Benjamin Porter (and ChatGPT 3.5)

AI Research Paper: Draft

The development of artificial intelligence (AI) from Alan Turing's conceptual ideas to modern large language model chatbots like me has been a fascinating journey spanning several decades. Alan Turing laid the groundwork for AI with his seminal paper "Computing Machinery and Intelligence" in 1950. He proposed the Turing Test as a measure of a machine's intelligence, suggesting that if a machine could exhibit behavior indistinguishable from that of a human, it could be considered intelligent. The 1950s and 1960s saw the development of early AI programs focused on tasks like symbolic reasoning, game-playing (like IBM's chess-playing computer in the 1960s), and logical inference. Expert systems emerged in the 1970s and 1980s, which were AI programs designed to mimic the decision-making abilities of human experts in specific domains. These systems used rules-based approaches and symbolic reasoning. The field of machine learning gained prominence, with algorithms like neural networks, decision trees, and support vector machines being developed. Neural networks, in particular, saw renewed interest with the development of backpropagation algorithms. NLP advanced significantly with techniques like statistical machine translation, part-of-speech tagging, named entity recognition, and sentiment analysis. Systems like IBM's Watson showcased the power of AI in understanding and processing human language. The 2010s witnessed a revolution in AI fueled by deep learning techniques, especially deep neural networks with many layers. This led to breakthroughs in computer vision, speech recognition, and NLP tasks. Large language models like OpenAI's GPT series (including me, GPT-3) represent the culmination of advances in deep learning, NLP, and large-scale data processing. These models leverage transformer architectures and vast amounts of data for tasks such as text generation, translation, summarization, and conversation. Throughout this evolution, AI has transitioned from rule-based systems to data-driven approaches, with increasing complexity and capability. Ethical considerations, bias mitigation, and responsible AI practices have also become integral parts of AI development and deployment.

AI Research Paper: Process

After entering the initial prompt into ChatGPT 3.5 and receiving a draft, I decided to break the paper apart into sections to ensure that ChatGPT could help me reach the proper word count (since the chatbot can only spit out sections ranging from 300-1000 words). I asked ChatGPT to formulate some proposed "sections" to focus on from Turing's developments all the way up to large language model chatbots, and it came up with seven: an introduction, early foundations of AI, the evolution of machine learning, advancements in natural language processing, the deep learning revolution, large language models and chatbots, and a conclusion that explored future prospects and challenges. Before beginning the process of crafting the first section, I made sure to give ChatGPT some general instructions and guidelines that I planned to modify and expand over time. I told ChatGPT to avoid general, vacuous writing, focusing on quality, in-depth, example-heavy analysis, and I emphasized the importance of providing examples that I could research, fact check, and source: I also requested a word count of around 400 to 600 words for each section. I was aware at the beginning of the process that ChatGPT can often spit out incorrect information or simply disregard instructions (especially about word count), so I was prepared to go through a long and arduous process. After getting a draft of the introductory section, I already spotted some issues that I wanted to eliminate for the remainder of the paper. ChatGPT already used examples, which I didn't need, so I told the chatbot to instead focus on providing a general "road map" for the rest of the paper, but to still follow my instructions for following sections. After getting a much better start for an introductory section, I went in myself and wrote in information that focused on the purpose of the paper and my thoughts on the unique collaborative perspective that ChatGPT and I could provide. I also edited it to fit a writing style that flowed better and fit my personal preferences.

Next, I moved on to the first main section, early foundations of AI. After receiving an initial draft from ChatGPT based on the instructions I gave earlier, I began making modifications. First, since the section was less than 300 words, I told ChatGPT to expand to my previous stated preference of around 500 words. Next, since the chatbot focused far too much on a biography of Turing, perhaps because I had emphasized the importance of citable examples, I told the chatbot to focus primarily on his contributions to the development of AI. ChatGPT then kept giving drafts that split the section into six or seven subsections, so I instructed the chatbot to keep most sections to a maximum of two large subsections. Next, I keyed in on specific examples from my knowledge that I wanted the chatbot to focus on in the early stages of AI development, since ChatGPT was giving around a sentence or two for each example it gave: I told ChatGPT to focus on the development of the Logic Theorist and the General Problem Solver, on top of its writing about Turing. After this, I was satisfied with the end result for the section, so I requested a bibliography for the section according to the Chicago Manual of Style, and the chatbot gave one that looked solid when I reviewed the guidelines, using mostly sources from historical literature rather than articles written in recent times.

