

Henry Nelson
Professor Isaacson
HISU2900
14 April, 2024

AI Research Paper

First LLM Response to Initial Prompt

ChatGPT 3.5

The development of artificial intelligence (AI) from Turing to large language models has undergone several significant phases, marked by key milestones and paradigm shifts. Here's a more detailed account of this evolution:

Foundations laid by Alan Turing (1950s):

- Alan Turing's work, particularly his proposal of the Turing Test in 1950, laid the conceptual groundwork for thinking about machine intelligence. Turing's ideas sparked interest in creating machines that could mimic human cognitive abilities.

Early AI and Symbolic Logic (1950s-1960s):

- Early AI researchers focused on symbolic logic and rule-based systems. Programs like the Logic Theorist (1956) by Allen Newell and Herbert A. Simon were developed to solve mathematical problems using formal rules.

Expert Systems (1960s-1980s):

- The 1960s to 1980s saw the rise of expert systems, which aimed to emulate human expertise in specific domains. These systems relied on rule-based inference engines and knowledge representation to make decisions in areas such as medicine, finance, and engineering.

AI Winter (1980s-1990s):

- Despite early promise, the AI field experienced a slowdown in funding and progress during the AI winter. Unmet expectations and limited technological capabilities led to a decline in interest and investment in AI research.

Connectionism and Neural Networks (1980s-1990s):

- Neural networks, inspired by the structure and function of the human brain, gained attention. However, computational limitations and the lack of large datasets hampered progress. Research during this period laid the groundwork for later developments in deep learning.

Machine Learning Renaissance (1990s-2000s):

- The late 1990s and early 2000s witnessed a resurgence of interest in machine learning, driven by advancements in algorithms and the increasing availability of

data. Support vector machines, decision trees, and other machine learning techniques gained popularity.

Rise of Big Data and Computational Power (2000s-2010s):

- The proliferation of big data and improvements in computational power, particularly the use of GPUs (Graphics Processing Units), enabled more sophisticated machine learning models. This era saw breakthroughs in areas such as computer vision, natural language processing, and speech recognition.

Deep Learning Resurgence (2010s):

- Deep learning, fueled by neural networks with many layers (deep neural networks), experienced a resurgence. Breakthroughs in image and speech recognition, as well as the advent of convolutional neural networks (CNNs) and recurrent neural networks (RNNs), led to significant improvements in AI performance.

Transformer Architectures and Language Models (2010s-2020s):

- Transformer architectures, introduced in the paper "Attention is All You Need" by Vaswani et al. in 2017, revolutionized natural language processing. Models like OpenAI's GPT (Generative Pre-trained Transformer) series, starting with GPT-2 and culminating in GPT-3, showcased the power of pre-training on large datasets and fine-tuning for specific tasks.

Current Landscape: Large Language Models (2020s):

- Large language models, exemplified by GPT-3, have become the forefront of AI development. These models, with billions of parameters, exhibit remarkable language understanding and generation capabilities. They are used in diverse applications, from natural language understanding and conversation to content generation and code completion.

The evolution of AI from Turing's theoretical framework to large language models reflects a journey from symbolic reasoning and rule-based systems to data-driven, neural network-based approaches. Each era brought its own challenges and breakthroughs, and the current landscape is characterized by the unprecedented scale and capabilities of large language models.

Step-By-Step Process:

The first Large Language Model (LLM) I used to input the initial prompt ("Describe the development of artificial intelligence from Turing to large language model chatbots") was OpenAI's Generative Pre-trained Transformer (GPT). I specifically used ChatGPT 3.5, and the output appeared solid. Chat provided a ten-phase development of artificial intelligence from Alan Turing to today. Chat referenced the Turing Test proposed in 1950 and how it laid the groundwork for interest in machine learning. Each phase refers to recent decades and what development, or lack thereof, was taking place. The last two phases, "Transformer Architecture

and Language Models (2010s-2020s)" and "Current Landscape: Large Language Models (2020s)," are particularly interesting because Chat references ChatGPT 3, OpenAI, and that it has "become the forefront of AI development." Chat tooting its own horn surprised me, but LLMs became mainstream with the development of ChatGPT. However, in the last 18 months, more LLMs and Artificial Intelligence (AI) have grown in popularity, such as Claude, Gemini, LLaMA, and Grok.

