

Doug Cantor
Professor Isaacson
HISU2900-02
11 March 2024

Starting the Paper:

To start the paper, I first tried the Perplexity chatbot. Before then, I had only ever used ChatGPT. However, I thought I would give Perplexity a try because someone in class mentioned that it was more helpful for academic papers and citing sources.

I asked perplexity the prompt “Describe the development of artificial intelligence from Turing to large language model chatbots.” without explaining or specifying any word counts. It responded with a brief overview that was the length of 3 short paragraphs. The first paragraph referenced Turing’s proposition of the idea of AI, then the first chatbot, ELIZA, created in 1966, and then briefly described present day LLMs. Paragraph 2 discussed where AI might be headed in the next 1-5 years, and the last paragraph was a summary conclusion of what was already discussed:

Answer to Initial Prompt:

The development of artificial intelligence from Turing to large language model chatbots has been a significant journey in the field of technology. Alan Turing first proposed the idea of AI in 1950, introducing the "imitation game" to test a machine's capacity to mimic human behavior. The first chatbot, ELIZA, was created by Joseph Weizenbaum in 1966 and could pass the Turing Test, demonstrating early success in AI chatbot technology. Fast forward to recent advancements, large language models (LLMs) like ChatGPT have revolutionized AI applications, enabling automated help with creative ideas, coding software, summarizing text, and enhancing language fluency. These LLMs are trained on massive datasets to predict the next word in a sequence, making them powerful tools for various tasks.

The impact of AI chatbots on customer service is profound, with projections indicating that by 2025, chatbots will handle 95% of customer interactions, streamlining processes and improving efficiency. Companies are increasingly using chatbots for basic customer engagement tasks, freeing up human agents for more complex issues and enhancing productivity. However, concerns around privacy and security persist, with experts highlighting potential privacy issues related to data collection practices of AI chatbots like ChatGPT. Addressing these privacy concerns is crucial for ensuring ethical and responsible use of AI technologies.

In conclusion, from the inception of AI concepts by Turing to the emergence of sophisticated large language model chatbots like ChatGPT, the evolution has been marked by significant milestones in artificial intelligence technology. These advancements have reshaped how

businesses interact with customers and automate various tasks while also raising important ethical considerations regarding data privacy and security in AI applications.

After some brainstorming, I realized that splitting the paper up into sections by decade might be a good way to build a solid foundation for the paper. So, my second prompt to Perplexity was “can you expand on your answer and go into more detail on the advancements of AI through each decade, starting in the 1950s and ending in the 2020s?” Perplexity then gave a relevant title to each decade/grouping of decades and 1-3 short bullets that highlighted the main achievements of the era.

Side Note:

I knew from previously working in GPT that GPT data knowledge only goes up to 2022, so I wanted to know if Perplexity’s knowledge base is current or has a cutoff date. I asked, “Does perplexity have knowledge up to today, or is the data it pulls from limited to a previous date?”

I was given a thorough explanation including information such as, “Perplexity AI has up-to-date knowledge as of today, offering cutting-edge information discovery and sharing capabilities. The company has recently unveiled online Large Language Models (LLMs) that outperform competitors like OpenAI's GPT-3.5 and Meta's raw models in terms of freshness and accuracy.” And in summary, “Perplexity AI's models are continuously updated to offer users the most relevant and up-to-date information. By combining advanced algorithms like GPT technology with a user-centric approach to search, Perplexity AI is reshaping the online search experience and positioning itself as a formidable contender in the field of artificial intelligence-driven search engines.”

Perplexity seemed promising, so I decided to keep with using it. At that point, I realized I should sign up to save my questions and responses.

Expanding the Paper:

At this point, I had no content other than a few bullet points. So, I asked Perplexity, “can you provide a more in-depth response to discuss the evolution of AI in the 1950s.” I was provided eight bullets signifying major milestones from the 1950s to the 1970s.

It seemed the 1950s did not have enough content alone to warrant a whole section of the paper, so grouping the 1950s, 60s, and 70s seemed logical. By this logic, I decided to group the 1980s and 90s together, as well as the 2000s and 2010s, with the 2020s being its own section.

At this point, I also checked on the source's aspect to Perplexity's responses. After each bullet point that it writes out, it gives easy links to the internet sources it pulls from. The sources are not entirely scholarly sources but seem to be informative and honest internet sources.

