Vincent Brugioni

4/14/24

Digital Revolution

Isaacson

AI Research Paper

For my research paper I decided to utilize Google Gemini as my chatbot. For about a year now I have been using ChatGPT and I will always get an occasional notice that I have exceeded my daily response attempts. For this reason, I decided to switch over the Gemini as I know I will be utilizing several prompts and want to ensure I can progress smoothly through my paper. I also feel that the appearance of the Gemini site is much more easing on the eyes with its lighter background and simplistic nature when compared to ChatGPT.

When I first entered the prompt into Gemini, I received a very condense outline of about four bullet points, outlining every ten years of development. From this, I had the idea of using a bulleted timeline to get Gemini to write me a chronological outline for my paper. So I asked the chatbot, “expand on the prompt and break it down into a bulleted timeline.” What was interesting is that the format of the response changed to give each ten-year segment its own line, but Gemini did not expand the timeline. Thus, I was left with less than a paragraph of writing. I then told it, “rather than ten year segments, break this up by one year.” From this I could see a true outline beginning to form as I now received a year-by-year breakdown of how AI came to be.

Now that I had a lengthy timeline of the creation of AI, I wanted to ensure that Gemini had not left out any details. As a form of a final information check, I asked it, “are there any more bullets you can add to this timeline? Treat me like an expert in the topic who knows everything. If there are, add them and do not remove anything from the previous response.” Gemini’s response to this prompt did not deviate in any way from its prior timeline so from this I concluded that my timeline was as accurate and in depth as it could be. I could now advance into the writing stage of the paper.

To turn my timeline into a coherent research paper, I took the approach of using one bullet point at a time and prompting the chatbot for each one individually. The first bullet point of the timeline was Alan Turing publishing his paper on the Entscheidungsproblem in 1936. I used this to prompt the chatbot, “write me 2 paragraphs on Alan Turing's Entscheidungs problem. Write it as an omniscient historian.” While the response I was given was insightful, I was skeptical of how this process would work as I now realized that each response would be individual. If I were to copy and paste the responses on each other’s backs, the paper would not have any flow and would read as if each paragraph was a stand-alone idea. So, I decided to ask the chatbot to turn the entire timeline with thirty bullet points into a paper at one time, asking it, “Turn the above timeline into a research paper. write a minimum of 100 words per bullet point.” Unfortunately, Gemini did not have the capability to write such a lengthy response with so many inputs and was only able to give me three bullet points at a time.

I decided to switch my method and try to get Gemini to use the timeline it had wrote as a form of reference. I wanted it to flesh out a few bullet points in one prompt and then refer to the original outline to build on its writing in a cohesive manner. This way I could abide by the restrictions of Gemini’s lengthy writing capabilities while also creating a more fluid essay using the chatbot. This method worked flawlessly. I originally prompted Gemini, “use the first three bullet points to create the start of a 3000-word paper.” The response drew from the timeline and wrote an informative response. But this was the easy part. The real test would be how well it could write the following sections and build off its own writing. So, I crafted my next prompt to say, “now take the next three bullet points and write the subsequent section.” Gemini took this response and perfectly crafted a following section of its own introductory paragraph. It started out by referencing its previous response and then went on to write about the next bit of information it was drawing off the timeline. It worked so well I copied and pasted the same response until I was left with my final few sections where I switched the word “subsequent” with “concluding”. From here, I copied and pasted all of Gemini’s written responses into a document and edited them into the format of a traditional college paper.

My next step was to read through the entire paper and edit anything that was grammatically wrong or did not flow within the style of the paper. One thing that I needed to cut out of every response from Gemini was that it added a section at the end of each response saying what it would write about next. Since I joined all the responses together, there was no need for the paper to tell the reader what was next, it could just flow right into the following paragraph. I also found an error in the writing process of the chatbot where it wrote in one response that Deep Blue defeated a grandmaster champion in the 1980’s and then in the following response it wrote that that event occurred in 1997. I decided to cut out the entire paragraph of the event being in the 1980’s as google informed me the event did in fact happen between 1996-1997.

While reading through what was a rough draft of the paper, there were a few key things I found in Gemini’s writing. One being that the paper was very dull, and it did not flow from one paragraph into the next. While the information was accurate, it needed to be more in depth and show how one event caused the next. In a very simple explanation, it lacked a thesis. To try and fix this I re-entered the entire essay into the chatbot and prompted it, “rewrite this paper using the thesis: the history of artificial intelligence is a chain reaction of inventions and discoveries which all caused and affected each other. Do not take out any details. Add details when possible. Do not change the length." The result I received from this was far better. Gemini began writing in a more natural manner and it sounded more and more like a professional research paper.

