William Bai

Walter Isaacson

HISU 2900

14 April 2024

AI Paper

The chatbot's first answer to the initial prompt

The development of artificial intelligence (AI) from the era of Alan Turing to the present day's large language model chatbots is a fascinating journey that spans over seven decades. It reflects a series of innovations, theoretical advancements, and technological breakthroughs. Here's a concise overview of this progression:

Alan Turing and The Foundations

- **1940s-1950s:** Alan Turing, often considered the father of theoretical computer science and artificial intelligence, laid the groundwork with his seminal paper "Computing Machinery and Intelligence" in 1950. He proposed the Turing Test as a criterion of intelligence, a test for a machine's ability to exhibit intelligent behavior indistinguishable from that of a human.

The Birth of AI and Early Excitement

- **1950s-1960s:** The term "Artificial Intelligence" was coined by John McCarthy in 1956, during the Dartmouth Conference. This period saw the development of early AI programs like ELIZA, a natural language processing computer program created by Joseph Weizenbaum, which could mimic conversation by pattern matching users' remarks.

Periods of Austerity and AI Winters

- **1970s-1980s:** Expectations of early AI were high, but the limitations of technology led to periods known as "AI winters," where funding and interest in AI research temporarily receded due to the overselling of AI's capabilities and underdelivery.

The Rise of Machine Learning

- **1980s-1990s:** The field began to shift focus from hardcoded rule-based systems to machine learning, where algorithms improve automatically through experience. This period saw the development of backpropagation for neural networks, which allowed AI to adjust its internal parameters to minimize errors.

The Internet Era and Big Data

- **2000s:** The explosion of the internet and the digitalization of information led to the availability of vast amounts of data. Machine learning techniques, especially deep learning, benefitted enormously from this, leading to significant improvements in AI capabilities.

Breakthroughs in Deep Learning

- **2010s:** AI research saw breakthroughs in deep learning, significantly impacting speech recognition, computer vision, and natural language processing (NLP). Notable developments include the success of deep neural networks in image recognition and the creation of AlphaGo by DeepMind.

The Era of Large Language Models

- **Late 2010s-2020s:** The development of large language models, such as OpenAI's GPT (Generative Pretrained Transformer) series and others like BERT (Bidirectional Encoder Representations from Transformers) by Google, marked a significant leap in AI's ability to understand and generate human-like text. These models are trained on vast datasets, allowing them to generate coherent, contextually relevant text across various domains.

Large Language Model Chatbots

- **2020s:** The culmination of this development is seen in the deployment of sophisticated AI chatbots based on large language models. These chatbots can engage in natural, nuanced conversations, provide information, answer questions, and even **(NB: response was too long for chatbot to load in its entirety)**

A step-by-step description of how you improved the paper:

My paper ended up coming together in three stages. Researching with ChatGPT (GPT-4), I created an outline of idea that I then "wrote" with Google's Gemini. I then edited the paper, finding several unforeseen issues with using these two large language models (LLMs).

The Research Phase

I inputted the initial prompt into OpenAI's Generative Pre-trained Transformer 4 (GPT-4). While it provided a solid introduction to the paper, I realized having GPT-4 generate an entire paper with just one highly engineered prompt would be unrealistically unwieldy. Furthermore, I realized that there were simply too many notable individuals and events in the AI revolution to fit into a 2500 word essay.

I figured that GPT-4 would be great tool to learn about the various events and periods within the AI revolution, from which I could select a few to highlight to depict its general progression. To delve deeper into each era that GPT-4's initial response, I prompted GPT-4 with targeted questions about each section, which enabled me to draft a preliminary outline consisting of specific events. For instance, looking to delve deeper than GPT-4's summarization on Alan Turing, I asked it to "tell me specifically about Alan Turing's contributions to AI," which yielded details on the Turing test and his belief on AI one day surpassing human intelligence.

Having grown up in Silicon Valley and been first introduced to machine learning in 2016, I have a personal interest in the development of AI. For my paper, I wanted to put more emphasis on the technical developments in the last several decades in AI. Thus, during this research phase, I took a deeper dive into the technical elements of AI that I had yet to fully understand.