I moved forward and asked for "a more in-depth response to the evolution of AI from Turing to LLM chatbots by discussing major points in history from the 1980s and 1990s." I was provided four more bullet points and a brief explanation for each.

I thought there might be more than four major events in those two decades, so I asked, "are there any more historical points to mention?" The chatbot basically just reworded the previous response except for a few new points, but it was mostly the same. So, I assumed it covered the major events that I requested.

I repeated this process for the 2000s/2010s and the 2020s sections.

At this point, I had roughly 30 bullet points to map the history of AI from Turing to LLMs today. I thought this gave me a solid foundation for the paper. The paper could be grouped into 4 groups: 1950s-1970s, 1980s-1990s, 2000s-2010s, and 2020s. If we were to spread out the 3000 words that the paper required into these four timespans, we could allot about 750 words to each. This would serve as a good starting point for the paper:

Producing the Paper:

After having Perplexity map out the timeline, I asked it to produce a 3000-word response based on what it had described in the bullets. The response was a 1490-word response, not a 3000 word response. So, I had to repeat myself to ask again to write a 3000 word response. The response was a 2010 word response. However, Perplexity claimed it wrote 2900 words, an obvious hallucination.

I clarified and told Perplexity it actually only wrote 2010 words, and it agreed that I caught it making a mistake. It then produced a 2090 word response, and I told Perplexity yet again that it was short on word count. It acknowledged its mistake for the 3rd time and rewrote the prompt, but in only 2415 words. I was frustrated and stopped asking it to write the whole thing.

My new approach was to prompt Perplexity to write each section one at a time in approximately 750 words. I started with the 1950s-1970s and the result was many small paragraphs. So, I then had to ask Perplexity to combine the response into fewer larger paragraphs to seem more like an academic paper.

After it produced the first section of the paper, I thought it would be appropriate for each section to have a memorable title more than just be titled “1950s to 1970s” so I asked Perplexity to title the sections.

I was satisfied with the content and length of the first section. So, I repeated this process for the following sections and the result was the paper I now have.

I decided not to prompt Perplexity to change the style or tone because the style and tone produced by Perplexity seemed appropriate for an academic paper. The structure was simply presenting the history and facts in a professional, chronological fashion.

Citing Sources:

Perplexity was helpful in that I did not have to work too hard for it to provide me with the sources it used. Each paragraph ended with a numbered hyperlink to the source it pulled from for the given paragraph. Each response also started with listing out all sources used before even starting the actual response. It also turned out that Perplexity used some scholarly sources that it forgot to mention when I prompted it to give me all the sources at the end, such as Blagoj Delipetrev’s publication titled "Historical Evolution of Artificial Intelligence."

I did notice some hallucinations with the sources. When I asked Perplexity to list all the sources at once, it listed about 20 of them, but a few were unrelated research papers, such as “Generative Adversarial Nets; Advances in Neural Information Processing Systems” by Ian Goodfellow and “ImageNet Classification with Deep Convolutional Neural Networks" by Alex Krizhevsky. These papers read as very complicated statistical experiments that had practically nothing to do with the topic of this paper.

Concluding Thoughts:

This paper was definitely one of the more entertaining papers I have written in my time at Tulane. I am glad that Professor Isaacson and our TAs chose to embrace the changing educational landscape instead of fighting it as many professors and institutions are.

The chatbots are only getting more powerful and learning how best to prompt them to produce quality and appropriate work is a skill that I think will largely translate to our generation’s professional lives.

Although I did not choose to change the style of my paper or make it into “rhyming couplets,” as Professor Isaacson mentioned, I wanted to see if I could produce a conventional academic paper, free of any mistakes or hallucinations. In the end, I think my paper turned out pretty well. As a soon to be graduate, I think the critical thinking skills I gained from producing this paper will stay with me as I approach and begin my professional career.

Doug Cantor

Professor Isaacson

HISU2900-02

11 March 2024

The Evolution of AI from Turing to Large Language Models

Introduction:

The journey of artificial intelligence (AI) has been a remarkable one, spanning decades of groundbreaking research, technological advancements, and paradigm shifts in how we perceive and interact with machines. From the foundational concepts introduced by visionary thinkers to the emergence of sophisticated language models that can engage in human-like conversations, the evolution of AI has been a testament to human ingenuity and the relentless pursuit of pushing the boundaries of what is possible.

