Charlie Bernstein

Walter Issacson

Digitial Revolution

April 14th, 2024

AI Research Paper

The Process:

When I sat down to begin this paper, I typed the following prompt into ChatGPT 4.0: "Describe the development of artificial intelligence from Turing to large language model chatbots in 2500 words." The response I received was suboptimal. The original AI-generated paper did not include anything about Ada Lovelace, Claude Shannon, John McCarthy, Sam Altman, Google, Microsoft, or OpenAI, so, I began to further prompt ChatGPT 4.0 to improve the paper. My next prompt for ChatGPT 4.0 stated, "Incorporate the following people's and companies' contributions to AI; Ada Lovelace, George Boole, John McCarthy, Claude Shannon, Bell Labs, Google, Microsoft, and Sam Altman." Again, the response I got was suboptimal. It was time to change my approach to this paper.

I developed a new strategy for tackling the paper after talking with peers about their experiences experimenting with AI. I learned that to get the best results and the most accurate information, I needed to feed the AI smaller bits of information and write the paper piece by piece.

I engaged the AI in a series of back-and-forth conversations with ChatGPT, Claude, Gemini, Perplexity, and more to gather information and insights on the key events, people, and ideas in the history of artificial intelligence. I started by asking the AI specific questions about important moments and thinkers, like Ada Lovelace in the 1800s and Alan Turing in the 1900s comparing answers to each prompt across each AI. The AI provided me with relatively clear, well-organized summaries that gave me a strong base of knowledge to work from. To my surprise, the answers I was receiving across the different models were astonishingly similar. As I dug deeper into the topic, I asked the AI more complex questions, encouraging it to make connections, highlight important developments and challenges, and think about the big-picture implications of AI for society and our understanding of intelligence. The AI's responses, while speculative, gave me new perspectives and sparked new ideas. Throughout this process, I took an active role in reviewing, fact-checking, and putting together the information the AI provided. I used my own knowledge and critical thinking skills to spot areas that needed more research or clarification, and I kept adjusting my questions to the AI to get more relevant and coherent responses. With all this curated research and AI-generated insight, I asked the AI to help me put it all together into a clear, compelling narrative. By giving the AI specific guidelines on the structure, tone, and logical flow I wanted for the piece, I was able to use its language generation

abilities to create a draft that effectively brought together all the key points from my research. Having an AI to collaborate with and bounce ideas off of challenged my thinking and helped me see this complex topic from new angles, which was invaluable. By learning to work with AI as a research assistant and writing partner, I believe I was able to create a more thorough and insightful piece. Looking back, I realize that the process of writing about AI became a small-scale example of the kind of human-machine collaboration that could shape our future. As I think about the future this technology could bring about, I'm filled with wonder, curiosity, a sense of responsibility, and a drive to do my part in shaping an AI-powered world that brings out the best in us as humans.

The Paper:

The Innovators or AI: A Tale of Genius, Perseverance, and Revolution

Once upon a time, long before we had smartphones in our pockets and computers we could wear as glasses, a group of brilliant thinkers began to ponder a fascinating question: Could machines ever be made to think like humans? This is the story of artificial intelligence, the quest to create machines that can learn, reason, and understand the world the way we do.

Our tale begins nearly two centuries ago with a remarkable woman named Ada Lovelace. Born in 1815, Ada was the daughter of the famous poet Lord Byron, but her true passion lay in mathematics and science. At a time when most women were not encouraged to pursue such interests, Ada's mother fostered her intellectual curiosity, arranging for her to receive tutoring in math and logic. It was through this unconventional education that Ada met Charles Babbage, a British mathematician and inventor. Babbage had conceived of a device he called the Analytical Engine - a mechanical computer that could perform complex calculations. Ada was fascinated by this idea and began corresponding with Babbage, eventually translating and annotating a description of his machine. But Ada's contribution went far beyond mere translation. In her extensive notes, she described how the Analytical Engine could be programmed to handle not just numbers, but any kind of symbol or information. She envisioned a machine that could compose music, produce graphics, and even model human cognition. Ada envisioned the first general-purpose computer, a hundred years before one was built.