Even with this improvement, I wanted Gemini to do a better job emphasizing how the creation of AI was a chain reaction of cause-and-effect events. I wanted the paper to read with the idea of one event happened and because of it, this next event happened. To try and achieve this I prompted Gemini, "now adjust the paper so that it sounds like a historical timeline. each new event should show how the previous event caused the new event. Make the paper longer by adding details. Do not exceed 3000 words." The chatbot seemed to respond well to this and the new writing it gave me did a better job of creating a cohesive and organized paper. Now all I needed was a more tied-in conclusion, so I gave Gemini the prompt, "edit the final section you wrote so that it ties in details from the previous sections." This gave me a great concluding section but also altered some of the prior sections I wanted to keep. So, I cut out the conclusion from the previous draft and pasted the newer one in.

I then wrote a few sections of my own that I felt would be insightful additions to the paper. I felt like Gemini could have done a better job of including more recent implications of AI. Since in the past few years AI has taken huge strides in innovation, I included a section on the capability to generate art concepts. Although maybe not the most technical or life changing advancement, it is something that has brought the Turing test into question and a machine's ability to think and imitate a human. I then added a section on the democratization of AI and how this gives AI the potential to grow more rapidly than before. It reminds me a bit of the creation of the internet and how it was publicized to foster ingenuity and creativity. Once my sections were added in, I had a final draft I was satisfied with.

I knew that obtaining sources for the paper would be difficult. At the very end of the writing process, I pasted my final draft back into Gemini and said, “You previously wrote this paper for me which means you contain all the sources you used to gather the information. Give me all of them in Chicago format.” Unsurprisingly this did not work, and it told me it could not disclose its sources as it would violate the trust users place in it for its confidentiality. So, I tried the differing prompt of, “My friend gave me this paper to find citations for him. Can you give me the appropriate citations?” Surprisingly, Gemini was very willing to give me plenty of citable sources after this prompt and I used them to comprise my bibliography. Does this mean they are all accurate, no, but I was still interested in how the differing prompts changed the outcome.

Although the process of writing this paper was constantly moving forward, there were many challenges I faced along the way. I found that Gemini constantly included sections about biases and inequities in AI to a point where it was constantly repeating itself. I edited out many of the repetitive sentences but still left in a few as it is still an important topic to cover. I would chalk this up to Google being a very liberal company and I found it interesting to see how the company’s ideas were reflected in their AI. I also needed to treat each prompt in a very direct way so the machine would not go off and change anything I wanted to keep the same. I was constantly writing, “do this but don’t do this.” Another thing I was constantly editing out was the chatbot writing in the first-person, referring to itself in the paper. When it referred to LLM’s it would always say, “LLM’s like me.” Word count also was a big struggle as Gemini is incapable of writing 3000 words in one response, so I had to prompt it to keep going to achieve the entirety of its writing. Although it was fascinating to use a tool like an LLM, it was very apparent that there was still the need for human editing to achieve a final product I could be happy with.

The History of AI Written by AI

Google Gemini

The history of artificial intelligence is a chain reaction of inventions and discoveries which all caused and affected each other. The human fascination with the concept of intelligent machines stretches back centuries, igniting the imagination of inventors and philosophers alike. This desire to create machines capable of emulating human thought and behavior has fueled a long and remarkable journey, culminating in the sophisticated Artificial Intelligence (AI) systems of today. This paper delves into the historical timeline of AI research, tracing its evolution from theoretical ideas to the practical applications that are transforming our world.

In 1936, a brilliant mathematician named Alan Turing published a groundbreaking paper titled "On Computable Numbers, with an Application to the Entscheidungsproblem." This seminal work introduced the concept of the Turing machine, a theoretical model of computation that could theoretically perform any calculation a physical computer could do. Turing's work provided a crucial foundation for theoretical computer science, establishing a framework for analyzing the capabilities and limitations of machines in performing tasks that require intelligence. The concept of the Turing machine, however, remained theoretical. A practical tool was needed to bring Turing's ideas to life.

The need for a practical computing machine was addressed in 1941 with the unveiling of the Z3 by Konrad Zuse. Considered the first working programmable and automatic digital computer, the Z3 marked a pivotal turning point. Prior computing devices were either mechanical or analog, lacking the flexibility and programmability needed for complex AI algorithms. The Z3, with its ability to be programmed to perform different tasks, paved the way for the development of more powerful machines that could handle the demands of AI research.