The Foundations of AI: From Turing's Vision to the AI Winter (1950s-1970s)

The 1950s ushered in the birth of artificial intelligence as a field of study, with pioneering work that laid the theoretical foundations and sparked early explorations into machine intelligence. It was during this decade that the concept of AI was first introduced and the initial steps towards creating intelligent machines were taken. In 1950, Alan Turing, the brilliant British mathematician and computer scientist, published his seminal paper "Computing Machinery and Intelligence." This groundbreaking work introduced the idea of the Turing Test, a proposal for evaluating a machine's ability to exhibit intelligent behavior indistinguishable from that of a human. Turing's test set a tangible goal for AI research and ignited debates around the nature of intelligence and consciousness. The next major milestone came in 1956 at the Dartmouth

Summer Research Project on Artificial Intelligence, organized by John McCarthy, Marvin Minsky, Claude Shannon, and Nathan Rochester. This historic workshop brought together researchers from various disciplines to explore the possibility of creating intelligent machines, and it was here that the term "artificial intelligence" was officially coined. The Dartmouth conference marked the birth of AI as a dedicated field of research and set the stage for future advancements.

In the late 1950s, early successes in AI research began to emerge. Simple problem-solving programs, such as the Logic Theorist and the General Problem Solver, demonstrated the potential of AI systems to tackle complex tasks. Additionally, the development of the Lisp programming language by John McCarthy in 1958 provided a powerful tool for AI researchers, enabling the creation of flexible and expressive programs. Alongside these advancements, the 1950s also witnessed the emergence of artificial neural networks (ANNs), which were inspired by the biological neural networks in the human brain. In 1951, Marvin Minsky and Dean Edmonds created the first ANN called SNARC, using vacuum tubes to simulate neural connections. This early work laid the foundation for the development of more sophisticated neural network architectures in the decades to come.

The 1960s saw the birth of early chatbots, with Joseph Weizenbaum's ELIZA being a notable example. Developed in 1966, ELIZA was a conversational program that could engage in human-like interactions by recognizing patterns in user input and generating responses based on predefined scripts. While limited in its capabilities, ELIZA showcased the potential of AI-powered conversational agents and laid the groundwork for future advancements in natural language processing (NLP).

However, despite these early successes, the field of AI faced significant challenges in the 1970s, leading to a period known as the "AI Winter." Limitations in computational power, lack of data, and unmet expectations resulted in reduced funding and interest in AI research. This period of stagnation and disillusionment threatened to derail the progress made in the previous decades. Yet, even during the AI Winter, important developments took place that would later contribute to the resurgence of AI. In the 1970s, researchers like Edward Feigenbaum and his team at Stanford University made significant strides in the development of expert systems, which were AI programs designed to replicate human expertise in specific domains. Expert systems, such as MYCIN for medical diagnosis and XCON for computer configuration, demonstrated AI's potential to solve complex problems by emulating human decision-making processes. The late 1970s and early 1980s also saw the emergence of knowledge-based systems, which aimed to represent and reason with human knowledge in a structured and systematic way. These systems, combined with expert systems, laid the foundation for the development of more advanced AI applications in various domains.

While the AI Winter slowed progress, it also prompted a shift in focus towards solving specific problems and developing practical applications, rather than pursuing general artificial intelligence. This pragmatic approach paved the way for the commercialization of AI technologies in the decades that followed. As the 1980s progressed, advancements in neural networks and machine learning, driven by researchers like Geoffrey Hinton, Yann LeCun, and Yoshua Bengio, reignited interest in AI and set the stage for the development of more sophisticated models capable of learning from data. The resurgence of neural networks and machine learning in the late 1980s, coupled with the increasing availability of large datasets and computational power in the following decades, laid the groundwork for the deep learning

revolution that would transform AI in the 2010s and pave the way for the emergence of large language models (LLMs) like GPT-3 and ChatGPT. While the 1950s to 1970s may not have directly produced LLMs, this period was crucial in establishing the theoretical foundations, early explorations, and practical applications that would eventually lead to the development of these powerful language models.