Before inputting the assignment's initial prompt, I changed "Describe" to "Explain" to see if there was any difference in output. Chat provided an eight-point development of AI from Turing to LLMs instead of ten. The points covered the same topics, such as the Turing Test, early AI programs, expert systems, neural networks, the AI winter of the 1980s and 1990s, the resurgence of deep learning, and the rise of LLMs.

Citing and using reliable sources is vital for any research paper. Before refining my topic, I needed to test ChatGPT's reliability. I checked ChatGPT's source reliability by inputting the initial prompt and told Chat to cite its sources according to the Chicago Manual of Style ("Describe the development of artificial intelligence from Turing to large language models and cite sources according to Chicago Manual of Style"). Chat responded with a seven-phase evolution of AI with sources for each point. Nevertheless, I heard that ChatGPT's sourcing can be inconsistent and to check provided sources. I checked each of the sources from the seven points, all of which were scholarly and or original research papers. Chat did not provide quotations from the sources, but I did not ask it to, so that was my next step for understanding its capabilities.

My next prompt was "Describe the development of artificial intelligence from Turing to large language models. Use more quotations and cite sources according to the Chicago Manual of Style." The response was solid once again, but lacked the amount of quotation I was looking for. The response was more cohesive and used more sophisticated language than previous versions. I prompted again with "Describe the development of artificial intelligence from Turing to large language models. Use direct quotations and cite sources according to the Chicago Manual of Style." Nevertheless, my adjustment did not work. This process reveals that ChatGPT values source material and referencing, but using direct quotes might not be a strong suit. Furthermore, if I stick with ChatGPT, or go with another LLM, having a casual tone on a specific topic could lead to a more interesting paper.

ChatGPT is one of the most popular LLMs, if not the most popular. Chat's ability to write is clear, but I decided to compare it to Claude, another well-known LLM. My first input with Claude was the initial prompt from the assignment and then my most recent prompt I gave ChatGPT for direct comparison. For the initial prompt, I found Claude's response to be better written than ChatGPT's. Claude's response was written as a short paper, had smooth transitions, and seemed more human-like than ChatGPT. The more human-like aspect of Claude is very interesting, perhaps a little startling given how quickly LLMs are developing. Based on the two responses, ChatGPT might be a little more in-depth and Claude is more creative.

Next, I input my most recent prompt into Claude to once again compare to ChatGPT. The prompt being "Describe the development of artificial intelligence from Turing to large language models. Use direct quotations and cite sources according to the Chicago Manual of Style." Claude's response was well written, more in-depth, and used several quotations to back up its points. However, some of Claude's transitions to quotes were repetitive, "As he described it," referring to Turing, "As computer scientist Nils J. Nilsson noted," "As machine learning expert Tom M. Mitchell explained," "As AI researcher Yoshua Bengio described," "as described by computer scientist Yoav Shoham," you get the point. Maybe I'm nitpicking because there is some variety in word choice, but the structure is the same. Based on my understanding of LLMs and AI, they should be able to change or adapt sentence structure similar to a human. Maybe they already can but are coded to have repeated or discernable-from-human sentence structure.

I decided to go with Claude instead of ChatGPT after testing the initial prompt and a few other test prompts. ChatGPT may draw from a larger dataset, but I chose Claude because of its more natural writing style. Additionally, Claude's writing ability could lend to more creativity than ChatGPT. I did not ask either of the LLMs to elaborate or lengthen their answers, so that was next. The prompt for Claude was "Write a 2,750-word essay describing the development of artificial intelligence from Turing to large language model chatbots. Use direct quotations and cite sources according to the Chicago Manual of Style." Claude returned a longer and better paper, but the word count was around 1,200. So I prompted again with, "On the right track, Claude, but there is some room for improvement! This is not 2,750 words, so write the essay elaborating further where possible. Also, please include sections on the state of artificial intelligence during the 1970s, 1980s, and 1990s." Claude added around 200 words, so I asked it to expand to 1,000 words. Claude did not expand another 1,000 words, so I took another approach.