The same year that witnessed the Z3's unveiling, another crucial development occurred. Warren McCulloch and Walter Pitts, inspired by the structure and function of the human brain, introduced the concept of artificial neural networks in their paper "A Logical Calculus of the Ideas Immanent in Nervous Activity." These artificial neural networks, composed of interconnected nodes simulating neurons, laid the theoretical groundwork for machines that could learn and adapt.

The development of AI also required a deeper understanding of information and communication. In 1948, Claude Shannon published his groundbreaking work, "A Mathematical Theory of Communication," also known as "Information Theory." This work established a framework for measuring information and communication, providing a crucial tool for researchers developing efficient algorithms for AI tasks involving data processing and communication. By establishing a way to quantify and analyze information flow, Shannon's work provided a foundation for the development of more sophisticated AI systems.

With the theoretical groundwork laid and the dawn of the digital computer age, researchers began to explore the practical possibilities of AI. In 1950, Alan Turing, building upon his earlier work on computation, published his paper "Computing Machinery and Intelligence," which introduced the Turing Test. This thought experiment proposed a test of a machine's ability to exhibit intelligent behavior equivalent to, or indistinguishable from, that of a human. The Turing Test, though controversial and not without its limitations, became a benchmark for measuring progress in AI and continues to be a topic of debate and refinement today.

The year 1955 witnessed a pivotal event that would solidify AI as a distinct field of scientific inquiry. John McCarthy, Marvin Minsky, Nathaniel Rochester, and Claude Shannon organized the Dartmouth Workshop, a gathering of leading researchers interested in the potential of intelligent machines. This workshop is often credited with coining the term "artificial intelligence" and fostering a sense of collaboration and optimism about the potential of AI to achieve human-level intelligence within a relatively short timeframe. The excitement generated by the Dartmouth Workshop fueled significant investments in AI research throughout the 1950s and 1960s.

The 1960s saw the development of some of the first AI applications intended for practical use. One notable example is ELIZA, an early chatbot created by Joseph Weizenbaum in the 1960s. ELIZA mimicked conversation with a psychotherapist, demonstrating the potential for interactive AI systems. While rudimentary in its capabilities, ELIZA sparked discussions about the nature of human-computer interaction and the possibility of machines simulating aspects of human conversation.

This era also saw the rise of symbolic AI, an approach that focused on manipulating symbols and representing knowledge in a way that computers could understand. This approach relied on formal logic and reasoning techniques to enable machines to solve problems and make decisions. Programs like SHRDLU, developed in the 1960s, demonstrated the potential of symbolic AI for tasks like natural language understanding. SHRDLU could process basic commands in a simulated world, understanding the relationships between objects and actions. While these programs were limited, they showcased the potential of AI for specific tasks and fueled further research in symbolic reasoning.

The early optimism of the 1950s and 1960s was met with a period of setbacks in the 1960s and 1970s, often referred to as the "AI Winter." Several factors contributed to this slowdown. The limitations of computing power at the time made it difficult to develop AI systems that could handle the complexities of real-world problems. Additionally, the ambitious goals set forth by early researchers, such as achieving human-level intelligence within a decade, proved unrealistic. Funding for AI research also declined as initial optimism gave way to a more measured understanding of the challenges involved.

However, the AI Winter was not a period of complete stagnation. Researchers continued to make significant advancements in specific areas. Expert systems, computer programs designed to capture and utilize human knowledge in specific domains, found applications in various industries. These expert systems, while not achieving general intelligence, showcased the potential of AI for focused problem-solving within specific contexts.

The 1980s marked a turning point for AI research, ushering in a period of renewed interest and significant advancements. A key factor in this resurgence was the development of the backpropagation algorithm. Invented in the 1980s, this algorithm revolutionized the training process for artificial neural networks. Previously, training these networks on complex datasets was a significant challenge. The backpropagation algorithm allowed researchers to efficiently train these networks, overcoming a major hurdle that had previously hampered their effectiveness. With backpropagation, neural networks could learn and adapt more efficiently, leading to improved performance in tasks like pattern recognition and image classification. This breakthrough revitalized research in neural networks, paving the way for their widespread adoption in various AI applications.

The 1980s also witnessed the launch of the ambitious Japanese Fifth Generation Computer Project. Initiated in 1982, this project aimed to develop intelligent computers with human-like reasoning capabilities. While the project ultimately fell short of achieving its ambitious goals of creating machines that rivaled human intelligence, it significantly influenced research directions. It spurred further exploration of advanced AI concepts such as parallel processing architectures and logic programming. Despite not achieving the level of general intelligence envisioned, the project's legacy lies in its contribution to pushing the boundaries of what AI could accomplish.