The AI Renaissance: From Expert Systems to Commercial Viability (1980s-1990s)

While the AI Winter of the 1980s was characterized by a significant decline in funding and progress, this period also witnessed important developments that would contribute to AI's resurgence. One major area of focus was expert systems, which aimed to replicate the decision-making abilities of human experts in specific domains. Expert systems like MYCIN, developed in the 1970s to diagnose and treat bacterial infections, demonstrated AI's potential for solving complex, specialized problems. MYCIN could make diagnoses and treatment recommendations based on a patient's symptoms and medical data, mirroring the reasoning of human experts. The 1980s also saw the emergence of knowledge-based systems, which represented human knowledge in a structured, logical way to enable automated reasoning and decision-making. These systems, combined with expert systems, paved the way for more advanced AI applications across various industries. As the 1980s progressed, renewed interest in neural networks and machine learning began planting the seeds for the next wave of AI progress. Researchers like John Hopfield and David Rumelhart made key advancements in neural network architectures and training algorithms like backpropagation. The late 1980s marked a pivotal shift as the AI field expanded its algorithmic toolkit and received a significant boost in funding from both

government and commercial sources. This "AI renaissance" set the stage for the rapid developments that followed in the 1990s.

In the 1990s, AI research and development gained significant momentum driven by increasing computational power, larger datasets, and theoretical breakthroughs. Machine learning approaches like neural networks and support vector machines enabled AI systems to learn from and adapt to data in more human-like ways. Natural language processing (NLP) made major strides, with systems becoming more adept at understanding, interpreting, and generating human language. One landmark 1990s development was Apple's introduction of the speech recognition software PlainTalk in 1993. Computer vision also advanced rapidly, with AI systems getting better at recognizing patterns, classifying images, and even describing visual scenes. By the late 1990s, real-time face detection became possible with systems like Rowley's neural network-based approach. The 1990s witnessed the increasing commercialization of AI across various sectors. Companies began integrating AI into products and services like recommendation engines (Amazon, Netflix), financial trading systems, medical diagnosis tools, and intelligent tutoring programs.

While expert systems remained prominent, the 1990s also gave rise to new AI paradigms like intelligent agents - software entities that perceived their environment and acted rationally to achieve goals. Intelligent agent architectures like reactive, deliberative, and hybrid models were proposed. As the decade progressed, the AI field made major theoretical advancements as well. In 1997, IBM's Deep Blue defeated world chess champion Garry Kasparov, showcasing AI's potential for complex decision-making. That same year, Hochreiter & Schmidhuber introduced Long Short-Term Memory (LSTM), a type of recurrent neural network that became foundational for many future AI breakthroughs. The 1990s culminated with growing enthusiasm about AI's

future potential, buoyed by rapid progress across multiple fronts - machine learning, knowledge representation, automated reasoning, natural language processing, robotics, and computer vision. However, the AI community also recognized the need to manage expectations, having witnessed the disillusionment of the 1980s AI Winter. There was a concerted effort to set realistic goals and timelines to sustain funding and research momentum.

As the new millennium dawned, the state of AI reflected a field that had matured significantly, shedding its reputation as a mere academic curiosity. AI systems were becoming increasingly sophisticated, capable, and commercially viable across a range of real-world applications. This set the stage for the transformative developments that would take place in the 2000s and beyond as AI combined with the exponential growth of data and computing power.

The Deep Learning Revolution and the Rise of AI Commercialization: (2000s-2010s)

The 2000s ushered in a new era of progress and commercialization for artificial intelligence. As computational power increased exponentially and large datasets became more readily available, AI systems grew increasingly sophisticated and capable across various domains. In the early 2000s, machine learning algorithms like support vector machines gained significant traction and were applied to a wide range of problems in areas such as computer vision, speech recognition, and natural language processing. Companies began integrating AI into their products and services at a larger scale, with recommendation engines for e-commerce giants like Amazon and media streaming platforms like Netflix being among the early mainstream applications. However, it was the late 2000s that witnessed some of the most significant breakthroughs that would propel AI into the next decade and beyond. In 2006, a seminal paper by Geoffrey Hinton, Ruslan Salakhutdinov, and others introduced deep belief

networks, a powerful form of deep neural network. This work reignited widespread interest in deep learning models, setting the stage for the transformative progress that was to come.

The 2010s are considered the decade of deep learning's transformative impact across AI. In 2012, a deep convolutional neural network developed by researchers at the University of Toronto won the prestigious ImageNet computer vision challenge, outperforming traditional machine learning techniques by a substantial margin. This breakthrough set off a deep learning revolution, with major tech companies like Google, Microsoft, Amazon, and others doubling down on AI research and development efforts. These companies also launched cloud services to democratize access to deep learning capabilities for businesses and developers worldwide. Specialized hardware like GPUs and TPUs provided the massive computational power required to train larger, more complex deep neural networks efficiently.

