Maxwell “Max” Levy
HISU 2900
Walter Isaacson
4/14/24

Original Post: Entered into ChatGPT

The development of artificial intelligence (AI) from its early conceptualizations to modern large language model chatbots represents a fascinating journey marked by significant milestones and breakthroughs. This evolution spans decades and involves the contributions of numerous researchers, engineers, and innovators. To describe this development comprehensively, let's explore the key stages and advancements in AI, starting from Alan Turing's foundational work to the emergence of large language models like GPT-3.

Early Foundations: Alan Turing and the Turing Test

The story of AI begins with Alan Turing, a British mathematician, logician, and computer scientist, who proposed the concept of a universal machine capable of performing any computation. In his seminal paper "Computing Machinery and Intelligence" published in 1950, Turing introduced the Turing Test as a measure of a machine's intelligence. The test involves a human evaluator interacting with both a human and a machine through a text-based interface without knowing which is which. If the evaluator cannot reliably distinguish between the human and the machine, then the machine is considered to have passed the Turing Test, indicating a level of artificial intelligence.

Early AI Approaches: Symbolic AI and Expert Systems

In the 1950s and 1960s, the focus of AI research shifted towards symbolic AI, which aimed to represent knowledge and solve problems using symbols and rules. This era saw the development of expert systems, which were computer programs designed to emulate the decision-making abilities of human experts in specific domains. One notable example is the MYCIN system developed in the 1970s for diagnosing bacterial infections.

The Rise of Machine Learning: Neural Networks and Deep Learning

Machine learning, particularly neural networks, gained prominence in the 1980s and 1990s. Neural networks are computational models inspired by the structure and function of the human brain, capable of learning from data and making predictions. However, progress in neural networks was limited by computational power and the availability of large datasets.

The AI Winter and Resurgence: Big Data and Computational Power

The late 20th century saw periods known as "AI winters," marked by reduced funding and interest in AI due to unmet expectations and technical challenges. However, the 2000s witnessed a resurgence of AI fueled by advancements in big data, computational power, and algorithms like deep learning. Deep learning, a subset of machine learning, involves neural networks with multiple layers, enabling them to learn complex patterns from vast amounts of data.

Breakthroughs in Deep Learning: Image Recognition and Natural Language Processing

Deep learning revolutionized various AI applications, including image recognition and natural language processing (NLP). In the mid-2010s, deep learning models achieved remarkable accuracy in tasks such as image classification, object detection, and speech recognition. This period also saw the emergence of neural networks specifically designed for NLP tasks, such as recurrent neural networks (RNNs) and long short-term memory (LSTM) networks.

Large Language Models: GPT and Transformer Architecture

The development of large language models represents a significant milestone in AI, particularly in the field of conversational AI and chatbots. The Generative Pre-trained Transformer (GPT) models, pioneered by OpenAI, introduced a new paradigm in natural language understanding and generation. GPT models, including GPT-3, are based on the Transformer architecture, which uses self-attention mechanisms to process input data in parallel, enabling efficient training on large datasets.

GPT-3 and Conversational AI

GPT-3, released in 2020, is one of the largest language models to date, with 175 billion parameters. It demonstrates impressive capabilities in understanding and generating human-like text across diverse domains. GPT-3's architecture allows it to perform a wide range of tasks, from answering questions and summarizing text to generating code and engaging in conversation. Its ability to generate coherent and contextually relevant responses has sparked interest and debate about the future of AI and human-machine interaction.

Ethical and Societal Implications

The development of AI, especially large language models like GPT-3, raises important ethical and societal questions. Concerns include bias in AI systems, privacy implications of data use, the impact on employment and labor markets, and the potential for misuse or malicious use of AI-generated content. Addressing these challenges requires interdisciplinary collaboration involving researchers, policymakers, ethicists, and industry stakeholders.

