands reading as an embodied interaction with a text that creates a complex, emergent experience. Machines do “read” text, but what we mean by this is quite different from what we mean when we talk about human reading. Nonetheless, machine reading offers interesting insights for narrative studies, and we find it augments our own understanding of narrative from a human-centered reading experience.

Fletcher points to the question “why?” as fundamental to exploring our experience of literature, and we agree. But unlike Fletcher, we investigate how AI tools can help us understand questions that ask “why?”; for example, why narrative may produce emotional effects. To the extent that it is scientific, this AI assisted approach has more in common with exploratory data analysis (Tukey)—which lets the data speak for itself—than with traditional scientific hypothesis testing. Although both analysis and generation are demanding tasks that require expert human supervision, analysis in particular presupposes an extensive specialized training. Story generation is far more permissive in admitting wide-ranging and creative forms at which Transformers excel.

While challenges remain, the major failings of current state-of-the-art AI narrative generation are areas of active research, and the insights we're discovering in our lab contribute valuable domain expertise. This is especially true for addressing the current limitations of generating long-form narrative, and we are actively researching why novels pose unique challenges in comparison to other literary forms. The cross-pollination between AI researchers and traditional narrative scholars can lend greater insight into narrative structure, thereby grounding work being done in theory with well-established practices (Castricato).

A Narrative Disproof of Fletcher's Argument

Let's now return to the narrative Fletcher tells. He offers many stories, but each stops short of the ending. In the case of Bertrand Russell, he pauses the story before we learn that Russell later retreated from his logicism claim that all math is logic and pure reason (Wagner)). This was in part a result of the works by Wittgenstein, Poincaré, and Gödel, who proved the inherent limitations of logical systems Russell initially favored. Russell instead wrote in his autobiographical work *My Philosophical Development*: "Wittgenstein maintains that logic consists wholly of tautologies. I think he is right in this although I did not think so until I read what he had to say on this subject" (88).

In his critique of AI, Fletcher stops his story in 1969 near the peak of symbolic logic dominance. AI hubris was at an all-time high, and Marvin Minsky had proclaimed two years earlier, "Within a generation, I am convinced, few compartments of intellect will remain outside the machine's realm—the problems of creating 'artificial intelligence' will be substantially solved" (2). Over the half century since then, the field of AI has experienced at least four more periods of AI winters and summers and has advanced in compute power, created massive training datasets, and invented entirely new algorithmic solutions to solve formerly intractable problems. Today's superhuman AI can solve problems in fuzzy domains such as vision and language that older paradigms like symbolic logic could not even formally define.

Finally, Fletcher stops with distant reading as the end of the story of digital humanities when in fact, many of us are working to expand the field in new and quite different directions. All the same, we share his belief that literary studies and scientific approaches should not be seen as opposites, and that literary studies must and should engage with the emerging questions inherent to the rise of AI. We also agree in spirit, if not in form, with Fletcher's reservations regarding the potential misuse, abuse, and exaggerated claims around AI, and especially with his identification of the novel as posing a unique challenge.

Our experience using AI, however, makes us optimistic that it can deepen our understanding of human creativity and language, and that our insights may even have much to offer researchers developing AI language models. It's a common practice in AI critique—which Alan Turing first described as "Arguments from Various Disabilities"—to draw a line in the sand beyond which one insists AI could never possibly advance (447). But as we've seen, those lines must constantly be redrawn. There is much to learn from our creations—our machines that can both "read" and write stories.

A Logical Disproof of Fletcher's Argument

FLETCHER'S LOGICAL PROOF THAT COMPUTERS CANNOT READ (OR WRITE) LITERATURE

1. Literature has a rhetorical function

2. Literature's full rhetorical function depends on narrative elements
3. Narrative elements rely on causal reasoning
4. Causal reasoning cannot be performed by machine-learning algorithms because those algorithms run on the CPU's Arithmetic Logic Unit, which is designed to run symbolic logic, and symbolic logic can only process correlation

QED: Computers cannot perform the causal reasoning necessary for learning to use literature.

The final claim of Fletcher's original proof (#4) above rests upon several fundamental misunderstandings of computer science and the state of AI today. Our disproof will focus on this claim, which is in fact a composition of several individual claims.

