

William Gelb and Brianna Kirsch

**Prompt:** *Describe the development of artificial intelligence from Turing to large language model chatbots.*

Step 1: Chat Gpt 3.5 Response

The development of artificial intelligence (AI) from Alan Turing's groundbreaking work to the emergence of large language model chatbots represents a journey marked by significant milestones, scientific breakthroughs, and technological advancements. This essay will explore the key stages in this evolution, highlighting the contributions of various researchers, and the transformative impact of AI on society.

Alan Turing, often regarded as the father of theoretical computer science and AI, laid the foundation for modern AI with his seminal paper "Computing Machinery and Intelligence" published in 1950. In this paper, Turing proposed the famous Turing Test, which assesses a machine's ability to exhibit intelligent behavior indistinguishable from that of a human. While Turing's ideas were visionary, the computational resources required to realize them were limited at the time.

The subsequent decades witnessed incremental progress in AI research, characterized by the development of expert systems, rule-based systems, and early neural networks. In the 1960s and 1970s, researchers explored symbolic AI, focusing on representing knowledge and reasoning symbolically. However, these approaches faced limitations in handling uncertainty and complexity, hindering their practical applicability.

The 1980s marked a shift towards connectionism, with researchers exploring neural network models inspired by the structure and function of the human brain. Pioneering works by Geoffrey Hinton, Yann LeCun, and Yoshua Bengio laid the groundwork for modern deep learning techniques. Despite initial challenges, such as the vanishing gradient problem, neural networks gradually gained traction due to advancements in computing power and the availability of large datasets.

The breakthrough moment for AI came in 2012 when AlexNet, a deep convolutional neural network, achieved unprecedented performance in the ImageNet Large

Scale Visual Recognition Challenge, significantly surpassing traditional computer vision methods. This success catalyzed the adoption of deep learning across various domains, including natural language processing (NLP).

In 2013, Tomas Mikolov and his colleagues at Google introduced Word2Vec, a neural network-based technique for word embedding, which revolutionized the field of NLP by enabling machines to capture semantic relationships between words. This was followed by the development of recurrent neural networks (RNNs) and long short-term memory (LSTM) networks, which excel at sequential data processing, making them well-suited for tasks like language translation and text generation.

The year 2018 witnessed a significant milestone with the introduction of OpenAI's GPT (Generative Pre-trained Transformer) model. GPT demonstrated remarkable proficiency in generating coherent and contextually relevant text by leveraging large-scale unsupervised learning on vast text corpora. The subsequent iterations, such as GPT-2 and GPT-3, further refined this capability, achieving human-level performance on various language tasks.

Large language model chatbots, powered by models like GPT-3, represent the culmination of decades of AI research and development. These chatbots are capable of engaging in natural and contextually relevant conversations, demonstrating a nuanced understanding of language and context. They find applications in customer service, virtual assistants, content generation, and educational tools, among others.

However, the rise of large language model chatbots also raises ethical concerns regarding biases, misinformation, and the potential misuse of AI-generated content. Ensuring responsible deployment and governance of AI technologies remains a pressing challenge for researchers, policymakers, and industry stakeholders.

In conclusion, the development of artificial intelligence from Turing to large language model chatbots has been a remarkable journey marked by scientific innovation, technological advancement, and societal impact. While the journey has been fraught with challenges and ethical considerations, the potential of AI to transform industries and enhance human experiences is undeniable. As we continue to push the boundaries of AI

research, it is imperative to pursue responsible AI development that aligns with ethical principles and societal values.

### Step 2: The Additional Prompts Given:

A. **What prompts would you give ChatGPT 3.5 to get the best essay, whether that be other ideas to interject into the essay, make it sound better, or make it more substantive? Give me ten ideas.** Interesting note- All 10 ideas were AI related as in nothing about focusing more on women, race, or any other specific stage of development.

B. **Implement your idea number three: Examine the societal implications of AI advancements, including workforce automation, job displacement, and socioeconomic inequality, and propose strategies for mitigating negative impacts. Do this while implementing facts about the Luddite movement.** While it added to the original prompt, it did not do so in essay format.

C. **I meant to add your idea into your original essay that you previously gave me.** They still did not give me the essay.

D. **I want you to write two paragraphs adding onto your response exploring Ada Lovelace and Charles Babbage's role in the original prompt of which I will remind you of "Write me an essay describing the development of artificial intelligence from Turing to large language model chatbots."** This was a complete failure and only gave me a two paragraph response which failed to 1. Write me an essay. 2. Forgot about the Luddite movement and their idea "three".