I had been listening the podcast "Acquired," which explores the histories and origins of prominent technology companies. I became particularly interested in a series of episodes on NVIDIA, where I learned how NVIDIA's evolved from a gaming-focused graphics processing unit (GPU) manufacturer to the market leader in server-grade hardware for AI training. The podcast referenced AlexNet, a groundbreaking innovation from 2012 that revolutionized image recognition. This multi-layered convolutional neural network, I discovered, was trained on a consumer-grade NVIDIA GPU. However, lacking a computer science background beyond basic coding, I struggled to grasp the concept of neural networks. I ask GPT-4 to "explain neural networks, then convoluted neural networks to a tenth grader." GPT-4 explained neural networks clearly and lucidly. It showed me how neural networks operate as interconnected layers of nodes making probability decisions and how convolutional neural networks (CNNs) are a specialized type of neural network designed to identify patterns within data. I also asked it specifically about the significance of AlexNet, learning that while CNNs were not that cutting-edge, AlexNet contributors Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton had taken a novel approach to combine multiple CNN layers with fully connected layers to revolutionize image recognition.

I also ventured a little too deep for relevance to this paper. For instance, I went down a rabbit hole after looking to NVIDIA and CUDA, asking GPT-4 to explain computer architecture and why ARM's architecture is superior to Intel's x86, which was a bit too technical for the scope of the paper. Nonetheless, I found many really interesting facets of computer science through to this

AI-guided exploration, without having to have beyond a surface level understanding of computer science.

The Writing Phase

Having previously used GPT-4 for proofreading and editing, I found its feedback produced a monotonous writing style, with an overuse of third-person active voice. Google's Gemini (then Bard), on the other hand, seemed to offer more flexibility and stylistically seemed more attuned to my style of prose. So, with the outline of topics I created from my GPT-4-aided exploration of the history AI, I turned to Gemini to translate this outline into paragraphs.

My prompts to AI generally followed this template: "Tell me about [xyz trend in AI during abc period] with [xyz and abc events from my outline]. Please keep this to [x] words. For the rest of this thread, please continue to write in the style of Michael Lewis. Please give enough context for the reader." I considered asking Gemini to mimic either Michael Lewis or Walter Isaacson, hoping to emulate a conversational, storytelling tone. While both Newman alums write with compelling style, Isaacson's in-depth narratives take a long time to simmer, while Lewis's style better matched the short length of this essay. To ensure that I could fit all topics within 2500-3000 words, I gave Gemini specific word counts, allotting greater lengths to contemporary to topics like LLMs and GPUs while keeping earlier developments in smaller sections just to provide enough context for the reader. For instance, my prompt for my section on the 40s through 60s started with "Tell me about how the Ideas of AI started with Alan Turing and the Dartmouth Summer Research Project. Please keep this to 400 words."

I found that Gemini was able to provide a narrative that, when read aloud, provided a storytelling flow similar to Michael Lewis's podcast "Against the Rules." However, the "witty" humor I asked Gemini to include was uncharacteristic of Lewis more tame and subtle humor. However, I decided to keep these elements in because it made the writing more enjoyable to read.

The Editing Phase

The Cons of Using AI

I found the editing process to be a lot more tedious than I expected. In fact, many of my edits were not included in the final revision, since I had to cut plenty of material that Gemini generated.

LLMs, while impressively accurate in crafting narratives with technical backing, are prone to some rather silly blunders. I asked Gemini to generate citations for points it mentioned in paragraphs, then turn those citations into a Chicago-style bibliography. As I proofread the bibliography, I realized that Gemini not only forgot to include certain citations but also included ones that did not even show up in the main text. I went back to manually check each citation in the paper and cross-reference it with the bibliography, discovering that I needed to manually add several citations. Gemini may have even fabricated citations as well. One in-text citation referenced a 2017 text regarding AlexNet written Andrej Karpathy, a founding member of OpenAI. I searched the internet for a corresponding publication after finding that it was not

included in the bibliography, and I could not find a one pertaining to AlexNet written by Karpthy in 2017.
Furthermore, I also found that Gemini would often cite too frequently. In one short paragraph about ELIZA developer Joseph Weizenbaum, Gemini created three identical citations in consecutive sentences; normally, only one citation would suffice. In addition, I noticed that Gemini provided citations in subjective statements, such as “the transformer was far superior,” which seems odd. Normally citations should be used in objective statements that back up subjective ones. Finally, I found the placement of certain citations odd. Gemini would often place citations right after a punchline, instead of closer to the sentence referencing a fact.