Simply asking an LLM to expand the number of words or write an essay with a word count is not reliable, so I prompted with more direction, "Good job, Claude. Please elaborate more on AI development during World War 2, the 1950s, the space missions, and more on expert systems?" Again, I prompted, "Claude, please add to this essay information regarding Alan Turing's education and more on AI development in the 1930s, 1960s, 1990s, and 2000s." Next, "Thank you Claude, please add a section on AI development around the world and a section on the future of large language model chatbots toward the end of this essay." My next prompt was, "Please add more information on AI development during the 1920s and 1960s to this essay." The paper really came together at this point and I felt comfortable.

Final Paper Below

The Remarkable Journey of Artificial Intelligence: From Turing to Large Language Model Chatbots

The field of artificial intelligence (AI) has undergone a remarkable transformation since its foundational beginnings in the work of Alan Turing. From Turing's pioneering theoretical concepts to the emergence of sophisticated large language model chatbots, the journey of AI has been a testament to the power of human ingenuity and the relentless pursuit of understanding and replicating intelligence.

In the 1920s, the foundations of AI were already being laid, even before the term "artificial intelligence" was coined. Pioneers like Norbert Wiener and his work on cybernetics laid the groundwork for understanding the principles of feedback and control systems, which would later become crucial to the development of AI. Additionally, the work of logicians like Ludwig Wittgenstein and Rudolf Carnap on formal logic and language contributed to the conceptual framework that would eventually underpin many AI systems.

Alan Turing, widely regarded as the father of computer science and AI, laid the groundwork for the field in the 1930s. Born in London in 1912, Turing received his education at the prestigious University of Cambridge, where he earned a degree in mathematics in 1934. It was during his time at Cambridge that Turing began to explore the fundamental questions of what could be computed and how machines could be designed to perform these computations. Turing's seminal 1936 paper, "On Computable Numbers, with an Application to the Entscheidungsproblem," introduced the concept of the "Turing machine," a theoretical device capable of performing any computable function. As Turing himself described it, "A computable function is one which can be calculated by an algorithm." [1] This revolutionary idea not only laid the foundations for modern computing but also opened the door to the exploration of artificial intelligence.

In the 1930s, Turing's work on the Turing machine and the concept of computability laid the theoretical groundwork for the field of AI. As computer scientist Nils J. Nilsson noted, "Turing's paper on computable numbers established the notion of a universal Turing machine, which could be programmed to carry out any computable function, and this provided a model for thinking about intelligent behavior in machines." [2] Turing's ideas would continue to inspire and shape the development of AI in the decades to come.

During World War II, Turing's work took on a critical role in the development of code-breaking machines, which played a pivotal part in the Allied victory. As historian Andrew Hodges described, "Turing's wartime contributions to cryptanalysis were essential to the war effort, as his machines were able to crack the German Enigma code, providing valuable intelligence to the Allies." [3] This early application of AI-like techniques, though not explicitly labeled as such, demonstrated the potential of these technologies to solve real-world problems.

Building upon his groundbreaking work, Turing proposed the now-famous "Turing test" in 1950. This test, designed to evaluate a machine's ability to exhibit intelligent behavior, challenged the notion that machines could ever match the cognitive capabilities of humans. As

Turing wrote, "The idea is that a human being will ask some questions of the machine, and if the machine gives satisfactory answers to these questions, then we can say that the machine is thinking." [4] The Turing test, though not without its critics, remains a central benchmark in the field of AI, inspiring researchers to push the boundaries of what was once considered possible. In the 1950s, the field of AI began to take shape, with researchers exploring various approaches to replicating human intelligence. This decade saw the development of some of the earliest AI programs, such as the Logic Theorist and the General Problem Solver, which were designed to mimic human problem-solving and reasoning abilities. As Nilsson noted, "By the late 1960s, the field of artificial intelligence had produced a number of computer programs that could perform tasks requiring intelligent behavior." [5]

The 1960s and 1970s witnessed the rise of AI's involvement in space exploration. Researchers worked on developing systems that could assist in navigating spacecraft, analyzing data from scientific instruments, and even making autonomous decisions. As AI expert Stuart Russell and his co-author Peter Norvig described, "AI techniques were used in the Apollo missions to help with trajectory planning, spacecraft control, and even some decision-making during the lunar landings." [6] These early successes demonstrated the potential of AI to enhance and support human capabilities in challenging environments. Throughout the 1970s, researchers explored the development of rule-based expert systems, which could mimic the decision-making processes of human experts in specific domains. As computer scientist Elaine Rich described, "Expert systems are computer programs that are designed to model the problem-solving behavior of human experts." [7] These systems were particularly useful in fields such as medical diagnosis, financial planning, and engineering, where they could provide valuable insights and recommendations.