E. **Please combine all of your responses into an essay. I did not believe this would work but chose to try despite that.** It gave me half an essay then devolved into short paragraphs summarizing what it already said. Looking back at this, Chat GPT forgetting previous responses and prompts was a huge frustration during this project.

F. **Did you get lazy? Why did you stop mid essay?** Out of frustration and morbid curiosity I asked this and ChatGPT apologized and finished the essay and even included information about the Luddite movement. Unfortunately, the essay was scattered and it seemed to have

forgot about Lovelace and Babbage. At this point I decided on hopefully a better approach in doing it section by section then attempting to string it together at the end. When I told chat gpt to **add the contributions of women back into the essay**, I didn't like how it added a paragraph saying "additionally, women also contributed..." instead of just including women into the essay with men. I believe this is because Chat GPT would never write an essay about women and then say "additionally, men also contributed..." I believe that the algorithm may be slightly sexist. I realized that I had my doubts about Chat GPT. While studying genetics, I realized that Chat GPT very often second guesses itself. I would ask it a question and give it answer choices knowing the answer ahead of time— I realized I need to be skeptical of hallucinating. So this promoted my next question.

G. **Are you sure about the women's contributions?** I was beyond shocked. Chat GPT said that the role of Ada Lovelace was essential in the previous essay but then corrected her role as overstated and that her role was exaggerated (Im sure Professor Isaacson would not appreciate that one!)

H. I took the current prompt I had and conversed with my partner. We decided it was time to end Chat GPT and move to Claude. My prompt was, **"I am going to show you my essay so far. I want you to add as much detail as possible...go in depth!"** I was more than pleased with the result. It has added things that were previously never mentioned, like the AI winter, impacts of neural networks on the development of AI, and more.

I. The tone felt a little too formal...I like easy to read and digested content. So... I asked **"rewrite this essay in a girl-boss tone without losing accuracy and detail"**. I immediately loved how it turned out. But it was only 1,400 words and I needed to pay in order to get more trials. I realized that I did not want to sacrifice tone for words. So, I told it, **"Take away the girlboss attitude and write in normal chat bot language."**

J. I took my work away from Claude and went to Chat GPT 4.0 for which my partner had a subscription to. I started with the prompt **"Integrate a timeline of the AI winter where it seems most logically prudent."**

K. While the essay was much improved with GPT 4.0, it went back to it's old ways of blocking the sections up in a robotic way. I also wanted the chat box to focus more on topic

of the original prompt. So, I fed it the prompt: “Change this more into essay format with organized paragraphs, intro conclusion and make sure its still answering the original prompt– ‘Describe the development of artificial intelligence from Turing to large language model chatbots.’” This failed horribly. The new essay was shorter, had subtitles, and I prompted “Try again by writing a cohesive essay more similar to what you had before with every topic covered and fill in any gaps of information that you see.”

L. “How do I make this a 2,500 word essay” after I gave it this response, it proposed a structure on how to expand on every topic and go more in depth. However, I wanted GPT 4.0 to incorporate this new structure into the essay itself so that I did not have to do it manually.. I fed the prompt, “can you take my essay that you just generated above and do this exact structured approach for me to produce approximately 2,500 words.”

That was a spectacular failure. One thing that I noted about the Chat GPT was that it said it would add words, however, it did not. Giving it a word count was basically useless.

M. I want you to expand on each point with additional paragraphs about women's contributions, ENIAC, Bell Labs, or anything from Walter Isaacson's Book "The Innovators" to make 4,000 words. Make every paragraph 3 times as long by adding details that you find relevant. Again, telling Chat GPT to include more words was useless but telling it to expand on paragraphs was slightly more useful.

N. For the first time, we had some success. We got the prompt within GPT 4.0 “Continued generating” twice. Still, we were about 800 words off from our target 2,500 word count. We gave the prompt “Are you forgetting anything? If so, add it.”

O. Add in a paragraph wherever you believe would be most logical about how war forced the advancement of technology...expand on that. It seemed like we hit another stride and wasn't planning on stopping there. Testing GPT 4.0 seeing where it would think adding the paragraphs where it would be most cohesive. Add those three paragraphs into your large essay above where you believe to be most logical.