Another issue was providing human-level detail. Perhaps due to space constraints, Gemini offered superficial information on each point. It either oversimplified complex concepts like GPT or parallel processing or tried (and failed) to inject forced humor. When I requested more detailed versions, Gemini gave overly technical responses—something more fit for entry-level software engineer. There is generally a certain level of detail that I write with, one that encompasses enough specificity to make the reader without overwhelming them. I could not seem to get Gemini to this sweet spot – engaging the reader with slightly technical information without academic density. However, I acknowledge that I could have provided better prompts, and had I practiced more prompt engineering, I might have received something closer to what I was envisioning.

The initial charm of LLM humor wore thin quickly. The "witty" style I requested mainly comprised of nonsensical analogies (a toddler trying to recite Shakespeare after a juice box and a nap) or strangely relatable ones (Funding disappeared faster than free pizza at a college dorm lobby). While the humor itself was engaging, I realized that it was fairly repetitive and at a certain point dry. I had to cut out a lot of stale punchlines; Gemini was trying too hard to make me laugh. Even across different chat threads, Gemini repeated the same catchphrases, like "buckle up, this ride's about to get wild" which I found at least 4 or 5 times as I edited. Possibly due to undiversified training data, Gemini equated "witty humor" with phrases like this. Furthermore, some of the humor seemed age-inappropriate. There are references to scotch and hangovers in the paper (I cut out many others), which might be a playful jab at the alcoholic consumption of certain AI innovators. However, I could not find a definitive link, and this off-beat reference could very well be taken the wrong way.

How I Edited the Paper

Despite these challenges, LLMs proved to be exceptional editing tools. The text you are reading right now (if you have made it this far) has been through several revisions with LLMs. Before LLMs, I would first type all of my thoughts into a very rough draft, and then spend several rounds of revision to make them cohesive and clear. LLMs automate a significant portion of that revision phase. I can now feed them a rough draft, get back a cohesive (albeit robotic) version, and then edit that into rhetoric that reflects my voice and ideas. This significantly streamlines my writing process, although having used this method for the last half year, I find I am significantly lazier with editing and perhaps a little over-reliant on what still is nothing compared to a human editor

I found that Gemini did a fantastic job of crafting the structure of the paper, and I saw little need to make substantial revisions. My edits centered around toning down the rhetoric in certain parts of the piece and inserting other bits of information where needed. For instance, each section begins with a creative and somewhat humorous title. As I edited sections, I realized that the titles had references to content areas I cut out; thus, I had Gemini come up with revised titles for several sections. Additionally, regarding the many comparisons found in the essay, I swapped out several of them for references that were more culturally or contextually relevant. For instance, I mentioned Prometheus in the first paragraph, alluding to the fear of humanity developing AI that leads to its own demise.

I inserted factual snippets that were not included in my initial outline from research using GPT-4. On Alan Turing, I added mentions of his contributions to the Second World War effort, which although tangential to AI gave better context for layperson readers and enhanced the general narrative of his contributions to AI. Additionally, when citing NVIDIA, I mentioned that first person shooter video games were instrumental in driving GPU development because the GPU was originally designed to handle tasks that created heavy graphics processing needs, such as games like Doom. I also alluded to a once underappreciated programming platform called CUDA, which allows developers to leverage GPUs for non-graphic applications. I also added a few more details on BERT to explain how it reads bidirectionally, and I included a few more mentions of contemporary fears of widespread AI usage.

AI Paper: Final Output (Human Edits in Green)

In 1956, a group of pioneers fueled by ambition (and maybe a little scotch) dared to ask a question that would send shivers down the spine of any modern venture capitalist: can we build a machine that thinks? This, my friends, is the origin story of Artificial Intelligence, a field fueled not by cold logic, but by a potent mix of moonshot bets, outsized egos, and a whole lot of scotch. This paper will be your DeLorean, hurtling you back to those heady first days, where the only limit was the fever dream of its creators. We'll explore the early victories – machines conquering checkers, solving logic puzzles – and the inevitable hangover of the AI winters. But through it all, one thing remains constant: the insatiable human urge to create a thinking machine in our own image, to surmount the limits of the human brain in what might be a Promethean journey towards untold excellence. We will find how this story is endlessly fascinating, and occasionally as terrifying as humanity itself.