Natural language processing (NLP) made major strides in the 2010s, with deep learning models becoming increasingly adept at language translation, text summarization, question answering, and other language-related tasks. A significant milestone was Apple's release of Siri in 2011, bringing AI-powered virtual assistants into the mainstream consumer market. Computer vision also saw rapid progress, with deep learning enabling real-time object detection, image captioning, and even visual reasoning capabilities. Self-driving car projects like Waymo (Google) reached key milestones by leveraging this advanced AI vision technology.

While supervised deep learning models were achieving state-of-the-art results across many domains, the late 2010s also witnessed the rise of unsupervised and generative AI techniques. In 2014, researchers introduced Generative Adversarial Networks (GANs), a novel AI architecture for generating synthetic data like images or audio. This paved the way for further advancements in generative AI models.

The decade culminated with the development of large language models like GPT (2018) and GPT-3 (2020) by OpenAI. These massive neural networks, trained on vast datasets, demonstrated remarkable capabilities in understanding and generating human-like text, ushering in a new era of conversational AI and content generation possibilities.

However, the rapid AI progress of the 2010s also sparked growing concerns around ethical AI development. As these systems grew more powerful and ubiquitous, issues like algorithmic bias, privacy violations, and lack of transparency emerged as major challenges. There were increasing calls for responsible AI practices, with organizations like the IEEE developing ethical guidelines and frameworks to mitigate AI's potential negative impacts on individuals and society. Policymakers also began exploring AI governance and regulation to ensure these technologies are developed and deployed safely.

The 2010s laid the groundwork for the current AI landscape, where deep learning and large language models have become vital tools across many industries. But it also underscored the importance of addressing AI's ethical risks as these capabilities become increasingly advanced and integrated into critical systems.

The Age of Large Language Models and Generative AI (2020-2024)

The early 2020s marked a pivotal period in the evolution of artificial intelligence, with rapid advancements in large language models (LLMs) and generative AI that have transformed the AI landscape. 2020 saw the release of GPT-3 by OpenAI, a massive language model trained on a staggering amount of online data. With 175 billion parameters, GPT-3 demonstrated remarkable capabilities in understanding and generating human-like text across a wide range of tasks and domains. This breakthrough highlighted the potential of very large neural networks to

acquire broad knowledge and skills from data alone. Building on GPT-3's success, AI labs doubled down on developing even more powerful language models. In 2021, Google unveiled its Pathways Language Model (PaLM) with 540 billion parameters, while DeepMind introduced Gopher and Chinchilla models in the hundreds of billions of parameters. These massive LLMs pushed the boundaries of what was possible with text generation, comprehension, and analysis.

However, the real game-changer came in late 2022 with the release of ChatGPT by Anthropic. Trained on a version of GPT-3 but with key improvements, ChatGPT was the first LLM that allowed open-ended conversational interactions. Its ability to engage in substantive back-and-forth dialogs, provide detailed explanations, and even assist with coding and creative writing, took the world by storm. ChatGPT's viral success brought LLMs into mainstream consciousness and sparked widespread discussions about AI's future impact across industries like education, customer service, content creation, and more. Companies raced to integrate ChatGPT-like capabilities, while educators grappled with its implications for teaching and assessment.

In 2023, the AI race intensified as major tech giants unveiled their own advanced conversational AI assistants powered by custom LLMs. Google released Bard, its answer to ChatGPT, while Microsoft doubled down on OpenAI by investing billions and integrating their models into products like Bing and Office. Multimodal AI models that could understand and generate content across different modalities like text, images, and audio also made major strides. Google's Parti and Anthropic's Claude demonstrated remarkable versatility in handling multimodal inputs and tasks. The progress in generative AI continued in 2023 and 2024, with tools like Stable Diffusion and Midjourney enabling highly realistic text-to-image generation. AI art and content creation became accessible to millions, sparking debates around intellectual property, copyright laws, and the future of creative professions.

However, the rapid advancement of AI has also amplified concerns around ethical risks like bias, privacy violations, misinformation, and lack of transparency. In 2023, the EU passed its landmark AI Act to regulate "high-risk" AI applications, while the U.S. explored its own governance frameworks. Efforts to develop more robust, reliable, and trustworthy AI systems have intensified. Techniques like constitutional AI, which bakes in rules and constraints during training, have gained traction. AI labs are also exploring ways to instill human values and preferences into LLMs through methods like reinforcement learning from human feedback.