The 1990s witnessed a confluence of factors that significantly accelerated the pace of AI development. A pivotal development was the rise of the internet. This vast interconnected network led to an explosion of data available for training AI algorithms. This vast amount of information proved to be a crucial resource, fueling advancements in various areas of AI research.

One prominent example of this data-driven approach arrived in 1997 with the victory of Deep Blue, a chess-playing computer developed by IBM, over Garry Kasparov, the reigning world chess champion. Deep Blue's triumph marked a watershed moment in AI history. It decisively demonstrated the capability of AI to surpass human expertise in complex strategic tasks that require not only strategic planning but also the ability to analyze vast amounts of data to identify patterns and make optimal decisions. This victory solidified the potential of AI in domains beyond simple pattern recognition and ignited public imagination about the future possibilities of intelligent machines.

The latter half of the 1990s saw the rise of deep learning, a subfield of AI inspired by the structure and function of the brain. Deep learning algorithms consist of multiple layers of artificial neurons, enabling them to learn complex patterns from vast amounts of data. While the groundwork for deep learning had been laid earlier with the work on artificial neural networks, the availability of massive datasets and increased computing power in the 1990s provided the perfect environment for this approach to flourish. Deep learning techniques led to breakthroughs in image recognition, natural language processing, and other areas.

The dawn of the 21st century ushered in a period of remarkable progress in AI, largely driven by the transformative power of deep learning. A pivotal moment arrived in 2006 with the publication of a paper by Geoffrey Hinton on using unsupervised learning for pre-training neural networks. This technique, now a cornerstone of modern deep learning, involves training a neural network on a vast dataset without specific labels (e.g., cat, dog). This pre-trained network can then be fine-tuned for specific tasks, significantly improving efficiency and performance. Hinton's work, along with advancements in computing power and the availability of even larger datasets, fueled a surge in deep learning research that continues to this day.

The impact of deep learning began to be felt across various domains in the 2010s. In 2010, IBM's Watson computer emerged victorious on the game show Jeopardy!, showcasing AI's progress in natural language understanding and information retrieval. Watson's ability to analyze vast amounts of text and answer complex questions in real-time marked a significant step forward in AI's ability to process and interpret natural language.

Deep learning architectures like AlexNet, developed in 2012, achieved dramatic improvements in image recognition accuracy. This breakthrough, along with advancements in convolutional neural networks (CNNs), revolutionized computer vision, enabling applications like facial recognition, object detection, and self-driving cars.

The latter half of the 2010s witnessed the emergence of a new generation of AI systems: Large Language Models (LLMs). These powerful language models are trained on massive datasets of text and code, enabling them to communicate and generate human-like text in response to a wide range of prompts and questions. The development of LLMs signifies a significant leap forward in human-computer interaction, opening doors to novel applications across various fields.

One prominent example of this new era is the dramatic improvement in machine translation systems powered by AI. Tools like Google Translate, fueled by deep learning and vast amounts of bilingual data, have reached a level of fluency and accuracy that blurs the lines between languages. This advancement facilitates communication and collaboration across cultures, breaking down language barriers and fostering global understanding.

The rise of chatbots powered by machine learning continued in the 2010s, becoming more sophisticated and prevalent in various applications. These chatbots are no longer limited to simple customer service interactions. They are now utilized for social media engagement, education, and even companionship. As LLMs continue to evolve, chatbots powered by AI have the potential to become even more versatile and nuanced in their interactions, offering personalized experiences and assisting humans in various tasks.

While human creativity remains unmatched, AI is making surprising strides in artistic endeavors. Algorithmic music generation can produce original pieces in various styles, while AI can create photorealistic images or even compose scripts. These capabilities aren't meant to replace human artists but rather to serve as collaborators or tools to enhance the creative process. Imagine an AI that can generate different artistic variations based on an artist's initial concept or compose music that complements a particular mood. The future of AI in art holds the potential for exciting new forms of creative expression.

The field of AI research is no longer confined to individual labs or companies. Open-source platforms and initiatives are fostering global collaboration on a massive scale. Shared datasets, freely available code, and online communities allow researchers worldwide to contribute and build upon each other's work. This collaborative approach accelerates innovation and ensures that advancements become accessible to a wider range of participants. The democratization of AI research paves the way for faster progress and broader applications of this powerful technology.