Bletchley Park to Botworld: How Codebreakers Cracked Nazis and Kickstarted the Robot Uprising (Just Kidding, Probably)

The origins of AI trickled in long before the chatbots of today. Our tale starts with WWII codebreaker Alan Turing. In the 30s, Turing dreamed up the preliminary ideas of a machine that could compute any computable problem. Cracking the Nazis' Enigma cypher at Britain's secretive Bletchley Park, Turing helped shorten the Second World War and saving thousands of lives in doing so. But he wasn't done there. He became convinced that one day machines would surpass human intelligence. To provide a benchmark for this futuristic idea, he proposed the Turing Test: a machine that fools a human in conversation is intelligent (Turing, 1950).

Inspired, across the pond John McCarthy coined "artificial intelligence" and rallied geniuses like Marvin Minsky and Claude Shannon (McCarthy et al., 1955). Their 1956 Dartmouth workshop aimed for a two-month AI brainstorm. It became the Big Bang of AI, the first time expertise from mathematics, psychology, and computer science came together to explore the possibility of machines that could simulate human intelligence.

These pioneers envisioned machines learning languages, composing music – all within a few years! Wildly optimistic, for sure, but it laid the groundwork for AI's fascinating, messy journey.

Talking to Toasters? The 1960s and AI's Quest for (Confusing) Conversation

Imagine the 1960s: a time of hippies, draft-dodgers, LSD, feminists and … AI pioneers? The first step to intelligent machines, many of these innovators believe, was something that could not only think but also hold a conversation – a real mind meld with your toaster, perhaps? Natural language processing, however, turned out to be a different animal altogether. Enter ELIZA, a program developed by Joseph Weizenbaum in the 60s that revolutionized the field, albeit in a way that would make Freud roll in his grave.

ELIZA wasn't omnipotent—it was essentially a glorified fortune cookie program, a master of emotional illusion. Feeling down? Vent to ELIZA, and it might throw back a line like, "Do you think you could bring yourself to talk about how you are feeling lonely?" (Weizenbaum, 1966). Spooky, right? Not quite. ELIZA was just scanning for keywords like "lonely" and regurgitating them in a generic, therapist-speak that wouldn't faze a self-help book (Russell, Stuart, Norvig, 2009).

But guess what? It fooled a lot of people. People poured their hearts out to this glorified thesaurus, a testament to our deep desire for connection – even with a machine that couldn't distinguish a metaphor from a grocery list (McCorduck, 2004). ELIZA was a wake-up call, though. While natural language processing (NLP) was complex and messy, it was also way more fascinating than anyone had anticipated. ELIZA was just the beginning.

The flower-power 60s also brought us SHRDLU developed by Terry Winograd. Forget therapy sessions; SHRDLU offered a virtual world filled with colorful blocks. The twist? You could interact with it in plain English. Tell SHRDLU to "put the red block on the green one," and it would manipulate the blocks accordingly (Winograd, 1972). Sure, SHRDLU's world was tiny, but it was a giant leap for AI. It showed that computers could not only process basic language but also reason about it to interact with the physical world.

ELIZA and SHRDLU –one a master of emotional mimicry, the other a manipulator of virtual Legos. Neither achieved true sentience, but they did something far more important: they opened a Pandora's Box of questions about language, intelligence, and the shrinking gap between humans and machines (Weizenbaum, 1966). Stay tuned, because this story, like a bad case of the hiccups, is far from over.

The 80s: AI Hits Puberty - Awkward, All Potential, and Powered by Backpropagation (Whatever That Is)

But by the roaring 80s, the shine of the AI dream team had come off a bit. Funding disappeared faster than free pizza at a college dorm lobby. AI entered its awkward teenage years: all potential, no real payoff.