P. Deciding to follow our previous winning strategy, we prompted “include in this same essay a paragraph or two about each woman and focus a lot on Lovelace. Try to make the

contributions from women at least 400 words more.” Continuing the strategy, our next prompt was “integrate these paragraphs into the essay that you previously generated.”

Q. At this point, the essay was good but needed tweaking to get what we wanted. The next prompt used was “Keep the word length but make it more of a cohesive blend of ideas rather than topic by topic. get rid of titles and subtitles.” This did not work as planned and instead of following the prompt, GPT 4.0 lost track of the original intentions of the paper while refusing to remove the subtitles.

R. So I decided to give GPT 4.0 one huge prompt. Those are all good things to do, however I do need the length to return. Can you add about the earliest successful AI programs written by Christopher Strachley, prototype IBM, GPS, IPL, PROLOG, expert systems and the chess game while keeping it cohesive and eliminating subtitles. Before you give me a response, give me a short response making sure you understand exactly what I am prompting you. While doing this, ensure that your word count is roughly 4,000 words. Do this by integrating the new prompts into your previously generated essay.

S. This was by far our most successful prompt. Reflecting on our experience with this project, my partner and I wish we had attempted to give it more detailed prompts from the beginning. Still, we needed to add just a little bit more, Thats great, thank you so much! Can you just expand on each paragraph by two sentences?

Overall, this project may have been more difficult than writing the essay normally (in some ways). Some of the biggest struggles were reaching the word count and the forgetfulness of Chat GPT. We overcame these struggles by deciding to focus on sections rather than just having it give us advice on how to make the essay longer and then telling Chat GPT to do it for us. When we gave it ideas to talk about, such as including a paragraph about women and things ranging from prototype IBM to AI Chess, the essay expanded incrementally. Some things Chat GPT hallucinated were the sources. It gave its best estimate of what sources it used to create the final essay, but some of the sources were incorrect. For instance, it made slight mistakes when the book “Grace Hopper: Admiral of the Cyber Sea” was published. It was published in 2013 by the specific book company but it noted that it was published in 2012. Subtle mistakes like this exemplified that chat GPT was either getting its information

from an unknown source or it was misfed this information in its creation. Chat GPT is simply a product of the programming it is fed. Also, we noticed that there was no significant difference between GPT 3.0 and GPT 4.0. When given the same prompts, it yielded the same results. The programming lacks good memory as to what prompts it is answering which means that we had to be extremely attentive to what it was taking out each response. The essay could not have possibly been the length that it was without constant input and some prior knowledge on the topic. The final essay was to our pleasure but it took a substantial amount of time to reach it.

### Step 3: Final Essay (Exactly 2500 words)

#### **Title: The Evolution of Artificial Intelligence: From Turing to Modern Chatbots**

Artificial Intelligence (AI) has evolved from a nascent theoretical concept into a cornerstone of modern technological innovation, impacting various sectors from healthcare to finance. This journey, from Alan Turing's foundational concepts to today's sophisticated AI systems that seamlessly interact with humans, showcases a dramatic transformation filled with groundbreaking milestones. The development of AI has not been linear but a complex array of innovations, challenges, setbacks, and remarkable recoveries, each contributing to the layers of this evolving technology. This essay explores these developments in detail, highlighting seminal events and the role of pioneering technologies and individuals that have significantly shaped the AI landscape. Through this exploration, the essay will provide an insightful look into how computational ideas became tangible technologies that are integral to daily life and will discuss the implications of these developments for future technological evolution, emphasizing the interdisciplinary nature of AI advancements and their profound impact on society.

Alan Turing, often celebrated as the father of modern computer science, laid the theoretical groundwork for what would become artificial intelligence. In his seminal 1950 paper, "Computing Machinery and Intelligence," Turing proposed the concept now known as the Turing Test—a criterion that assesses a machine's ability to exhibit intelligent behavior equivalent to or indistinguishable from that of a human. This was a provocative challenge to the capabilities of then-current technology and a significant philosophical

inquiry into the nature of intelligence. Turing's contributions during the mid-20th century were profound, proposing that machines could one day think independently. His work sparked global interest and catalyzed subsequent research in machine learning and AI, setting a theoretical benchmark that continues to influence AI goals today. Furthermore, Turing's broader contributions to computational theory laid the groundwork for the binary system of zeroes and ones that underpins all modern computer processes, illustrating his integral role in the development of computing technology. His ideas not only initiated the computational revolution but also set the stage for ethical discussions regarding the future of intelligent machines.