- 4(a). Current AI run on CPU Arithmetic Logic Units
- 4(b). CPU Arithmetic Logic Units (ALU) can (only) process Symbolic Logic
- 4(c). Symbolic Logic can only process correlation
- 4(d). Correlation is not causation

This first sub claim 4(a) is logically false: AI algorithms can be run on any computational architecture that is Turing complete (Immerman). This includes not only the familiar CPU/ALU-based von Neuman architecture Fletcher cites, but also other architectures including those based on quantum properties and data flow. However, in practice this claim is largely true as most AI runs on some type of von Neuman architecture that includes CPUs as well as GPUs and TPUs.⁴

The most problematic claim is 4(b): CPU ALU (equated with current AI) can (only) process Symbolic Logic. AI algorithms can run on any CPU ALU von Neuman architecture with varying degrees of efficiency. These include symbolic logic, statistical machine learning, deep neural networks, cellular automata, probabilistic, and causal graphical models (Schölkopf) among others. In fact, Judea Pearl received his Turing Award "for fundamental contributions to artificial intelligence through the development of do-calculus for probabilistic and causal reasoning" ("Judea Pearl"). As of May 2021, there are 158 open source software repositories on GitHub under the topic 'causality' that run on von Neumann CPU ALU architectures ("Topic: Causality"). .

The correction of claim 4(b) renders claim 4(c) less relevant to the overall argument. Given that von Neuman architectures (CPU/ALU) are Turing complete universal compute platforms, they can process any statistical relationship and not just correlations. The limitations of causal models like Pearl's do-calculus relate to practical implementation details and not computer architecture or inherent theoretical limitations (Greenland).

Reservations

What we mean when we say computers can read and write literature is not as clear as either critics like Fletcher or AI enthusiasts who trumpet the latest advances claim.

While we share many of Fletcher's concerns, we have a more optimistic outlook on how recent advances in AI can be integrated with traditional scholarship to advance research on narrative. This outlook is firmly based in both our research and an understanding of current AI developments.

Our two biggest reservations working at the nexus of AI and narrative are theoretical and existential, not technological. AI NLP techniques continue to rapidly evolve and advance. The most salient question is: if language, narrative and art can be modeled by machines, does that imply human thought and creativity are simply the emergent properties of vast statistical properties of atoms and neuronal spike trains? Where is humanity in such radically reductionist materialism?

This leads to our second, perhaps more intractable, concern about the use of AI to analyze and generate narrative. The current trajectory of AI development makes it clear that the concern should not be *if* AI can analyze and generate narrative, but rather *how* we as humans prepare for such a future. This concern is not unique to narrative studies, as AI achieves superhuman ability in a growing number of human domains from playing Go to protein folding to medical diagnosis. A consolation to the narrative scholar is that many AI experts hold that "solving" language may be one of the last milestones on the route to Artificial General Intelligence (AGI)—an AI that surpasses humans by every conceivable measure. Not only do most AI experts believe that AGI lies nowhere in the immediate future (Grace), but if AGI is indeed achieved, humanity will face far greater problems than AI narrative analysis and generation (Oxford).

While these two concerns, theoretical and existential, remain unresolved, we choose the path of active engagement—rather than passive critique—for several reasons. First, as a tool, we've found AI technologies provide leverage and insight to augment traditional human scholarship, not replace it. Second, we see ourselves as part of a larger interdisciplinary movement to use technology in ways that improve our understanding of human experience and creativity. Perhaps most importantly, as AI continues to expand its reach, experts from all domains including the humanities should help articulate a vision for the future based upon both the wisdom of our field and an informed understanding of the promises and perils of this technology.

Endnotes

1. While Fletcher uses the term "literature," we focus here on narrative. Elsewhere we discuss poetry and theater in more detail (2020, 2021).
2. Beijing Academy of Artificial Intelligence, June 2021 and <https://www.baai.ac.cn/>
3. For our work on plot, emotion, and narrative arc, consult our early work ("Can Sentiment Analysis") and a forthcoming work in Cambridge UP's Digital Literary Studies series, *The Shapes of Stories*. For the past year, our lab has been working on applying cutting-edge AI tools to analyze other components of narrative (forthcoming).
4. The development of more powerful AI models requires more powerful computers. In addition to the dramatic increase in CPU power over decades, Graphical Processing Units and Tensor Processing Units have been used to parallelize processing and substantially increase the ability to train and implement very large AI models.