But there was flicker of hope: artificial neural networks. These networks were inspired by the human brain, connected in layers resembling human neurons, creating a way for computers to learn by example, just like a toddler picking up on the difference between a cat and a dog. But there was a catch: training these networks required a monstrous amount of computational power, and backpropagation, the algorithm that made it all work, was notoriously tricky.

Enter a trio of AI renegades: David Rumelhart, Geoffrey Hinton, and Ronald Williams. They didn't invent backpropagation, but in their groundbreaking 1986 paper (Rumelhart, 1986), they did something far more important: they showed these neural networks could actually learn cool stuff. They were the Michelangelos of this messy code, sculpting it into something that mere mortals could (sort of) understand. Suddenly, AI could tackle problems that were child's play for humans but a nightmare for computers – things like recognizing faces in images or translating languages. It was far from perfect, though. These neural networks were still like temperamental toddlers, requiring constant training and hand-holding. But for the first time in a long time, AI seem like it could become something truly powerful.

AI's Hardware Hangover: From Chugging Pixels to Parallel Processing Powerhouse - Thanks to Fancy Video Game Gadgets (and Maybe a Little Caffeine)

Remember those beige boxes we called computers back in the day? The ones that needed a pot of coffee just to load a single image. Yeah, those wouldn't stand a chance against the data-hungry beasts of modern AI. The dream of neural networks – those computer brains inspired by our own – was limited by a hardware bottleneck (Minsky, Marvin, Papert, 1969). Central processing units (CPUs), the workhorses of those old computers, were about as exciting as watching paint dry when it came to training neural networks. Imagine trying to teach a sloth how to juggle chainsaws – that's basically what CPUs were like for this job.

Enter the cavalry, riding in on a wave of pixels and polygons: the Graphics Processing Unit (GPU). Originally designed to churn through the massive calculations needed for smooth graphics on first-person shooter video games, GPUs turned out to be a perfect match for AI. Why? Because unlike CPUs, built for tackling tasks one after another like a meticulous accountant, GPUs were all about parallel processing. Imagine a room full of researchers, each assigned a single data point to analyze. That's the serial processing approach of a CPU. Now, picture the same scenario, but with each researcher collaborating on smaller pieces of the data simultaneously. That's the parallel processing superpower of a GPU (Mitzenmacher & Upfal, 2005).

But a GPU alone wasn't the magic bullet. NVIDIA, the company that invented the GPU, in a stroke of genius created CUDA, the catalyst that set the gunpowder ablaze. CUDA unlocked the GPU for non-graphical tasks. Programmers could now unleash GPU's power for general computing, not just fancy graphics (Che et. al, 2008).

Training a massive neural network involves processing an eye-watering amount of data. But luckily, parallel processing on GPUs enables us to digest this data down in two ways.

- **Data Parallelism:** Imagine you have a giant pot of coffee (dataset) that needs brewing. With data parallelism, you split the coffee grounds (data) into smaller mugs (batches). Each overworked grad student (GPU) gets a mug and brews their coffee (trains the model) using the same recipe (model architecture). Finally, everyone pours their coffee back into the pot (combines gradients) to make a stronger brew (better model). This speeds things up because everyone brews at the same time (Dean & Ghemawat, 2008).

- **Model Parallelism:** Rather than letting every grad student brew a small amount of coffee, you break down the process into components like the grinder, filter holder, and brewer (sub-networks). Each overworked grad student (GPU) gets a component and uses the entire pot of coffee grounds (dataset) to brew their part. Then, they combine their partially brewed coffee (sub-network outputs) to create the final, deliciously caffeinated beverage (model output). This lets you brew much stronger coffee (train a larger model) than any single machine could handle on its own (Larochelle et al., 2011).

By combining the raw power of GPUs and the magic of parallel processing, AI researchers finally had the tools to tackle the large-scale models and datasets that were the key to true breakthroughs (Goodfellow, Bengio, & Courville, 2016).

From Pixels to Meaning: The 2010s AI Boom Where Machines Learned to See and Speak (*Almost* Fluently)

The early 2000s were a time of slow internet connections, flip phones, and a lull in AI research. There was exciting potential, but progress felt about as fast as a dial-up connection. Then came a surge of innovation. The groundwork laid by Rumelhart, Hinton, and Williams in the 80s with backpropagation, coupled with powerful GPUs, blossomed into an AI renaissance in the 2010s.