The field of computing and artificial intelligence owes a significant debt to pioneering women who contributed foundational ideas and implementations, often in the face of considerable obstacles. Ada Lovelace, an English mathematician in the 19th century, is often celebrated as the world's first computer programmer. Working closely with Charles Babbage on the development of the Analytical Engine, Lovelace described an algorithm intended to be processed by the machine, effectively creating the first computer program. Her foresight into the capabilities of computing extends to her speculation that such machines could one day manipulate symbols and create music or art, predicting the potential uses of computers in graphics and other digital arts. Grace Hopper, another trailblazer, developed the first compiler, a program that translates written code into computer language, which was foundational for the creation of COBOL, one of the first high-level programming languages. Her efforts democratized computing, making it more accessible and significantly influencing the development of user-friendly programming languages that are essential in AI development. Hopper's innovations in making programming languages closer to natural language opened new avenues for non-specialists to engage with computing, thus expanding the pool of intellectual contributions to the field.

During World War II, Jean Bartik and a group of skilled female mathematicians were among the first to program the ENIAC (Electronic Numerical Integrator and Computer), setting up the machine to solve complex ballistic trajectories manually. Katherine Johnson, known for her work in mathematics and aerospace engineering, also played a crucial role in the use of computers for space missions. Her precision in



calculating flight trajectories and pioneering use of computers to validate her manual calculations were critical in a time when computers were still new and untrusted. These women provided critical foundations for the development of computer programming and AI, expanding the scope of what computers could accomplish and helping to shape the future of machine capabilities. Their work not only broke technological barriers but also paved the way for more inclusive work environments in the tech fields, challenging the prevailing norms of their times.

The journey of AI in the 1950s and 1960s included significant contributions from various early AI programs and languages that set the stage for later advancements. Christopher Strachey, a British computer scientist, wrote the first successful AI program to play a complete game of draughts (checkers) in 1951. This program, running on a prototype of the Ferranti Mark I computer, demonstrated the potential for machines to engage in complex decision-making tasks, a foundational concept in AI. Strachey's work laid the groundwork for the exploration of AI in gaming and problem-solving contexts, illustrating early on the potential for machines to perform tasks traditionally associated with human intelligence. Around the same time, the development of the IBM 701 and subsequent models marked the beginning of commercially available computers that could handle symbolic information, not just numerical calculations. This capability was crucial for the development of AI as it allowed for the processing and manipulation of symbolic data, essential for AI applications. The IBM machines' ability to process vast amounts of data and perform complex calculations more efficiently opened new possibilities for AI research, providing the necessary hardware to support increasingly sophisticated algorithms.

In the realm of AI-focused languages, John McCarthy's development of Lisp in the late 1950s, and the subsequent creation of the Information Processing Language (IPL), provided the necessary tools for coding complex AI algorithms. Lisp, in particular, became synonymous with AI research due to its features that supported symbolic computation and recursive algorithms, features that are particularly suited for AI problem-solving. This period also saw the introduction of the Logic Theorist and General Problem Solver (GPS), developed by Allen Newell and Herbert Simon. These programs were designed to mimic

human problem-solving skills and were among the first to attempt to generalize AI applications beyond specific tasks. This was followed by the development of Prolog in the 1970s, a programming language designed specifically for AI, which allowed for more efficient handling of "rules" about the world, which are fundamental for developing expert systems. Expert systems, a type of AI program that emulates the decision-making ability of a human expert, became prominent in the 1970s. These systems were designed to solve complex problems by reasoning through bodies of knowledge, represented mainly as if-then rules rather than through conventional procedural code. The expert system known as DENDRAL, which was used for chemical analysis, and MYCIN, developed to diagnose blood infections, were among the first successful implementations in this field. The development and success of these early systems demonstrated AI's potential not only in replicating human thought processes but also in surpassing human ability in specific domains, thereby expanding the scope of what could be achieved with AI.

AI's capabilities were further showcased in the domain of games, particularly chess. The development of chess-playing programs served as an important benchmark for AI capabilities and progressed significantly in the 20th century. The match between IBM's Deep Blue and world chess champion Garry Kasparov in 1997, where a computer defeated a reigning world champion under standard chess tournament conditions, marked a significant milestone. This event demonstrated the potential of AI to not only replicate but also surpass human cognitive abilities in specific domains. Deep Blue's victory was not just a demonstration of raw computing power; it also showcased the sophisticated integration of algorithms, machine learning techniques, and tactical understanding that could evaluate millions of positions. This historic win by Deep Blue helped to dispel many doubts about AI's potential and sparked a new wave of interest in what AI could achieve in the future.