Ada's insights were so far ahead of their time that they went largely unrecognized during her lifetime. But her vision of a machine that could manipulate symbols according to rules - the essence of what we now call programming - set the stage for everything that would follow in the world of computing and AI.

Fast forward to the 1930s, and we encounter another pivotal figure in our story: Alan Turing. A brilliant British mathematician, Turing made groundbreaking contributions to the fields of computer science and artificial intelligence. In 1936, he published a paper outlining what came to be known as the Turing Machine - a theoretical device that could perform any computation that could be done by a human following a set of rules.

Turing's work provided a blueprint for the modern computer. He showed that a machine could be designed to handle any task that could be expressed as an algorithm - a series of logical steps. This was a crucial insight that would pave the way for the development of programmable computers. But Turing's interests went beyond mere computation. He was also deeply curious about the nature of intelligence, both human and artificial. In 1950, he published a seminal paper called "Computing Machinery and Intelligence," in which he posed a simple but profound question: Can machines think? To explore this idea, Turing proposed an experiment that came to be known as the Turing Test. In this setup, a human judge would engage in natural language

conversations with both a human and a machine designed to generate human-like responses. If the judge could not reliably tell the machine from the human, Turing argued, then the machine could be said to be "thinking." The Turing Test established a benchmark for artificial intelligence that remains influential to this day. It suggests that the key to creating intelligent machines lies not just in their computational power, but in their ability to communicate and reason in a way that is indistinguishable from a human.

In the years following Turing's groundbreaking work, the field of artificial intelligence began to take shape. The term "AI" itself was coined in 1956 at a conference at Dartmouth College, where a group of scientists gathered to discuss how machines could be made to simulate aspects of intelligence.

One of the key figures at this event was John McCarthy, a young mathematician who would go on to become one of the founding fathers of AI. McCarthy believed that any aspect of learning or intelligence could be so precisely described that a machine could be made to simulate it. He envisioned a future in which machines could not only perform calculations, but could reason, learn, and interact with the world. To pursue this vision, McCarthy and his colleagues developed a new programming language called Lisp, which became a cornerstone of early AI research. Lisp allowed programmers to manipulate symbols and lists in a way that was well-suited to the kind of symbolic reasoning that McCarthy saw as the key to artificial intelligence.

Throughout the 1960s and 70s, AI researchers made significant strides in areas like problem-solving, natural language processing, and computer vision. They developed programs that could solve mathematical problems, understand simple English sentences, and recognize basic shapes and patterns. One of the most influential figures during this period was Marvin Minsky, a cognitive scientist who co-founded the AI laboratory at MIT. Minsky was a proponent of what came to be known as "symbolic AI" - the idea that intelligence could be achieved by manipulating symbols according to formal rules. Under Minsky's leadership, the MIT lab produced some of the most advanced AI systems of the era, including programs that could solve geometry problems and engage in simple dialogues. However, despite these successes, symbolic AI began to run into limitations in the 1970s and 80s. Researchers found that it was extremely difficult to encode all the knowledge and rules needed for a machine to navigate the complexities of the real world.

A new approach began to emerge in the 1980s, inspired by the structure and function of the human brain. This approach, known as "connectionism" or "neural networks," involved building systems composed of many simple processing units, or "neurons," that could learn from experience. One of the pioneers of this approach was Geoffrey Hinton, a British-Canadian researcher who began working on neural networks in the 1970s. Hinton and his colleagues developed a technique called "backpropagation," which allowed neural networks to adjust their internal connections based on the errors they made during training. This was a crucial breakthrough that opened the door to much more powerful and flexible AI systems.

In the 1990s and 2000s, the rise of the internet and the explosion of digital data provided a massive boost to AI research. With vast amounts of information to train on, neural networks began to achieve remarkable results in areas like speech recognition, image classification, and natural language processing. One milestone came in 1997, when IBM's Deep Blue computer defeated world chess champion Garry Kasparov. This was a significant achievement, as chess had long been considered the pinnacle of human intellectual prowess. Deep Blue's victory showed that machines could outperform humans in certain domains of strategic thinking. An even more impressive display came in 2011, when IBM's Watson computer competed on the TV quiz show Jeopardy! against top human champions. Watson used natural language processing and machine learning to understand the often-cryptic clues and buzz in with the correct answers, ultimately winning the competition.