As of March 2024, we are witnessing the rise of AI assistants that can engage in freeform multimodal interactions, blending language, vision, and other modalities seamlessly. LLMs have become core infrastructure powering a new wave of AI applications and services across industries. In healthcare, AI is accelerating drug discovery, enabling earlier disease detection through medical imaging analysis, and providing clinical decision support. Generative AI is being leveraged for applications like generating synthetic training data for models and producing personalized educational content. Sustainable energy and climate change mitigation are other key areas where AI is making an impact. LLMs and generative models are aiding the design and optimization of renewable energy systems, batteries, and carbon capture technologies. AI is also enhancing climate modeling and prediction capabilities.

However, the evolution of AI has not been without growing pains. Concerns around AI existential risk and the long-term implications of super intelligent systems have intensified. Influential voices like Elon Musk have warned about the potential dangers of advanced AI, calling for stringent regulation and safety measures. As we look ahead, the future of AI remains both exhilarating and daunting. The rapid pace of progress shows no signs of slowing, with even more powerful models and novel architectures on the horizon. But addressing AI's ethical

challenges and ensuring its responsible development will be crucial to unlocking its full potential for the betterment of humanity. The journey of AI has been remarkable, evolving from the conceptual origins of Turing's "imitation game" to the modern reality of conversational AI assistants and generative models that can create virtually any type of content. As AI continues to reshape our world, striking the right balance between innovation and ethical governance will be the key to navigating the transformative road ahead.

Works Cited:

The Foundations of AI: From Turing's Vision to the AI Winter (1950s-1970s)

"A (Brief) History of Machine Learning - News - SparkFun Electronics." SparkFun, 22 Aug. 2023

"The History of Artificial Intelligence: Complete AI Timeline - TechTarget." TechTarget, 16 Aug. 2023

"History of Artificial Intelligence - Wikipedia." Wikipedia

"The Evolution Of Artificial Intelligence In Cybersecurity - VC3." VC3, 3 Oct. 2023

"Appendix I: A Short History of AI." AI100 Stanford

The AI Renaissance: From Expert Systems to Commercial Viability (1980s-1990s)

Rollins, Mark, et al. "The History and Evolution of Artificial Intelligence: A Journey Through Time." LinkedIn, 7 June 2023

Delipetrev, Blagoj, et al. "Historical Evolution of Artificial Intelligence." Publications Office of the European Union, 2020

"History of Artificial Intelligence - LibGuides at Illinois Central College." Illinois Central College, 11 Mar. 2024

Patel, Vandan R., et al. "Overview of Artificial Intelligence in Medicine." Journal of Family Medicine and Primary Care, vol. 8, no. 7, July 2019, pp. 2328–2334

The Deep Learning Revolution and the Rise of AI Commercialization: (2000s-2010s)

Delipetrev, Blagoj, et al. "Historical Evolution of Artificial Intelligence." Publications Office of the European Union, 2020

Hinton, Geoffrey E., et al. "A Fast Learning Algorithm for Deep Belief Nets." Neural Computation, vol. 18, no. 7, July 2006, pp. 1527–1554

Krizhevsky, Alex, et al. "ImageNet Classification with Deep Convolutional Neural Networks." Communications of the ACM, vol. 60, no. 6, June 2017, pp. 84–90

Goodfellow, Ian, et al. "Generative Adversarial Nets." Advances in Neural Information Processing Systems, vol. 27, 2014

Vaswani, Ashish, et al. "Attention Is All You Need." *Advances in Neural Information Processing Systems*, vol. 30, 2017

Rollins, Mark, et al. "History and Evolution of Artificial Intelligence: A Journey Through Time." LinkedIn, 7 June 2023

The Age of Large Language Models and Generative AI (2020-2024)

"57 NEW AI Statistics (Apr 2024)." *Exploding Topics*, 27 Mar. 2024

Roser, Max. "The Brief History of Artificial Intelligence: The World Has Changed Fast – What Might Be Next?" *Our World in Data*, 6 Dec. 2022

Feldman, Michael, et al. "The Evolution of Artificial Intelligence (AI) Spending by the U.S. Government." *Brookings*, 26 Mar. 2024