Works Cited

- Association for Computing Machinery. “Judea Pearl—A.M. Turing Award Laureate.” [amturing.acm.org](https://amturing.acm.org/award_winners/pearl_2658896.cfm), 2011. https://amturing.acm.org/award_winners/pearl_2658896.cfm.
- Broan, Noam and Tuomas Sandholm. “Superhuman AI for Multiplayer Poker.” *Science* 365.6456 (July 2019): 885–90.
- Brown, Tom B., Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, et al. “Language Models Are Few-Shot Learners.” [Arxiv.org](https://arxiv.org/abs/2005.14165), 28 May 2020. <https://arxiv.org/abs/2005.14165>.
- Castricato, Louis, Stella Biderman, Rogelio Cardona-Rivera, and David Thue. “Towards a Model-Theoretic View of Narratives.” [ArXiv.org](https://export.arxiv.org/pdf/2103.12872), 23 March 2021. <https://export.arxiv.org/pdf/2103.12872>.
- Cowie, Fiona. “Innateness and Language (Stanford Encyclopedia of Philosophy).” [Stanford.edu](https://plato.stanford.edu/entries/innateness-language/), 16 Jan. 2008. <https://plato.stanford.edu/entries/innateness-language/>.
- Dennen, Jim, Jon Chun, and Katherine Elkins. “Can a Machine Play? AI, Theatre, and Improvisation.” *Arts at Denison*, 25 Jan. 2021. <https://tinyurl.com/39kxddy9>.
- DiResta, Renée. “The Supply of Disinformation Will Soon Be Infinite.” *The Atlantic*, 20 Sept. 2020. <https://www.theatlantic.com/ideas/archive/2020/09/future-propaganda-will-be-computer-generated/616400/>.
- Elkind, Landon D C, and Gregory Landini. *The Philosophy of Logical Atomism. A Centenary Reappraisal*. New York: Springer International Publishing, 2018.
- Elkins, Katherine, and Jon Chun. “{K}DH Colab & Research Fellows.” *Kenyon Digital Humanities Collaboratory*, 1 Jan. 2017. <https://www.kenyon.edu/digital-humanities/kdh-colab/>.
- . “Can GPT-3 Pass a Writer’s Turing Test?” *Journal of Cultural Analytics* 9 (2020): <https://doi.org/10.22148/001c.17212>.
- . “Can Sentiment Analysis Reveal Structure in a Plotless Novel?” *ArXiv 1910.01441 [Cs]* (31 Aug. 2019).
- . *The Shapes of Stories*. Cambridge UP, forthcoming.
- Fletcher, Angus. “Why Computers Will Never Read (or Write) Literature: A Logical Proof and a Narrative.” *Narrative* 25.3 (January 2021): 1–28.
- Ghi, Trung, and Abhishek Srivastava. “The Global AI Arms Race—How Nations Can Avoid Being Left Behind.” *Arthur D. Little*, March 26, 2021. <https://www.adlittle.com/en/GlobalAI>.
- GPT-3. “A Robot Wrote This Entire Article. Are You Scared Yet, Human? | GPT-3.” *The Guardian*, 8 Sept. 2020. <https://www.theguardian.com/commentisfree/2020/sep/08/robot-wrote-this-article-gpt-3>.
- Grace, Katja, John Salvatier, Allan Dafoe, Baobao Zhang, and Owain Evans. “When Will AI Exceed Human Performance? Evidence from AI Experts.” [arXiv.org](https://arxiv.org/abs/1705.08807) 1705.08807, 2017. <https://arxiv.org/abs/1705.08807>.
- Greenland, Sander, and Mohammad Ali Mansournia. “Limitations of Individual Causal Models, Causal Graphs, and Ignorability Assumptions, as Illustrated by Random Confounding and Design Unfaithfulness.” *European Journal of Epidemiology* 30.10 (February 17, 2015): 1101–10.
- Hao, Karen. “A College Kid Created a Fake, AI-Generated Blog. It Reached #1 on Hacker News.” *MIT Technology Review*, 14 Aug. 2020. <https://www.technologyreview.com/2020/08/14/1006780/ai-gpt-3-fake-blog-reached-top-of-hacker-news/>.
- Immerman, Neil. “Computability and Complexity.” *Stanford Encyclopedia of Philosophy*, (24 June 2004) 18 Oct. 2021. <https://plato.stanford.edu/entries/computability/>.