One of the most transformative developments in recent AI history has been the rise of deep learning. This is a neural network-based technique that involves training systems on massive amounts of data. Deep learning has achieved stunning results in fields like computer vision, speech recognition, and machine translation. Some of the most prominent applications of deep learning have come from the major tech companies. In 2012, Google's X Lab developed a system that could learn to recognize cats in YouTube videos, without being explicitly told what a cat looks like. This showcased the power of deep learning to discover patterns in raw data. In 2016, Google DeepMind's AlphaGo system defeated world champion Lee Sedol at the ancient Chinese board game Go. This was a landmark achievement, as Go had been considered an even greater challenge for AI than chess, due to its vast number of possible moves.

More recently, OpenAI, a research company co-founded by tech luminaries Elon Musk and Sam Altman, has been pushing the boundaries of language AI. In 2020, they released GPT-3, a deep learning system that can generate strikingly human-like text, from news articles to poetry to computer code. GPT-3's ability to understand and generate natural language at such a high level has been recognized as the biggest breakthrough in the field.

As AI continues to advance at a rapid pace, it is raising profound questions about the future of work, creativity, and even humanity itself. Some worry about the potential for AI to displace human jobs or be used for malicious purposes. Others see it as a tool that could help solve some of the world's greatest challenges, from disease to climate change. One thing is clear: AI is no longer just a subject of science fiction or academic speculation. It is a real and growing presence in our lives, shaping everything from the ads we see online to the medical diagnoses we receive. As such, it is crucial that we grapple with the ethical implications of this powerful technology. Some of the key issues include bias and fairness (ensuring that AI systems do not discriminate based on race, gender, or other factors), transparency (being able to understand and explain how AI systems make decisions), and safety (preventing unintended consequences or malicious use).

There are also philosophical questions to ponder. As machines become increasingly intelligent and autonomous, will they ever achieve something akin to consciousness or self-awareness? Will they be able to truly understand and experience emotions? And what will be the nature of our relationship with them - will they be our servants, our partners, or something else entirely? These are not easy questions to answer, but they are ones that we must engage with as AI becomes ever more integrated into the fabric of our society. The story of artificial intelligence is, in many ways, the story of our own future - a future that will be shaped by the choices we make and the values we uphold.

Looking back over the history of AI, from Ada Lovelace's prophetic visions to the cutting-edge research of today, a few key themes emerge. One is the power of imagination and curiosity - the drive to ask "what if?" and to pursue ideas that may seem far-fetched or even impossible. Another is the importance of collaboration and diversity. Many of the greatest breakthroughs in AI have come from the meeting of minds across disciplines - mathematicians working with psychologists, computer scientists teaming up with neuroscientists. The more perspectives we can bring to bear on the challenges of AI, the richer and more robust our solutions will be. A third theme is the iterative nature of progress. The history of AI is not a straight line from Ada Lovelace to Alan Turing to today's deep learning systems. Rather, it is a series of fits and starts, dead ends and breakthroughs, each building on and learning from what came before. This is how science and technology often advance - through a process of trial and error, hypothesis and experimentation.

As we stand on the threshold of a new era in artificial intelligence, it is worth keeping these lessons in mind. The future of AI will not be determined by any single person or breakthrough, but by the collective efforts of researchers, engineers, policymakers, and citizens working together to shape this technology for the benefit of all. One important aspect of this collective effort will be addressing the societal impacts of AI. As the technology becomes more powerful and pervasive, it has the potential to greatly influence fields like healthcare, education, finance, and governance. We must consider how AI can be developed and deployed in a way that promotes fairness, accountability, and the public good. It will necessitate the development of new legal and regulatory frameworks to govern the use of AI, ensuring that the technology is used responsibly and in service of human values.

Another key challenge will be preparing for the economic and workforce disruptions that AI may bring. As machines become capable of performing more and more tasks currently done by humans, there is a risk of widespread job displacement. This will require proactive policies to support worker retraining, ensure a strong social safety net, and promote the creation of new jobs in the AI economy. At the same time, AI also has the potential to greatly augment and empower human capabilities. By taking over routine or dangerous tasks, AI can free humans to focus on higher-level, more creative and fulfilling work. For example, the invention of the loom, printing press, and other seminal inventions were met with backlash in fear they would displace jobs. As

we look backwards, we now realize how important these inventions were to the development of humanity.