I then moved onto the next section, evolution of machine learning. After a first draft, I noticed that ChatGPT was writing a significant portion about the relationship between machine learning and two concepts later in the paper, deep learning and natural language processing. I added instructions to the chatbot, telling it to have the paper flow in a generally chronological order, instead of skipping around mid-paper between different periods in the development of AI. However, since I gave a general timeline for the section to be from the 1950s to the 1980s, the chatbot's next draft was very date-heavy, mentioning those years constantly, perhaps in an effort to make the chronology more obvious. I told ChatGPT to be "more subtle" with this chronological timeline and re-emphasized my earlier instructions to provide specific, concrete historical examples. Additionally, I gave instructions to give a longer conceptual description of machine learning, and asked ChatGPT to connect it to the development of AI overall, as I thought the section seemed too much like a short summary of the history of machine learning, rather than part of the larger paper. This was an issue I encountered multiple times over the course of my experience constructing the paper, as ChatGPT seemed to sometimes forget instructions I gave for previous sections that I intended to apply to the paper as a whole, so in later sections I restated the instructions at the beginning when I initiated the first draft. After getting some more specific examples from the development of machine learning and their relevance to the development of AI, I moved to the next section.

For the section on advancements in natural language processing, I had an easier time because of the strategies I learned from previous sections. However, my instructions on continuing the flow of the paper (rather than simply summarizing each aspect in a disconnected manner) backfired a bit, as the first draft mentioned Turing's contributions, which clearly was already stated in the first main section of the paper. After instructing ChatGPT to fix this issue by having a clear focus on the topic of the section, I also added an instruction that made the process significantly less difficult: I told ChatGPT to focus about 50 to 60 percent of each section on the historical evolution of each relevant topic, and 40 to 50 percent on its importance to the development of AI. This resulted each subsequent section (and the previous ones after revision) to have a consistent tone and flow throughout the paper. The next few sections to be much easier to process, with my strategy of restating the instructions and adjusting based on the results being extremely effective. The deep learning revolution and LLM/chatbot sections both came out very good on first draft, and I was satisfied with the results (although I still planned on adding material of my own to bolster the AI's writing.

For the conclusion, I had different instructions, as I wanted the paper to wrap up in a very thoughtful and creative manner. I requested ChatGPT to focus on the future prospects and ethical challenges posed by AI, the relationship between AI's growth and the potential decline of human identity, and ways that AI could develop (in some potentially comedic and humorous trajectories), leading to interesting results. The first result was sloppy, as ChatGPT used the word "humorous" far too often, which makes sense since humor can be difficult for machines to fully understand and replicate. However, I refocused the chatbot's conclusion on specific developments that could benefit humanity, like healthcare, and also asked ChatGPT to keep the conclusion more grounded in wrapping up the initial prompt. I was satisfied with the end result.

Finally, before adding my own material and research, I put the whole drafted essay into ChatGPT to identify repetitive writing patterns (such as those that make papers seem "AI-generated" which can lead to reader boredom), and to also examine areas of the paper that seemed thinner on content, where I could expand with further research to synergize the collaboration. ChatGPT specifically pointed out phrases like multifaceted development, transformative power, ethical considerations, and advancement as being repetitive throughout the paper, which was unsurprising to me because I went section by section, which could lead to similar patterns of writing throughout the combined paper. I then split the paper up into 500-word sections and asked ChatGPT to write with more of an active voice, since I have a habit of writing in the passive voice too often (which can weaken a paper), and I figured that the combined efforts would sound professional yet natural.

Fact-checking the paper and adding my own elements was an easier process than I expected. All the incorrect or ahistorical references in the examples given in ChatGPT's writing were easy to spot simply by reading through the paper and researching each example. There were some developments incorrectly attributed to different people, and some insignificant developments that seemed propped up more than necessary in the paper, so I deleted some and corrected others. I added significant portions of my own writing once I researched most topics in order to make sure that the important historical landmarks were hit. The best writing that the AI did was on deep learning and large language model chatbots, which I had the hardest time understanding conceptually. My guess is that due to the recency of these developments, as well as their extreme relevance to ChatGPT and other chatbots, the AI was well-prepared to describe the concepts in detail, so I added the least to these sections (with the exception of deleting some bloated examples). After adding these elements to each section based on my own research on examples and topics ChatGPT already mentioned and doing a final fact-check of the information in each section, I was satisfied with the final product.