- Kriegeskorte, Nikolaus, and Pamela K. Douglas. "Cognitive Computational Neuroscience." *Nature Neuroscience* 21.9 (September 1, 2018): 1148–60.
- Marche, Stephen. "The Computers Are Getting Better at Writing." *The New Yorker*, 30 April 2021. <https://www.newyorker.com/culture/cultural-comment/the-computers-are-getting-better-at-writing>.
- Minsky, Marvin Lee. *Computation: Finite and Infinite Machines*. Englewood Cliff: Prentice-Hall, 1967.
- McCarthy, John, M.L. Minsky, N. Rochester, and C.E. Shannon. "A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence." *Website of Professor John McCarthy*, (21 Aug. 1955) <http://jmc.stanford.edu/articles/dartmouth/dartmouth.pdf>.
- Moretti, Franco. *Distant Reading*. London: Verso, 2013.
- NeurIPS. "Causality-NeurIPS 2020." *Conference on Neural Information Processing Systems 2020*. <https://www.cmu.edu/dietrich/causality/neurips20ws/> (accessed 1 May 2021).
- OpenAI. "GPT-3 Powers the Next Generation of Apps." *OpenAI Blog*, March 21, 2021. <https://openai.com/blog/gpt-3-apps/>.
- Oxford University. "The Ethics in AI Institute." *Schwarzman Cewntre*, 16 Feb.2021. <https://www.schwarzmancentre.ox.ac.uk/ethicsinai>.
- Peiser, Jaclyn. "The Rise of the Robot Reporter." *The New York Times*, 5 Feb. 2019. <https://www.nytimes.com/2019/02/05/business/media/artificial-intelligence-journalism-robots.html>.
- Radford, Alec. "Better Language Models and Their Implications." *OpenAI*, 14 Feb. 2019. <https://openai.com/blog/better-language-models/>.
- Russell, Bertrand, and Alan Wood. *My Philosophical Development*. New York: Simon And Schuster, 1959.
- Sarkar, Tirthajyoti. "Why Math Is Easy for AI but Gardening Is Not: Moravec's Paradox." *Towards Data Science*, 9 March 2020. <https://towardsdatascience.com/why-math-is-easy-for-ai-but-gardening-is-not-moravecs-paradox-99994b201d10>.
- Schölkopf, Bernhard. "Causality for Machine Learning." *ArXiv:1911.10500*, 23 Dec. 2019. <https://arxiv.org/abs/1911.10500>.
- Sejnowski, Terrence. "The Unreasonable Effectiveness of Deep Learning in Artificial Intelligence." In *Proceedings of the National Academy of Sciences of the United States of America*, 177:30033. Washington, DC: National Academy of Sciences of the United States of America, 2020. <https://www.pnas.org/content/117/48/30033>.
- Sharp, Oscar, and Ross Goodwin. "Sunspring | a Sci-Fi Short Film Starring Thomas Middleditch." YouTube video. *Ars Technica*, June 9, 2016. <https://www.youtube.com/watch?v=LY7x2Ihqjmc>.
- "Topic: Causality." *GitHub*. <https://github.com/topics/causality>. (Accessed 1 May 2021).
- Tukey, John W. "We Need Both Exploratory and Confirmatory." *The American Statistician* 34.1 (February 1980): 23–25.
- Turing, Alan M. "COMPUTING MACHINERY and INTELLIGENCE." *Mind* 59. 236 (1950): 433–60.
- Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aiden Gomez, Łukasz Kaiser, and Illia Polosukhin. "Attention Is All You Need." In *Proceedings of NIPS 2017 Symposium on Interpretable Machine Learning*. Cornell University: 2017. <https://arxiv.org/abs/1711.09889>.
- Wagner, Roi. *Making and Breaking Mathematical Sense: Histories and Philosophies of Mathematical Practice* Princeton: Princeton Univ. Press, (1973) 2017.