Future Directions: AI Ethics, Explainability, and Beyond GPT-3

Looking ahead, the field of AI is poised for further advancements in AI ethics, explainable AI (XAI), and the development of even more sophisticated language models. Researchers are exploring techniques to enhance the transparency and accountability of AI systems, improve their ability to understand and explain decisions, and mitigate biases and unintended consequences. Beyond GPT-3, the quest for AI systems that can truly understand and reason across diverse domains remains a central goal of AI research.

Conclusion

The development of artificial intelligence from Alan Turing's pioneering work to modern large language models like GPT-3 represents a journey of continuous innovation, fueled by advances in computational power, algorithms, and data. While AI has made remarkable progress in areas such as image recognition, natural language processing, and conversational AI, challenges remain in ensuring AI systems are ethical, transparent, and beneficial to society. The ongoing evolution of AI promises exciting opportunities and responsibilities as we navigate the complexities of an AI-enabled future.

Final Paper:

The journey of AI, a remarkable tale, Visionary ideas that never fail, Groundbreaking discoveries, a boundless grail. From the concept's inception, a spark ignited, To language models' advancements, minds united, A transformative evolution, foresight delighted.
Brilliant minds, collective efforts converge, Computational power, a relentless surge, Pushing boundaries, machines' abilities to emerge. Relentless pursuit of what machines can achieve, No limits, no bounds, we shall believe, AI's potential, a future we conceive. This remarkable journey, a testament to human ingenuity, Discoveries and ideas, fueling our community, Transformative evolution, AI's destiny.

Alan Turing, a pioneering mathematician and computer scientist, laid the groundwork for modern artificial intelligence with his visionary ideas in the mid-20th century. His seminal 1950 paper, "Computing Machinery and Intelligence," posed a provocative question that challenged the prevailing paradigm: "Can machines think?" This debate divides students, scholars, and academics, with the prevailing perception that machines cannot currently “think”, and especially couldn’t during the time of Turing .Turing's boldness in questioning such a fundamental assumption set the stage for the exploration of artificial intelligence as a viable field of study. His paper served as a catalyst, igniting a revolution in how we conceptualized and approached the nature of intelligence itself. Turing's audacious proposition opened up new frontiers, inviting researchers and scientists to venture beyond the confines of human-centric notions and investigate the potential for intelligence to manifest in non-biological systems, paving the way for the transformative advancements that would follow.

Turing's groundbreaking contribution was the concept of the "Imitation Game," later known as the Turing Test. This test proposed a scenario where a human evaluator engages in a text-based conversation with both a human and a machine, without knowing which is which. If the evaluator cannot reliably distinguish between the two, the machine is considered to have passed the test, demonstrating a level of intelligence that is indistinguishable from a human's. The Turing Test not only provided a framework for evaluating machine intelligence but also sparked a fundamental shift in how we perceive and approach the concept of intelligence itself. It challenged the traditional anthropocentric view and paved the way for a new era of inquiry into the nature of intelligence and its potential manifestations beyond biological systems.

Early Explorations:

Symbolic AI and Expert Systems Inspired by Turing's visionary ideas, the field of AI gained traction in the 1950s and 1960s with the emergence of symbolic AI. This approach focused on representing knowledge and solving problems through the manipulation of symbols and rules. Researchers sought to capture human expertise and reasoning processes in computer programs, giving rise to the development of expert systems.

While the contributions of pioneers like Alan Turing and the advancements in symbolic AI, neural networks, and machine learning often take center stage, it is crucial to acknowledge the pivotal role played by women in the early development of computing technology – a foundation upon which modern AI was built.

In the midst of World War II, a group of remarkable women, known as the "ENIAC Six," were instrumental in programming the ENIAC (Electronic Numerical Integrator and Computer), one of the earliest general-purpose electronic computers. These women – Kathleen McNulty, Marlyn Wescoff, Ruth Lichterman, Betty Snyder, Frances Bilas, and Betty Holberton – were tasked with manually programming the ENIAC to perform complex calculations, paving the way for the era of modern computing. Another pioneering figure was Grace Hopper, a computer scientist and United States Navy rear admiral, who was instrumental in the development of the first compiler for a computer programming language. Hopper's work on the UNIVAC I computer and her contributions to the creation of the COBOL (Common Business Oriented Language) programming language laid the groundwork for the widespread adoption of user-friendly programming languages, which would later facilitate the development of AI systems. Also, while Hopper was not a member of the ENIAC six, her work on programming during the Cold War Era cannot be understated. Her significant contributions to the western cause was a factor in their triumphs against the USSR, both during her military career and afterwards. Ada Lovelace, often regarded as the world's first computer programmer, made significant contributions to the field of computing even before the advent of modern computers. In the mid-19th century, Lovelace recognized the potential of Charles Babbage's Analytical Engine and wrote an algorithm for it to compute Bernoulli numbers, effectively creating the first computer program.