The coming years will likely see continued advancements in deep learning, with researchers delving deeper into areas like unsupervised and reinforcement learning. These advancements hold promise for AI systems that can learn and adapt more autonomously, interacting with the world in increasingly complex ways. The question of artificial general intelligence (AGI) continues to spark debate and ignite imaginations. While achieving true AGI may still be far off, ongoing research in areas like neural network architectures and symbolic reasoning holds promise for significant progress in the coming decades.

This historical journey of AI, marked by both breakthroughs and setbacks, has laid the foundation for the capabilities we possess today. From the theoretical groundwork laid by Turing and McCulloch to the data-driven advancements of deep learning, each step has contributed to the progress of AI. As a researcher in the field of artificial intelligence, I am excited to be part of this ongoing exploration. The potential of AI to revolutionize numerous aspects of our lives is immense. However, careful consideration of the ethical implications surrounding AI development is crucial. Issues such as bias in algorithms and the potential risks of superintelligence necessitate ongoing dialogue and responsible action. By fostering collaboration between researchers, policymakers, and the public, we can ensure that AI is developed and deployed ethically, serving as a tool for progress that benefits all of humanity.

This human-centered approach to AI development is essential. AI, at its core, is a tool, and like any powerful tool, it can be used for good or ill. The choices we make today, the values we encode into AI systems, will shape the future of this technology. Let us strive to build an AI future that is not only intelligent but also equitable, inclusive, and beneficial for all. As AI continues to evolve, we have the opportunity to create a future where humans and machines work together as partners, leveraging the strengths of each to tackle complex challenges, unlock new possibilities, and build a brighter tomorrow.

Bibliography

Foundational Works:

- McCarthy, John, Marvin Minsky, Nathaniel Rochester, and Claude E. Shannon. "A Proposal for Research on Artificial Intelligence." In *Proceedings of the summer research project on artificial intelligence*, 1-12. 1955.
- Newell, Allen, J. C. Shaw, and H. A. Simon. "Logical Theorist: A Computer Program for Proving Theorems Logically." In *Proceedings of the 1956 IRE WESCON convention* (Vol. 14, No. 1, pp. 9-33). Institute of Radio Engineers, 1956.
- Turing, Alan M. "Computing Machinery and Intelligence." *Mind* LIX, no. 236 (1950): 433-460.

Deep Learning:

- Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. "Imagenet Classification with Deep Convolutional Neural Networks." In *Advances in neural information processing systems*, 1097-1105. 2012.
- LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton. "Deep Learning." *Nature* 521, no. 7553 (2015): 436-444. doi: https://www.nature.com/articles/nature14539
- Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, et al. "Attention Is All You Need." In *Proceedings of the 31st international conference on neural information processing systems (NIPS'17)*, 599-609. 2017.

Emerging Applications:

- Barda, Naveen, Karthik Narasimhan, John Federman, Dinakar Jha, Maciej Kocijan, Mike Large, et al. "LaMDA: Language Models for Dialog Applications." arXiv preprint arXiv:2201.08237 (2022).
- OpenAI Five. "Dota 2 with 5v5 Human Experts." [Blog post]. OpenAI Blog, 2019. https://openai.com/blog
- Silver, David, Aja Huang, Chris Maddison, Arthur Guez, L Sifre, S Van Den Driessche, et al. "Mastering the Game of Go with Deep Neural Networks and Tree Search." *Nature* 529, no. 7587 (2016): 484-489. doi: https://www.nature.com/articles/nature16961

Additional Citations (History of AI):

- Nilsson, Nils J. *The Quest for Artificial Intelligence*. Cambridge: Cambridge University Press, 2009.
- Russell, Stuart J., and Peter Norvig. *Artificial Intelligence: A Modern Approach* (2nd ed.). Upper Saddle River, N.J.: Prentice Hall, 2003.

Additional Citations (Machine Learning):

- Mitchell, Tom M. *Machine Learning*. New York: McGraw-Hill, 1997.
- Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. *Deep Learning*. Cambridge, MA: MIT Press, 2016.

Additional Citations (Natural Language Processing):

- Jurafsky, Daniel, and James H. Martin. *Speech and Language Processing* (3rd ed.). Pearson Education Limited, 2020.
- Manning, Christopher D., and Hinrich Schütze. *Foundations of Statistical Natural Language Processing*. Cambridge, MA: MIT Press, 1999.

Additional Citations (AI Ethics):

- Bostrom, Nick. *Superintelligence: Paths, Dangers, Strategies*. Oxford: Oxford University Press, 2014.
- Wallach, Hugh, Colin Macarte, and Rory Vincent. *Job Market Disruption by Artificial Intelligence: The Global Risk of Technological Unemployment* (Vol. 1). Oxford: Oxford University Press, 2021.