Ultimately, the impact of AI will depend on the choices we make as a society. Will we use this technology to enhance human potential and address global challenges, or will we allow it to exacerbate existing inequalities and create new risks? Will we develop AI systems that are transparent, accountable, and aligned with human values, or will we rush forward without adequate safeguards and ethical considerations? These are the defining questions of the AI era, and they will require ongoing engagement from all of society. The story of AI is still being written, and we all have a role to play in shaping its future trajectory.

As we navigate this uncharted territory, we can draw inspiration from the pioneers who have brought us to this point - visionaries like Ada Lovelace, Alan Turing, John McCarthy, and Geoffrey Hinton, who dared to imagine a world where machines could think and learn like humans. Their legacy reminds us of the power of human ingenuity and the endless possibilities that emerge when we push the boundaries of what is possible. At the same time, we must proceed with caution and humility, recognizing that AI is a tool - a powerful one, but a tool nonetheless. It is up to us to ensure that this tool is used wisely and for the benefit of all. This will require not just technological innovation, but also moral imagination and a deep commitment to human values. In this sense, the story of AI is not just a tale of machines and algorithms, but a reflection of our own humanity. As we create intelligent systems in our image, we have the chance to reflect on what it means to be human - our creativity, empathy, hopes and fears. And as we shape the future of AI, we are also shaping the future of our society and our world. So let us approach this great frontier with both excitement and responsibility. Let us draw on the wisdom of the past and the potential of the present to create a future where AI empowers rather than replaces us, where it narrows rather than widens the gaps between us, and where it helps us to build a world that is more just, more prosperous, and more humane. For if we can rise to this challenge, the story of AI will not just be a technological triumph, but a human one - a testament to our boundless capacity for learning, for growth, and for using our most powerful tools in service of our highest ideals. And that, in the end, may be the most important chapter of all.

Post Paper Thoughts:

After writing this paper, it posed a question that I wanted to explore further: What does the development of AI mean for the future of education?

I have been pondering this question for quite some time. I have always wondered why people pay hundreds of thousands of dollars for education when you can go online and learn just about anything.

After hearing Sal Khan's talk at the New Orleans Book Festival on Khan Amigo, I knew the education sector is ripe for disruption. I wanted to explore this idea with AI so I have shared the results of my research on this topic below. I askedPperplexity to research this topic. After prompting several times and iterating my process the following paper was produced by Perplexity.

The Disruptive Potential of AI in Transforming Education

Sal Khan, the founder and CEO of Khan Academy, believes that artificial intelligence (AI) has the potential to transform education for the better.[1][2] In his 2023 TED talk, Khan shared his vision for how AI can be leveraged to provide personalized learning experiences for students and support teachers.[1][2]

Khan highlighted how AI-powered tools like Khanmigo, Khan Academy's new AI assistant, can act as a personal tutor for every student and a teaching assistant for every teacher.[1][2][3] Khanmigo is designed to detect students' mistakes, identify misconceptions, and provide effective feedback to help them learn.[2][3] It can also assist teachers by explaining answers, teaching methods, helping with lesson planning, and creating progress reports.[2][3]

This aligns with the findings of Benjamin Bloom's "Two Sigma" study in 1984, which showed that one-to-one tutoring can result in a two standard deviation improvement in student performance.[2] However, providing personalized tutoring at scale has long been a challenge due to cost and scalability issues. Khan believes that AI has the potential to address this "Two Sigma Problem" by making one-to-one tutoring economically viable and accessible to students globally.[2]

In addition to personalized tutoring, Khanmigo is designed to encourage Socratic dialogue and collaborative learning, helping students fine-tune their arguments and improve their writing abilities.[2][3] The platform also includes safety features, such as limiting daily interactions, making chat histories visible to parents and teachers, and using moderation technology to detect inappropriate interactions.[2]

Khan emphasized that while AI in education comes with limitations and risks, the positive use cases should be the focus. He called on everyone to "fight like hell" for the responsible and innovative application of AI in education, as it has the potential to accelerate learning and foster human potential like never before.[2]