These pioneering women, among others, played a crucial role in the early stages of computing technology, overcoming societal barriers and stereotypes to make invaluable contributions. Their work laid the foundation for the development of modern computers and programming languages, which ultimately enabled the advancements in artificial intelligence that we witness today.

The Late 20th Century:

Knowledge-Based Systems and Expert Systems While the early enthusiasm for AI waned in the 1970s, this period witnessed the emergence of knowledge-based systems and expert systems. These systems were designed to capture and codify the knowledge and decision-making processes of human experts in specific domains, such as medical diagnosis, geological analysis, and financial planning.

One notable example was the MYCIN system, developed at Stanford University in the mid-1970s. MYCIN was an expert system for diagnosing bacterial infections, capable of providing recommendations for appropriate antibiotics based on a vast knowledge base of rules and patient data. Though limited in scope, MYCIN demonstrated the potential of AI in emulating human expertise and reasoning within specialized fields.

The 1980s: Rise of Neural Networks and Parallel Distributed Processing The 1980s marked a resurgence of interest in neural networks, a biologically-inspired approach to AI that had been largely dormant since the 1960s. Neural networks are a type of AI that takes similarities to the human brain function. This renewed interest was fueled by the development of parallel distributed processing (PDP) models, which enabled neural networks to learn and process information in a more distributed and parallel manner, mimicking the structure of the human brain. Researchers like Geoffrey Hinton, David Rumelhart, and James McClelland made significant contributions to the field of neural networks during this period, developing algorithms and architectures that paved the way for future advancements in deep learning.
The 1990s: Resurgence of AI and Intelligent Agents The 1990s witnessed a resurgence of interest and investment in AI, driven by the increasing availability of computational power and the potential for practical applications. This period saw the development of intelligent agents, software programs designed to perceive their environment and take actions to achieve specific goals. While not directly related to artificial intelligence, the emergence of the World Wide Web cannot be understated as well. It was breakthroughs in computational powers that helped breakthroughs occur over the next three decades. Influential works like "Rational Agents" by Stuart Russell and Peter Norvig, published in 1995, laid the theoretical foundations for the development of intelligent agents, while practical applications emerged in areas such as robotics, gaming, and virtual assistants.

Additionally, the 1990s saw the emergence of data mining and machine learning techniques for extracting patterns and insights from large datasets. These advancements laid the groundwork for the subsequent explosion of big data and deep learning in the early 21st century. While progress during the 1970s-1990s was incremental compared to the breakthroughs that followed, these decades played a crucial role in refining AI techniques, exploring new approaches, and setting the stage for the remarkable achievements we have witnessed in recent years.

2000s and Beyond:

The advent of deep learning, a subset of machine learning that involves neural networks with multiple layers, further accelerated the progress of AI. Deep learning models, with their ability to learn hierarchical representations of data, achieved remarkable accuracy in tasks that were previously considered insurmountable for machines, such as object detection, facial recognition, and language translation. Nowadays, there are a few examples of this that we see. For example, recurrent neural networks use sequential data. One of the lectures, which was based on Siri, is a prime example of Apple bursting onto this new sphere of technology. Nowadays, products like the aforementioned Siri, as well as Google Home and Amazon Alexa are vital to current technology, and are used by consumers worldwide.

The AI Winter and Resurgence: Big Data and Computational Power Despite the promising advancements, the field of AI experienced periods of disillusionment and reduced

funding, known as the "AI winters." These setbacks were often attributed to overly ambitious promises, unmet expectations, and technical challenges that hindered progress.