Let's talk AlexNet, the spark that ignited this revolution. Imagine a dog show – how do you tell a pug from a poodle? That's the challenge of image recognition. Super simple for humans, incomprehensibly hard for a machine. In 2012, however, researchers at the University of Toronto tackled this with AlexNet, a deep convolutional neural network (CNN) inspired by the brain (LeCun et al., 2015).

AlexNet wasn't superhuman AI. It was not even that technologically revolutionary. It was simply a massive neural network, trained on a giant dataset of images (think millions of cute dog pictures) (LeCun et al. 2012), with a couple of bells and whistles like ReLU (Rectified Linear Unit) to solve issues like the "vanishing gradient problem" that previous models struggled with. AlexNet learned patterns – edges, shapes – to distinguish a Chihuahua from a Dachshund with impressive accuracy.

Winning the 2012 ImageNet competition was a turning point. AlexNet wasn't just a trick; it proved that deep learning, with the right hardware and data, could tackle complex human tasks (Bengio, 2009).

Fast forward a few years, and the king of natural language processing (NLP) emerged. Developed by Google in 2017, the transformer emerged, making its ancestor ELIZA look like a toddler trying to recite Shakespeare after a juice box and a nap (Vaswani et al., 2017).

Unlike traditional NLP models that processed language word by word, the transformer looked at entire sentences—even entire paragraphs—at once. It used a mechanism called "attention" to focus on the most relevant parts, like underlining key words in a sentence. NLPs now could understand relationships between words, sentence context, and even sarcasm – things that left ELIZA flummoxed. Computers weren't just mimicking language; they were starting to grasp the meaning behind the words.

From Shakespearean Sonnets to Gibberish News Articles: The Hilarious (and Scary) Potential of LLMs

Forget AlexNet, the image-recognition rockstar of the 2010s. The hottest ticket in AI town these days is the Large Language Model, or LLM. Imagine a machine that can not only understand human language but also chat you up with fluency that could almost pass the Turing Test. That's the basic idea behind LLMs. These are colossal neural networks trained on a mountain of text data (Browne et al., 2022) – think books, articles, code, the entire internet – like about the biggest brainiac you know spent 5,000 years learning everything under the sun. They can translate languages, write different kinds of creative content, and even answer your questions in a way that's eerily human (hey, that's kind of what I'm doing here!).

The LLM game has two big players: BERT and GPT. BERT (Bidirectional Encoder Representations from Transformers), developed by Google AI in 2018, is like a master exam-taker who aced every language multiple-choice question ever written (Devlin et al., 2018). It reads sentences both backwards and forwards, excelling at understanding the context of words and their relationships in a sentence. Need to know the sentiment of a text? BERT's your man (or should we say, machine?).

Then there's GPT (Generative Pre-trained Transformer), developed by OpenAI (Radford et al., 2018). Think of GPT as the Eminem of the AI world, spitting out creative lines (that can rhyme) with remarkable fluency.

But here's the catch: LLMs are still under development, and with great power comes great responsibility (cue the Spiderman quote). These machines are fantastic tools, but they can also be tricked into generating nonsensical text or even harmful content. Imagine a news article written by an LLM that starts out factual but devolves into complete gibberish by the third paragraph. Or a news article that misleads readers, leading to perilous consequences. Scary, right?

That's why researchers are grappling with issues of bias and factual accuracy in LLMs. Training data is never perfect, and these models can pick up on the biases and misinformation that exist in the real world (Brundage et al., 2020). It's like that old saying: garbage in, garbage out.

But despite the challenges, the potential of LLMs is, frankly, undeniable. They are revolutionizing how we interact with computers, how we access information, and even how we create art and literature. Every day presents a new application for LLMs and AI, with many excited for the best and others fearing the worst of this paradigm shifting creation. The future of AI is wide open, and LLMs are poised to play a starring role in this ongoing saga.