AI Research Paper: Final Product

The evolution of AI from its conceptual origins with Alan Turing to the advent of large language model chatbots represents a journey marked by significant milestones and technological advancements. This paper examines the multifaceted development of AI, tracing its progression through key stages of innovation and breakthroughs. As a unique collaboration between myself and a large language model chatbot, ChatGPT 3.5, the paper represents an insight into the many ways that AI has been shaped by individuals and historical events, the ways it already is shaping the world we live in, and how it may completely change the idea of life as we know it. First, we will delve into the early foundations of AI, examining Turing's groundbreaking contributions and the emergence of early AI programs focused on symbolic reasoning and logical inference. We then explore the evolution of machine learning, from neural networks to decision trees, and their impact on AI's capabilities. Next, we shift our focus to advancements in natural language processing (NLP), discussing the development of NLP techniques and their applications in systems like IBM's Watson. The paper then delves into the deep learning revolution, reviewing the transformative power of deep neural networks in computer vision, speech recognition, and NLP tasks: eventually leading into the development of large language models. We analyze the capabilities of these models in text generation, translation, summarization, and conversation, while also addressing the vast array of ethical considerations. Finally, we consider the future in AI, including the potential for explainable AI, ethical AI development, and the integration of AI in diverse sectors such as healthcare and finance. At the end of this paper, my hope is that whoever reads this understands how AI became what it is today, and perhaps have a better capability of collaborating and controlling it for the betterment of humanity.

Alan Turing, a pioneering mathematician and computer scientist, played a pivotal role in shaping the early foundations of artificial intelligence. Turing's groundbreaking contributions during World War II, where he worked on breaking the German Enigma code using early computing machines, highlighted the potential for machines to perform complex tasks previously deemed exclusive to human intelligence. His seminal paper from 1950, "Computing Machinery and Intelligence", introduced the Turing Test as a benchmark for evaluating machine intelligence, sparking extensive research and debate within the AI community. Further activity on the AI front occurred in the late 1950s, where significant developments began with the creation of early AI programs like the Logic Theorist. Developed by Allen Newell, Herbert A. Simon, and J.C. Shaw in 1956, the Logic Theorist demonstrated AI's potential for logical reasoning tasks. It utilized symbolic logic and heuristic search algorithms to prove mathematical theorems, showcasing early AI's ability to mimic human reasoning processes. The Logic Theorist's groundbreaking approach marked a shift in AI research towards problem-solving using computational methods. By representing problems as logical statements and employing search algorithms, the Logic Theorist achieved notable success in theorem proving tasks, laying the groundwork for future AI systems' logical reasoning capabilities. One year later, Newell, Simon, and Shaw developed the General Problem Solver (GPS). GPS was designed to tackle a wide range of problems by representing them as states, operators, and goals, showcasing AI's ability to solve complex problems through computational methods. GPS utilized heuristics and problem-solving strategies, laying the groundwork for future AI systems' problem-solving capabilities. The most collaborative aspect of early AI research and development was the Dartmouth workshop in 1956, which combined minds from across the world to build a foundation for how to approach the issue of artificial intelligence. This uniquely human activity of putting together various perspectives who shared a common goal (avoiding Luddism and focusing on forging our future by any means necessary) would become a common theme throughout the development of AI. The Dartmouth workshop, Logic Theorist, and GPS, as well as Turing's early contributions, laid the groundwork for all other aspects of AI development throughout the rest of the century and beyond. They demonstrated AI's potential for logical reasoning, problem-solving, and symbolic manipulation, setting the stage for the rapid evolution of AI capabilities in the decades to come. These foundational contributions paved the way for advancements in machine learning, natural language processing, and intelligent systems, shaping the modern landscape of artificial intelligence.