However, the dawn of the 21st century witnessed a resurgence of AI, fueled by the convergence of three key factors: big data, increased computational power, and advances in algorithms like deep learning. The availability of vast amounts of data and the exponential growth in computing power provided the necessary fuel for deep learning models to thrive, enabling them to process and learn from unprecedented volumes of information. Even over the past decade, the information online has grown quickly. This sphere of data continues to rapidly expand, and the technology developed during this timeframe further allowed companies to "keep up" with the expansion.

This synergy between data, computing resources, and sophisticated algorithms paved the way for groundbreaking achievements in various AI applications, from self-driving cars and intelligent assistants to advanced medical diagnostics and scientific research.

The Rise of Large Language Models: Transformers and GPT One of the most significant milestones in the evolution of AI has been the development of large language models, particularly the Generative Pre-trained Transformer (GPT) models pioneered by OpenAI, which was founded by Sam Altman, Greg Brockman, among others. These models are based on the Transformer architecture, which employs self-attention mechanisms to process input data in parallel, enabling efficient training on massive datasets.

GPT models, such as GPT-3, have revolutionized the field of natural language understanding and generation. With their vast knowledge bases and ability to generate human-like text across a wide range of domains, these models have demonstrated impressive capabilities in tasks like answering questions, summarizing information, and even generating creative content.

The release of GPT-3 in 2020, with its staggering 175 billion parameters, marked a significant leap in the capabilities of AI systems, as well as its predecessors, GPT 1 and 2. Its ability to engage in coherent and contextually relevant conversations has sparked a new era of human-machine interaction, opening up exciting possibilities for conversational AI and intelligent assistants. The release of GPT 3.5 and 4, which happened in March 2022 and 2023, goes farther than GPT3 in many ways. These include multimodality, larger context windows, as well as general knowledge that outpaces GPT3 significantly. For example, GPT4 has 220 billion parameters, which makes the staggering mark of GPT3 seem small.

Ethical and Societal Implications:

Navigating the AI Future As AI continues to advance at an unprecedented pace, it is crucial to address the ethical and societal implications that accompany these technological advancements. Concerns surrounding bias in AI systems, privacy implications of data use, the impact on employment and labor markets, and the potential for misuse or malicious use of AI-generated content have become increasingly pressing. AI systems can reflect and amplify human biases present in the data and algorithms used to train them, raising concerns about

perpetuating discrimination and unfair treatment based on characteristics like race, gender, or age. Addressing these biases is crucial to ensure AI systems are fair and equitable.
Privacy risks loom large with the massive data collection required for training AI models. Personal information, browsing histories, and sensitive data, if not properly safeguarded, could be misused or exploited, violating individual privacy and security.
As AI automates an increasing number of tasks, fears of job displacement and economic disruption pervade. Industries from manufacturing to customer service face potential upheaval, with concerns about the impact on livelihoods and the need for widespread reskilling and retraining.

The opaque nature of many AI systems, often referred to as "black boxes," hinders transparency and accountability. Understanding how these systems arrive at their decisions and ensuring they are free from unintended biases or errors is a significant challenge that erodes public trust and adoption. The potential for malicious use of AI-generated content poses severe dangers. Misinformation campaigns, deep fakes, and cyber threats could exploit AI's ability to generate highly convincing text, images, and videos, undermining trust and causing widespread harm.

These ethical challenges demand a multifaceted approach involving diverse stakeholders, including researchers, policymakers, ethicists, and industry leaders. Developing ethical codes, frameworks, and guidelines is crucial to ensure the responsible development and deployment of AI.

Researchers are actively exploring Explainable AI (XAI) techniques to enhance transparency and interpretability, enabling humans to understand the decision-making processes and rationale behind AI systems. This understanding is vital for identifying and mitigating biases or unintended consequences. As such, the implementation and production of XAI is important for the development of AI itself, especially as AI advances, especially in fields such as business. Policymakers grapple with the complex task of crafting regulations that balance innovation and economic opportunities with public obligations and safeguarding individual rights. Establishing robust governance frameworks and legal protections is essential for upholding ethical principles and accountability.