From Booze to Bots: Looking at the Raucous Ride of AI and its Future (with Existential Dread Included)

AI's journey started with a boozy brainstorming session in 1956, fueled by wild dreams and questionable funding. Early attempts were as impressive as a glorified fortune cookie program (think ELIZA, the 60s chatbot). Then came the awkward teenage years of AI research, until a glimmer of hope arrived: neural networks. These were like trainable toddlers, needing constant attention but holding immense potential.

The key to unlocking this potential turned out to be a video game hero – the Graphics Processing Unit (GPU). Its parallel processing power allowed researchers to build hungry models that could devour data, paving the way for breakthroughs. Deep learning superstars like AlexNet emerged in the 2010s, capable of recognizing a pug from a poodle with mind-blowing accuracy. Now, we have Large Language Models (LLMs) – the Shakespeares of the AI world, generating creative text formats.

The future of AI is a tightrope walk between thrilling possibilities and unsettling dangers. From collaborating with AI to grappling with its potential sentience, the questions we face are enormous. As we navigate this uncharted territory, let's remember the lessons of the past: bold vision, unwavering collaboration, and a healthy dose of caution. Buckle up, because the innovations we make today will shape the incredible, and maybe slightly terrifying, future of AI.

Sources Cited:

Bengio, Yoshua. "Learning deep architectures for AI." Foundations and trends in Machine Learning, vol. 2, no. 1 (2009).

Brundage, Miles, et al. "The malicious use of artificial intelligence: Forecasting, prevention, and mitigation." arXiv preprint arXiv:1802.07228 (2020).

Browne, Julian, et al. "Language models are few-shot learners." arXiv preprint arXiv:2201.08237 (2022).

Che, Shuai, et al. "A Performance Study of General Purpose Applications on Graphics Processors." Journal of Parallel and Distributed Computing, vol. 70, no. 1, 2008, pp. 10–22.

Dean, Jeffrey, and Sanjay Ghemawat. "MapReduce: Simplified Data Processing on Large Clusters." Communications of the ACM, vol. 51, no. 1, 2008, pp. 107-113.

Devlin, Jacob, et al. "BERT: Pre-training of deep bidirectional transformers for language understanding." arXiv preprint arXiv:1810.04805 (2018).

Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. Deep Learning. MIT Press, 2016.

Larochelle, Hugo, Django Mangan, and Andrew Courville. "Learning in Divided Neural Networks." Artificial Neural Networks and Machine Learning - ICANN 2011, Springer, Berlin, Heidelberg, 2011.

LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton. "Deep learning." Nature, vol. 521, no. 7553, pp. 436-444, 2015.

LeCun, Yann, et al. "ImageNet Classification with Deep Convolutional Neural Networks." Advances in Neural Information Processing Systems, vol. 25, Curran Associates, Inc., 2012, pp. 1097-1105.

McCarthy, John, et al. "A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence." August 31, 1955.

McCorduck, Pamela. Machines Who Think: A History of Artificial Intelligence. W. H. Freeman and Company, 2004.

Mitzenmacher, Michael, and Eli Upfal. Probability and Computing: Randomized Algorithms and Probabilistic Methods. Cambridge University Press, 2005.

Minsky, Marvin, and Seymour Papert. Perceptrons: An Introduction to Computational Geometry. MIT press, 1969.

Radford, Alec, et al. "Improving Language Understanding by Generative Pre-Training." arXiv preprint arXiv:1801.06146, 2018.

Russell, Stuart, and Peter Norvig. Artificial Intelligence: A Modern Approach. Third Edition, Pearson Education Limited, 2009.

Rumelhart, David E., Geoffrey E. Hinton, and Ronald J. Williams. "Learning representations by back-propagating errors." Nature, vol. 323, no. 6088, pp. 533-536, 1986.

Turing, A. M. "Computing Machinery and Intelligence." Mind, vol. 59, no. 236, 1950, pp. 433-460.

Vaswani, Ashish et al. "Attention Is All You Need." Neural Information Processing Systems, 2017.

Weizenbaum, Joseph. "ELIZA - A Computer Program for the Study of Natural Language Communication." Communications of the ACM, vol. 9, no. 1, pp. 36-45, 1966.

Winograd, Terry. "Understanding Natural Language." Cognitive Psychology, vol. 3, no. 1, 1972, pp. 1-191.