The profound impact of war on technological advancement, particularly in the realm of computing and AI, is a significant aspect of technological history that often goes underappreciated. Historically, many crucial technological innovations have been propelled by the exigencies and pressures of wartime needs, which typically prioritize rapid, efficient solutions to complex logistical and strategic problems. The development of ENIAC, one of the first electronic general-purpose computers, serves as a quintessential example. Initially

funded and constructed during World War II, ENIAC was designed specifically to calculate artillery firing tables for the United States Army. This task, critical for the accuracy of long-range weapons, required the rapid processing of large amounts of data, a need that significantly accelerated the development of computing technology. Moreover, the impact of war on AI and computing extends beyond World War II. During the Cold War, the intense rivalry between the superpowers spurred further advancements in computing technology, which were vital for intelligence gathering, cryptography, and missile control systems. The space race, part of this broader geopolitical struggle, led to substantial investments in technology that required complex computational capabilities, pushing forward the boundaries of what was technically feasible at the time. The demands for robust, fast, and reliable processing systems led directly to innovations in semiconductor technology, such as integrated circuits, which were essential for miniaturizing and enhancing the performance of military equipment and later became fundamental components of personal computers and AI systems. These periods of intense conflict necessitated the development of technologies that could perform under pressure, manage vast datasets, and execute complex computational tasks reliably. The results of these military-driven innovations have had far-reaching implications. Once developed, these technologies found applications beyond the military, in fields ranging from healthcare to finance, and significantly influenced the development of artificial intelligence. By providing robust hardware platforms and complex problem-solving frameworks, wartime technological advancements laid critical groundwork for the development of AI, demonstrating how conflict-driven needs can inadvertently foster broader technological progress that benefits civilian and scientific domains. This dynamic underlines the complex interplay between societal needs, conflict, and innovation, highlighting how periods of strife can accelerate technological development with lasting impacts on artificial intelligence and beyond.

The periods known as AI winters were marked by a significant reduction in funding and interest in artificial intelligence research. The first AI winter occurred in the mid-1970s, triggered by the Lighthill Report, which criticized the inflated expectations of AI and the lack of progress in accomplishing the grandiose promises of the 1960s. A second AI winter occurred in the late 1980s after the collapse of the Lisp machine market

and a series of failed AI projects, which led to a general disenchantment with AI among government and funding agencies. Despite these setbacks, the foundational research conducted during these periods helped pave the way for future successes. For instance, during these winters, researchers refined algorithms, improved models, and addressed some of the limitations identified during earlier explorations. The commitment to underlying research during these periods kept the field slowly advancing, setting the stage for the resurgence of interest that followed as computational power increased and new algorithms were developed. The resilience and continued evolution of AI through these periods highlight the cyclical nature of scientific advancement and the importance of sustained investment in research and development, even in the face of setbacks and disillusionment.

The early 21st century marked a pivotal change in AI research with the revival of neural networks through deep learning. This new approach to neural networks, characterized by layers of algorithms and vast interconnections, was inspired by the architecture of the human brain. Key to this resurgence was the development of backpropagation and convolutional neural networks by researchers such as Geoffrey Hinton, Yann LeCun, and Yoshua Bengio. These techniques allowed neural networks to adjust their internal parameters in ways that minimized errors, and they were particularly effective for tasks involving large amounts of data, such as image and speech recognition. The availability of large datasets and powerful GPU computing fueled these advancements, enabling the training of larger and more complex models. This period also saw the integration of neural networks with reinforcement learning, which led to the development of systems capable of superhuman performances in complex games like chess and Go, symbolizing the significant advancements in AI capabilities. The success of deep learning has reinvigorated the AI field, leading to new applications and a broader acceptance of AI technologies in everyday life.

Natural Language Processing (NLP) has seen significant advancements due to deep learning technologies, dramatically improving machines' ability to understand and generate human language. The introduction of sequence-to-sequence models and attention mechanisms enabled more effective handling of long-range dependencies within text, crucial for applications such as machine translation and text summarization. The

development of the Transformer model by researchers at Google introduced an architecture that relies entirely on attention mechanisms without the need for recurrent layers. This model set the stage for the development of more sophisticated systems, such as OpenAI's GPT and Google's BERT, which utilize vast amounts of text data to generate models capable of understanding context, sentiment, and nuance in language. These models have transformed how machines handle language, providing the backbone for AI applications in content creation, customer service, and even in generating written works, pushing the boundaries of AI's capabilities in understanding and replicating human language. The advances in NLP not only enhance the functionality of AI in various applications but also deepen our understanding of language as a human attribute, bridging the gap between human and machine communication.