The quest for ethical AI development requires a sustained, collaborative commitment from all stakeholders. Prioritizing human values, rights, and well-being must be at the forefront as we navigate the transformative power of AI. As technology is developing at a rapid pace, laws are necessary in order to ensure all groups, especially historically marginalized ones have protections. By addressing these ethical concerns proactively and responsibly, we can mitigate risks and ensure the benefits of AI are equitably distributed and aligned with the greater good of society. Ethicists contemplate the profound moral dimensions of AI, grappling with questions of autonomy, agency, and the implications of increasingly capable systems. Their insights and guidance are invaluable for developing ethical principles and frameworks to govern responsible innovation.

During the Biden Administration, there was a statement made by the White House called the AI Bill of Rights to help remedy some of these concerns. In particular, it focuses on five concerns, some of which have already been highlighted within the paper. The five concerns include: Protection from safe and effective systems, algorithmic discriminatory protections, data privacy, notice and explanation, and lastly, human alternatives, consideration, and fallback. This document is a strong start to further protect United States citizens from malpractice involving AI, especially as it further permeates into current daily lives. Despite this, legislators have discussed AI and AI protection has been discussed little in Congress. It is yet to be seen whether the Biden Administration can come to a deal with the Republican led 118th Congress over AI protections.

Addressing these challenges requires a multidisciplinary approach, involving researchers, policymakers, ethicists, and industry stakeholders. Efforts are underway to enhance the transparency and accountability of AI systems, develop techniques for explainable AI (XAI), and mitigate biases and unintended consequences. Moreover, the quest for artificial general intelligence (AGI) – systems that can match or exceed human-level intelligence across a broad range of domains – remains a central goal for many AI researchers. While current AI systems excel in specific tasks, achieving true general intelligence presents formidable challenges that will require breakthroughs in areas such as reasoning, learning, and knowledge representation.

The Future of AI:

Limitless Possibilities and Shared Responsibility Looking ahead, the future of AI holds limitless possibilities and transformative potential. From revolutionizing healthcare and scientific research to tackling global challenges like climate change and energy sustainability, AI has the power to reshape our world in profound ways. As AI becomes more and more powerful, it has the opportunity to shape society in ways humans currently do not understand. If you had told me that five years ago I would write a paper using AI, I would not know what to say to you. However, with great power comes great responsibility. As we continue to push the boundaries of AI, it is imperative that we embrace a shared commitment to ethical development, responsible innovation, and inclusive decision-making processes. By fostering interdisciplinary collaboration and promoting transparency and accountability, we can ensure that the benefits of AI are equitably distributed and its risks are carefully managed.

From Turing's dream to GPT-3's gleam, A saga of minds, a visionary stream. Human brilliance, relentless and keen, In AI's journey, a visionary theme. On giants' shoulders, we dare to dream, The future calls with a powerful beam. Ethical heights, in every scheme, For humanity's sake, we'll always redeem.

Sources: Turing, Alan M. "Computing Machinery and Intelligence." Mind 59, no. 236 (1950): 433-460. https://academic.oup.com/mind/article/LIX/236/433/986238.

1. Shortliffe, Edward H., et al. "An Artificial Intelligence Program to Advise Physicians Regarding Antimicrobial Therapy." Computers and Biomedical Research 6, no. 6 (1973): 544-560. https://doi.org/10.1016/0010-4809(73)90011-6.