Citations:

[1] https://www.ted.com/talks/sal_khan_how_ai_could_save_not_destroy_education?language=en

[2] https://blog.khanacademy.org/sal-khans-2023-ted-talk-ai-in-the-classroom-can-transform-education/

[3] https://www.khanmigo.ai

[4] https://www.youtube.com/watch?v=2ShAkmhG-_E

[5] https://www.youtube.com/watch?v=hJP5GqnTrNo

Citations:

"5 Pioneers in the Field of Artificial Intelligence." Computer Science Degree Hub, February 22, 2021. https://www.computersciencedegreehub.com/lists/5-pioneers-in-the-field-of-artificial-intelligence/.

Artificial Intelligence: Short history, present developments, ... Accessed April 19, 2024. https://www.ll.mit.edu/doc/artificial-intelligence-short-history-present-developments-and-future-outlook-final-report.

Bernard Marr. "14 Incredible Artificial Intelligence Pioneers Everyone Should Know About." Bernard Marr, July 13, 2021. https://bernardmarr.com/14-incredible-artificial-intelligence-pioneers-everyone-should-know-about/.

"The Birth of Artificial Intelligence (AI) Research." The birth of Artificial Intelligence (AI) research | Science and Technology. Accessed April 19, 2024. https://st.llnl.gov/news/look-back/birth-artificial-intelligence-ai-research.

"Evolutionary Computing." Encyclopædia Britannica. Accessed April 19, 2024. https://www.britannica.com/technology/artificial-intelligence/Evolutionary-computing.

Gold, Edem. "The History of Artificial Intelligence from the 1950s to Today." freeCodeCamp.org, May 5, 2023. https://www.freecodecamp.org/news/the-history-of-ai/.

Greenemeier, Larry. "20 Years after Deep Blue: How Ai Has Advanced since Conquering Chess." Scientific American, February 20, 2024. https://www.scientificamerican.com/article/20-years-after-deep-blue-how-ai-has-advanced-since-conquering-chess/.

"Heroes of Deep Learning: Geoffrey Hinton." DeepLearning.AI. Accessed April 19, 2024. https://www.deeplearning.ai/blog/hodl-geoffrey-hinton/.

"John McCarthy." Encyclopædia Britannica. Accessed April 19, 2024. https://www.britannica.com/biography/John-McCarthy.

Karjian, Ron. "History and Evolution of Machine Learning: A Timeline." WhatIs, September 22, 2023. https://www.techtarget.com/whatis/A-Timeline-of-Machine-Learning-History.

Kumar, Anurodh. "Who Are Some of the Pioneers in the Field of Ai?" LinkedIn, July 17, 2023. https://www.linkedin.com/pulse/who-some-pioneers-field-ai-anurodh-kumar.

Marr, Bernard. “The Most Amazing Artificial Intelligence Milestones so Far.” Forbes, December 10, 2021. https://www.forbes.com/sites/bernardmarr/2018/12/31/the-most-amazing-artificial-intelligence-milestones-so-far/?sh=6f67bc867753.

“Marvin Minsky.” CHM, February 12, 2021. https://computerhistory.org/profile/marvin-minsky/.

Muggleton, Stephen. “Alan Turing and the Development of Artificial Intelligence.” *AI Communications* 27, no. 1 (2014): 3–10. https://doi.org/10.3233/aic-130579.

Press, Gil. “114 Milestones in the History of Artificial Intelligence (AI).” Forbes, October 10, 2023. https://www.forbes.com/sites/gilpress/2021/05/19/114-milestones-in-the-history-of-artificial-intelligence-ai/?sh=4057bc8674bf.

Singh, Nidhi, and Namrata Narsinghani. “The Timeline of Artificial Intelligence - from the 1940s.” Verloop.io, October 18, 2023. https://verloop.io/blog/the-timeline-of-artificial-intelligence-from-the-1940s/.

Teneo.ai. “Women in AI: From Pioneers to Modern Leaders.” Teneo.Ai - Transforming every phone call to a love story with your brand, August 1, 2023. https://www.teneo.ai/blog/the-evolution-of-womens-contributions-in-ai-from-pioneers-to-modern-leaders.