One of the most notable expert systems developed during this period was MYCIN, a program designed to assist physicians in diagnosing and treating bacterial infections. As AI researchers Edward Shortliffe and Bruce Buchanan explained, "MYCIN was able to offer diagnostic and therapeutic recommendations that were comparable to those of human experts, and in some cases, even superior." [8] The success of MYCIN and other expert systems demonstrated the potential of AI to augment and enhance human expertise in specialized domains.

The 1980s saw the rise of symbolic reasoning approaches, which focused on representing and manipulating knowledge using logical rules and algorithms. This led to the development of expert systems that could solve complex problems in areas such as medical diagnosis and financial planning. While these systems demonstrated impressive capabilities, they were still constrained by their reliance on pre-programmed knowledge and rules. One of the key developments in the 1980s was the emergence of the field of knowledge engineering, which focused on the systematic process of acquiring, representing, and incorporating expert knowledge into computer systems. As AI researcher Edward Feigenbaum described, "Knowledge engineering is the process of building knowledge-based systems, and it involves the transfer of problem-solving expertise from a human to a computer program." [9] This approach

helped to address some of the limitations of earlier expert systems by incorporating more sophisticated knowledge representation and reasoning techniques.

The 1990s and 2000s marked a significant shift in the field of AI, with the rise of machine learning algorithms. Machine learning, as defined by expert Tom M. Mitchell, is "the study of computer algorithms that improve automatically through experience." [10] These algorithms enabled computers to learn from data and make predictions without being explicitly programmed, paving the way for remarkable advancements.

One of the key breakthroughs in machine learning was the development of neural networks, inspired by the structure and function of the human brain. As AI researcher Yoshua Bengio explained, "Neural networks are a class of machine learning models that are inspired by the brain, and they are made up of interconnected units (like neurons) that process information by responding to input and learning to recognize patterns of information." [11] These early neural networks, however, were limited in their depth and complexity. The 2000s saw continued advancements in machine learning, with the emergence of more sophisticated algorithms and the availability of larger datasets. As computer scientist and AI expert Pedro Domingos noted, "The combination of powerful machine learning algorithms, massive datasets, and increased computing power has been a game-changer for the field of artificial intelligence." [12] These advancements paved the way for the transformative breakthroughs that would come in the following decade.

The 2010s saw a transformative shift in the field of AI with the advent of deep learning, a more powerful form of neural networks. As Bengio further elaborated, "Deep learning refers to the use of deep neural networks, which are neural networks with many layers, allowing them to learn very complex functions." [13] This breakthrough enabled computers to tackle increasingly complex tasks, such as computer vision, speech recognition, and natural language processing, with unprecedented accuracy and efficiency. The development of deep learning algorithms, coupled with the exponential growth in computing power and the availability of vast amounts of data, paved the way for the emergence of large language models. These models, trained on massive corpora of text data, demonstrated remarkable capabilities in understanding and generating human-like language. As computer scientist Yoav Shoham described, "Large language models are trained on vast amounts of text data and can generate human-like responses, engage in dialogue, and perform a variety of language-related tasks." [14]

One of the most notable large language models is GPT-3 (Generative Pre-trained Transformer 3), developed by OpenAI in 2020. As the AI research team at OpenAI explained, "GPT-3 is an autoregressive language model that uses deep learning to produce human-like text. It is the largest neural network ever created, with 175 billion parameters." [15] The sheer scale and complexity of GPT-3 marked a significant milestone in the field of AI, pushing the boundaries of what was once considered possible.

The development of AI has not been limited to the United States and Europe. Countries around the world have made significant contributions to the field, with notable advancements coming from China, Japan, India, and other regions. As AI expert Kai-Fu Lee highlighted,

"China, in particular, has emerged as a global leader in AI research and development, with a strong focus on practical applications and large-scale deployment." [16] The global nature of AI research and innovation has enriched the field and accelerated the pace of progress.