2. Rumelhart, David E., James L. McClelland, and the PDP Research Group. Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Vol. 1. Cambridge, MA: MIT Press, 1986.
3. Russell, Stuart J., and Peter Norvig. Artificial Intelligence: A Modern Approach. Englewood Cliffs, NJ: Prentice Hall, 1995.
4. Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. "ImageNet Classification with Deep Convolutional Neural Networks." Advances in Neural Information Processing Systems 25 (2012): 1097-1105. https://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf.
5. Vaswani, Ashish, et al. "Attention Is All You Need." Advances in Neural Information Processing Systems 30 (2017): 5998-6008. https://papers.nips.cc/paper/7181-attention-is-all-you-need.pdf.
6. Brown, Tom B., et al. "Language Models Are Few-Shot Learners." Advances in Neural Information Processing Systems 33 (2020): 1877-1901. https://proceedings.neurips.cc/paper/2020/file/1457c0d6bfcb4967418bfb8ac142f64a-Paper.pdf.
7. Mehrabi, Ninareh, et al. "A Survey on Bias and Fairness in Machine Learning." ACM Computing Surveys 54, no. 6 (2021): 1-35. https://doi.org/10.1145/3457607.
8. Gunning, David, et al. "XAI—Explainable Artificial Intelligence." Science Robotics 4, no. 37 (2019): eaay7120. https://doi.org/10.1126/scirobotics.aay7120.
9. White House. "FACT SHEET: Biden Administration Launches the AI Bill of Rights to Protect the Rights of American Families." The White House, October 4, 2022. https://www.whitehouse.gov/briefing-room/statements-releases/2022/10/04/fact-sheet-biden-administration-launches-the-ai-bill-of-rights-to-protect-the-rights-of-american-families/.
10. Light, Jennifer S. "When Computers Were Women." Technology and Culture 40, no. 3 (1999): 455-483. https://www.jstor.org/stable/25147356.
11. Ensmenger, Nathan L. "Making Programming Masculine." In Gender Codes: Why Women Are Leaving Computing, edited by Thomas J. Misa, 115-141. Hoboken, NJ: Wiley, 2010.
12. Nilsson, Nils J. The Quest for Artificial Intelligence: A History of Ideas and Achievements. Cambridge: Cambridge University Press, 2010.
13. Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. Deep Learning. Cambridge, MA: MIT Press, 2016.
14. Amodei, Dario, et al. "Concrete Problems in AI Safety." arXiv preprint arXiv:1606.06565 (2016). https://arxiv.org/abs/1606.06565.
15. Whittaker, Meredith, et al. "AI Now Report 2018." AI Now Institute at New York University (2018). https://ainowinstitute.org/AI_Now_2018_Report.pdf.
16. LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton. "Deep Learning." Nature 521, no. 7553 (2015): 436-444. https://doi.org/10.1038/nature14539.
17. Bostrom, Nick. Superintelligence: Paths, Dangers, Strategies. Oxford: Oxford University Press, 2014.

18. Bryson, Joanna J. "The Artificial Intelligence of the Ethics of Artificial Intelligence: An Introductory Overview for Scientists." In Fundamental Issues of Artificial Intelligence, edited by Vincent C. Müller, 211-233. Cham: Springer, 2016.
19. Partnership on AI. "Report on Cooperative AI." (2022). https://www.partnershiponai.org/cooperative-ai-report/

Step 1: Original ChatGPT model. My first order of business was to put the original prompt into ChatGPT and see what response the model gave me, knowing it would be well under the 2500 word minimum. Due to the fact GPT3 usually doesn't go above 1000 words, I wanted to see the bare bones response the initial ChatGPT prompt gave. I liked my initial response, but it was lacking in a few areas, specifically garnered towards course materials we have studied during our time in HISU 2xxx. I liked some of the bones the ChatGPT response gave me, however, specifically the parts focused on ethical concerns, GPT3, and its importance. From there, I knew that I needed to work in two separate ways. The first of my two steps would be to expand my response using a second AI website, to attempt to reach my allotted word limit. After that, I would like to fill in holes in the assignment, especially to ensure topics from the course are talked about as well as personal points of order I would have liked to have seen mentioned.

Step 2: Expansion. In this context, I used ClaudeAI to expand my 800 word draft towards the final goal of 2500-3000 words. I initially decided to see the outcome of placing the initial prompt into Claude, but was unhappy with how different it was than using ChatGPT. So the next course of action from here was to use the initial ChatGPT document and expand on it. I specifically asked Claude to "expand on the initial response with the same word limit as the previous response" (in the response before, I had prompted Claude to hit the 2500-3000 word limit) Despite my prompting, it unfortunately did not reach that specific limit - falling to around 1400 words or so.