The introduction of Generative Pre-trained Transformers (GPT) by OpenAI has marked a significant evolution in the field of AI, particularly in the domain of natural language processing. The GPT models, beginning with GPT-1 and extending to GPT-3, have demonstrated remarkable language understanding and generation capabilities. These models are trained on a diverse range of internet text and are fine-tuned for various tasks without task-specific data training, making them incredibly versatile. They can generate coherent and contextually relevant text across many domains and styles, from conversational AI in chatbots to complex problem solving in academic fields. This level of performance has made GPT models central to many current AI applications, offering a glimpse into future possibilities where AI can assist in creative processes or provide detailed technical assistance. The development of GPT and other advanced models signifies a major leap forward in AI's capability to interact naturally with humans, enhancing the user experience and expanding the ways in which AI can be utilized in everyday activities.

The evolution of AI from the theoretical proposals of Alan Turing to the sophisticated capabilities of modern chatbots is a testament to the remarkable progress in computational technologies and algorithms. This journey highlights not only the technological advancements but also the collaborative efforts of individuals from diverse backgrounds and disciplines. As AI continues to advance, it presents new challenges and

opportunities that require thoughtful consideration. The integration of AI in various sectors of society necessitates discussions about ethics, policy, and the future role of human labor. This evolution of AI is not merely a technological narrative but a profound shift in the interface between humans and machines, promising to redefine the boundaries of what machines are capable of achieving. Looking forward, the continued evolution of AI will likely usher in new innovations that will further blur the lines between human and machine capabilities, reshaping the technological landscape in unimaginable ways.

### References:

Alan Turing's seminal work and the Turing Test:

- Turing, A. M. (1950). "Computing Machinery and Intelligence." *Mind*, LIX(236), 433-460. <https://doi.org/10.1093/mind/LIX.236.433>

Contributions of Ada Lovelace:

- Toole, B. A. (1998). "Ada, the Enchantress of Numbers: Prophet of the Computer Age." Strawberry Press.

Grace Hopper and the development of compilers:

- Williams, K. (2012). "Grace Hopper: Admiral of the Cyber Sea." Naval Institute Press.

Jean Bartik and the ENIAC programmers:

- McCartney, S. (1999). "ENIAC: The Triumphs and Tragedies of the World's First Computer." Walker Publishing Company.

Katherine Johnson's contributions to spaceflight and computing:

- Shetterly, M. L. (2016). "Hidden Figures: The American Dream and the Untold Story of the Black Women Mathematicians Who Helped Win the Space Race." William Morrow & Company.

Development of early AI languages and systems (LISP, Prolog):

- McCarthy, J. (1960). "Recursive Functions of Symbolic Expressions and Their Computation by Machine, Part I." *Communications of the ACM*, 3(4), 184-195.
- Roussel, P. (1993). "Prolog: Manual de référence et d'utilisation." Groupe d'Intelligence Artificielle, Marseille-Luminy.

Expert systems and their applications:

- Buchanan, B. G., & Shortliffe, E. H. (1984). "Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project." Addison-Wesley.

Chess AI (Deep Blue vs. Kasparov):

- Hsu, F. H. (2002). "Behind Deep Blue: Building the Computer that Defeated the World Chess Champion." Princeton University Press.

Developments in modern AI (Deep Learning, NLP, GPT):

- Goodfellow, I., Bengio, Y., & Courville, A. (2016). "Deep Learning." MIT Press.
- Vaswani, A., et al. (2017). "Attention Is All You Need." In *Advances in Neural Information Processing Systems*.
- Brown, T. B., et al. (2020). "Language Models are Few-Shot Learners." In *Advances in Neural Information Processing Systems*.

General references on AI history and development:

- Russell, S., & Norvig, P. (2016). "Artificial Intelligence: A Modern Approach." 3rd Edition. Pearson Education.
- Isaacson, W. (2014). "The Innovators: How a Group of Hackers, Geniuses, and Geeks Created the Digital Revolution." Simon & Schuster.