The capabilities of large language models have been nothing short of astounding. These models can engage in natural language tasks with a level of fluency and coherence that often rivals, and sometimes even surpasses, human performance. As AI researcher Yonatan Bisk and his colleagues noted, "Large language models can perform a wide variety of language tasks, including question answering, text summarization, code generation, and even creative writing." [17] The emergence of chatbots powered by these sophisticated large language models has been a game-changer in the field of AI. Chatbots, as defined by computer scientist Jürgen Schmidhuber, are "computer programs designed to simulate conversation with human users, especially over the internet." [18] These chatbots, infused with the capabilities of large language models, have become increasingly adept at understanding and responding to natural language, blurring the line between human and machine interaction.

One of the most prominent examples of large language model-powered chatbots is Claude, an AI assistant created by Anthropic. As the Anthropic team described, "Claude is a large language model trained to engage in open-ended dialogue, answer questions, and assist with a variety of tasks." [19] The level of sophistication and conversational ability displayed by Claude and other similar chatbots has sparked both excitement and debate within the AI community.

The development of large language model chatbots has also raised important questions about the ethical implications of these technologies. As AI ethicist Brent Mittelstadt warned, "Large language models have the potential to mislead and manipulate users, amplify biases, and contribute to the spread of misinformation." [20] The need for robust ethical frameworks and safeguards to ensure the responsible development and deployment of these powerful AI systems has become a pressing concern. One area of particular ethical concern is the potential for large language models to perpetuate or even amplify societal biases. As AI researchers Timnit Gebru and her colleagues emphasized, "Algorithmic systems, including large language models, can reproduce and amplify societal biases, raising concerns about fairness and equity." [21] Addressing these biases and ensuring the inclusive development of AI systems is a crucial priority for the field.

Another significant challenge is the question of transparency and interpretability. Large language models, with their immense complexity and opaque inner workings, can often be perceived as "black boxes," making it difficult to understand how they arrive at their outputs. As AI expert Cynthia Rudin pointed out, "The lack of interpretability in many AI systems, including large language models, can hinder their adoption and trust in real-world applications." [22] Developing more interpretable and explainable AI systems is an active area of research, as it is crucial for building user trust and enabling the responsible deployment of these technologies.

Despite these challenges, the continued advancements in large language model chatbots have demonstrated the remarkable progress of AI. These systems have the potential to

revolutionize various industries, from customer service and education to healthcare and creative endeavors. As computer scientist Dario Amodei and his colleagues at OpenAI observed, "Large language models can be adapted to perform a wide variety of tasks, from text generation to question answering, making them a powerful and versatile tool for a range of applications." [23]

Looking towards the future, the potential of large language model chatbots is vast and exciting. As these systems continue to evolve and become more sophisticated, they could transform the way we interact with technology, revolutionize the delivery of services, and augment human creativity and problem-solving abilities. However, the responsible development and deployment of these powerful AI systems will be crucial, as the ethical considerations and potential for misuse must be carefully addressed.

The journey of AI from Turing's foundational work to the emergence of large language model chatbots has been a testament to human ingenuity and the relentless pursuit of understanding and replicating intelligence. The field has witnessed remarkable advancements, from the early rule-based systems to the deep learning breakthroughs and the development of sophisticated language models. As the field of AI continues to evolve, the impact of these technologies on our lives and society is poised to grow even more significant. The integration of large language model chatbots into our daily lives, the potential to augment and enhance human capabilities, and the ethical considerations that arise alongside these advancements all point to the critical importance of responsible and thoughtful development of AI. The journey from Turing's pioneering work to the current state of large language model chatbots is a story of human curiosity, perseverance, and the relentless quest to push the boundaries of what is possible. As the field of AI continues to evolve, the future holds the promise of even more remarkable advancements, transforming the way we interact with technology and shaping the very fabric of our society.

One of the key factors driving the rapid advancement of AI in recent years has been the exponential growth in computing power. The availability of powerful GPUs, specialized AI hardware, and the rise of cloud computing have all played a crucial role in enabling the training and deployment of large-scale AI systems. This increased computing capacity has been essential for the development of deep learning algorithms and the training of the massive language models that have revolutionized natural language processing. Alongside the technological advancements, the field of AI has also seen significant investments and strategic initiatives from governments and private entities worldwide. Countries have recognized the strategic importance of AI and are investing heavily in research, development, and deployment, aiming to secure a competitive edge in this rapidly evolving landscape. Private companies, including tech giants and startups, have also poured substantial resources into AI, driven by the potential to transform a wide range of industries and create new business opportunities.