From here, my next step was to add in a few topics that I felt like we're not talked about in the initial draft. Specifically, there were three topics I realized needed to be expanded, or talked about entirely. For me, there were three topics: Alan Turing, Women in AI, as well as the 1970s, 1980s, and 1990s. The first one was Turing, who was mentioned for 150 words in the initial response. It was funny, however, because I needed to goldy-locks this response to make sure the length of Turing's section was adequate, but not too long. Claude gave me around a 700 word response when prompted "expand on this section" and then a copy-paste of the initial paragraph - which I had to whittle down to around 250 words.

The next thing to add would be the role of women into the development of AI. The prompt I asked Claude was "add paragraphs describing the role of women in AI, with emphasis towards initial development." From there, I was able to get my desired paragraphs about Lovelace, Hopper, and the ENIAC six. I felt like these women were extremely important in getting to the large language models this paper was assisted by, as already mentioned in course lectures and the texts. However, it did seem like a gap to go straight from the time of Turing straight into the 1950s and 1960s, and I felt like using the women we discussed in this course would be a good addition. Speaking of gaps, in my first prompt, there was also a gap in time

going from the 1960s to the 2000s. While these thirty years may not have been the most important, to truly follow the prompt it felt necessary to add knowledge and advancements from every decade or so, especially as in 1971 technology was extremely different from 1999. With that in mind, I decided to add the prompt "what advancements happened between 1970 and 1990, while mentioning something in every decade."Going from the 1960s all the way to the 2000s would do a disservice to all of the technological developments that happened over those decades. Many of my favorite lectures were focused on this timeframe in course, so, similar to the women in the early development, it was important to further write on this topic.

Lastly, the final "expansion" I made was on the ethics paragraph. With ethics being in the limelight so much over the last few years. Ethics is a key interest of mine and something that I wanted to learn more about, even while working on my paper. I asked Claude to expand on the "ethical concerns of AI, up to your current knowledge." which allowed a few further paragraphs on the topic. It added about 600 words, but I cut out a bit by re-asking the prompt and limiting it to about 500 words. With these expansions done, I was at about 2200-2300 words and needed to add my green to the assignment to ensure I added details the Claude and ChatGPT missed.

Step 3: Details: From here, there were two types of edits that I made when adding additional content to supplement my assignment, which were either A) adding details and my own commentary occasionally to demonstrate course knowledge, or B) adding relevant information that Claude missed that I think would be necessary. The first, was being a bit more detailed in certain places regarding names, dates, and definitions. For example, I added in a few definitions, which included one about neural links. Another thing here I often added occurred regarding the prompt and it being sometimes a bit short handed in a few places. I felt like at times, the paper needed a few words here and there about the importance of something, or who something was founded by. The second thing was to directly add a couple topics into the paper I felt were missed. A few topics in particular I added were about GPT 4, as well as the Biden Administrations ethical concerns around AI, and what the administration is currently doing to help remedy this concern, as well as details on Hopper. These were probably some of my longer additions to the paper, as I felt like paragraphs or similar lengths were important to further comprehend their importance to the prompt or add relevant details to the paper. Secondly, it is interesting to note that both of these additions are more recent news, and something that happened over the last two years (though both are inside of the knowledge base for Claude.)

Step 4: Fun. I wanted to be a bit funny here. I decided to have the first and last paragraphs rhyme. This was just for fun, but I still wanted to have important knowledge demonstrated across both paragraphs. I had Claude for the first paragraph, and ChatGPT for the last paragraph. I used the exact same prompt which was "make these paragraphs rhyme while ensuring a similar amount of words from the initial paragraphs." The reasoning for this was to make a symbolic move, with ChatGPT being the first large language model I used to help write the last paragraph. Claude did a good job with the prompt, but ChatGPT struggled with the "word count" part, and cut about 40% of the words I had. This is still reflected as my last paragraph.