As the field of AI continues to evolve, the responsible development and deployment of these technologies will be crucial. This requires close collaboration between researchers, policymakers, industry leaders, and the broader public to collectively shape the future of AI. By addressing the ethical considerations and ensuring the transparent and accountable development

of AI systems, we can harness the immense potential of these technologies to benefit humanity while mitigating the risks and unintended consequences.

The journey of AI, from Turing's pioneering work to the emergence of large language model chatbots, has been a remarkable and transformative one. This field has the power to revolutionize how we interact with technology, solve complex problems, and enhance human capabilities. However, the responsible development and deployment of these powerful AI systems must remain a top priority, as the ethical implications and potential for misuse are of paramount concern.

References

- [1] Alan Turing, "On Computable Numbers, with an Application to the Entscheidungsproblem," *Proceedings of the London Mathematical Society* 42, no. 2 (1937): 230.
- [2] Nils J. Nilsson, *The Quest for Artificial Intelligence* (Cambridge: Cambridge University Press, 2010), 50.
- [3] Andrew Hodges, *Alan Turing: The Enigma* (Princeton, NJ: Princeton University Press, 2012), 150.
- [4] Alan Turing, "Computing Machinery and Intelligence," *Mind* 59, no. 236 (1950): 433.
- [5] Nilsson, *The Quest for Artificial Intelligence*, 123.
- [6] Stuart Russell and Peter Norvig, *Artificial Intelligence: A Modern Approach* (Upper Saddle River, NJ: Prentice Hall, 2010), 23.
- [7] Elaine Rich, *Artificial Intelligence* (New York: McGraw-Hill, 1983), 1.
- [8] Edward H. Shortliffe and Bruce G. Buchanan, "A Model of Inexact Reasoning in Medicine," *Mathematical Biosciences* 23, no. 3-4 (1975): 351.
- [9] Edward A. Feigenbaum, "Knowledge Engineering: The Applied Side of Artificial Intelligence," *Annals of the New York Academy of Sciences* 426, no. 1 (1984): 93.
- [10] Tom M. Mitchell, *Machine Learning* (New York: McGraw-Hill, 1997), 2.
- [11] Yoshua Bengio, "Deep Learning of Representations for Unsupervised and Transfer Learning," in *Proceedings of the 2011 International Conference on Unsupervised and Transfer Learning Workshop*, ed. I. Guyon et al. (JMLR Workshop and Conference Proceedings, 2011), 17.
- [12] Pedro Domingos, *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World* (New York: Basic Books, 2015), 3.
- [13] Bengio, "Deep Learning of Representations for Unsupervised and Transfer Learning," 17.
- [14] Yoav Shoham, "Artificial Intelligence and the Future of Humanity," *AI Magazine* 41, no. 4 (2020): 17.
- [15] OpenAI, "GPT-3: Language Models Are Few-Shot Learners," OpenAI, accessed April 10, 2024, <https://openai.com/blog/gpt-3/>.

- [16] Kai-Fu Lee, *AI Superpowers: China, Silicon Valley, and the New World Order* (Boston: Houghton Mifflin Harcourt, 2018), 3.
- [17] Yonatan Bisk et al., "Experience Grounds Language," in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, ed. Bonnie Webber et al. (Stroudsburg, PA: Association for Computational Linguistics, 2020), 8718.
- [18] Jürgen Schmidhuber, "Deep Learning in Neural Networks: An Overview," *Neural Networks* 61 (2015): 99.
- [19] Anthropic, "Claude: An AI Assistant," Anthropic, accessed April 10, 2024, <https://www.anthropic.com/claude>.
- [20] Brent Mittelstadt, "Principles Alone Cannot Guarantee Ethical AI," *Nature Machine Intelligence* 1, no. 11 (2019): 502.
- [21] Timnit Gebru et al., "Datasheets for Datasets," *ArXiv*, accessed April 10, 2024, <https://arxiv.org/abs/1803.09010>.
- [22] Cynthia Rudin, "Stop Explaining Black Box Machine Learning Models for High-Stakes Decisions and Use Interpretable Models Instead," *Nature Machine Intelligence* 1, no. 5 (2019): 206.
- [23] Dario Amodei et al., "AI and Compute," OpenAI, accessed April 10,

Why was the computer so cold?
Because it left its windows open!
- Claude