The next foundational development in the field of AI was machine learning. Unlike traditional programming paradigms that rely on explicit instructions, machine learning enables computational systems to learn from data, identify patterns, and make decisions autonomously. Arthur Samuel is generally credited with coming up with the concept of machine learning, as he developed a checkers program while working at IBM in the 1950s. His program used a search algorithm known today as alpha-beta pruning (due to computer limitations) to learn the possible winning positions from a given game position. Hypothetically, this program could learn and improve over time to become better at the game: in later years, similar programs like Deep Blue and AlphaGo were even able to surpass humans in games like chess and Go. Although seemingly fixated on games, the ability for a computer to learn at a game is applicable to complex problems in a variety of ways. Neural networks, inspired by biological neurons, form the foundation of machine learning. Early models like the perceptron, introduced by Frank Rosenblatt in 1957, laid the groundwork for neural network architectures. Through a process called training, neural networks learn to recognize patterns and relationships in data, making them highly effective for tasks such as image recognition, natural language processing, and predictive analytics. The These networks gained renewed interest in the 1980s with the introduction of backpropagation algorithms by Rumelhart, Hinton, and Williams, which allowed neural networks to adjust their internal parameters and optimize performance based on feedback. Backpropagation allowed efficient learning and adjustment of internal parameters, overcoming limitations of earlier models and enabling neural networks to tackle more complex tasks. The next evolution of machine learning was Ross Quinlan's development of decision tree algorithms. In 1986, Quinlan developed the ID3 (Iterative Dichotomiser 3) algorithm, providing a transparent and interpretable framework for decision-making. By recursively partitioning data based on features, decision trees generate simple yet effective rules for classification and regression tasks. This interpretability makes decision trees valuable in domains where understanding the reasoning behind predictions is essential, such as healthcare diagnostics and financial risk assessment. Support vector machines (SVMs), primarily developed by Vladimir Vapnik, are another powerful machine learning technique that excels in classifying data points into distinct categories. SVMs work by finding the optimal hyperplane that separates different classes in a high-dimensional space, making them particularly effective in scenarios with complex and non-linear relationships. SVMs have been successfully applied in diverse fields, including image recognition, bioinformatics, and sentiment analysis. Machine learning techniques have transformed AI capabilities and applications. By harnessing the power of machine learning, AI systems can learn from vast amounts of data, adapt to changing environments, and make intelligent decisions in real time.

The next development in the evolution of AI is Natural Language Processing (NLP), which enables machines to understand, interpret, and generate human language. The development of symbolic NLP, which required rule-based systems, began with the Georgetown experiment with IBM in 1954, which showed the capability of a computer to automatically translate Russian sentences to English. However, these early systems were limited in handling complex language structures and nuances. As years passed, research shifted towards statistical NLP, which did not require the long rule-based systems of symbolic NLP, primarily due to the developments of machine learning. Researchers explored techniques like Hidden Markov Models (HMMs) for speech recognition and syntactic parsing, enhancing accuracy and robustness in language processing tasks. The development of probabilistic context-free grammars (PCFGs) by Eugene Charniak and colleagues further improved parsing accuracy, laying the groundwork for more sophisticated syntactic analysis (Charniak et al., 1993). Advancements in machine learning algorithms, such as Conditional Random Fields (CRFs) introduced by Michael Collins and colleagues, revolutionized sequence labeling tasks like part-of-speech tagging and named entity recognition. These statistical models incorporated context dependencies and achieved state-of-the-art performance in various NLP applications (Collins et al., 2002). The turn of the millennium saw a surge in research on sentiment analysis, opinion mining, and text summarization. Researchers like Bo Pang and Lillian Lee pioneered sentiment analysis methodologies using machine learning algorithms, enabling the classification of textual data based on sentiment polarity (Pang & Lee, 2008). Techniques like Latent Dirichlet Allocation (LDA) for topic modeling and word embeddings such as Word2Vec and GloVe gained prominence, allowing machines to capture semantic relationships and context in textual data. Importantly, the historical development of NLP laid the foundation for deep learning techniques in the field. While traditional NLP approaches focused on rule-based systems and statistical models, deep learning models like recurrent neural networks (RNNs) and transformers revolutionized language understanding and generation. The introduction of the Transformer architecture, popularized by the BERT (Bidirectional Encoder Representations from Transformers) model, achieved state-of-the-art performance in tasks like language translation, sentiment analysis, and question answering, showcasing the power of deep learning in NLP (Devlin et al., 2018). Statistical NLP is vitally important to the evolution of AI from Turing to chatbots, as it overarchingly gave computers the ability to generate, understand, and process human language, which was a huge improvement from basic machine learning and symbolic NLP. NLP facilitates human-machine communication, information retrieval, sentiment analysis, and knowledge extraction from textual data. This continuous evolution of NLP techniques and their integration with deep learning models drives innovation in AI applications and propels advancements in language processing capabilities.

The Deep Learning Revolution represents a transformative phase in artificial intelligence, building upon the foundations laid by Natural Language Processing (NLP) and statistical modeling techniques. Deep learning, a subset of machine learning, involves training deep neural networks with multiple layers to learn hierarchical representations of data, enabling machines to perform complex tasks and make decisions based on learned patterns and features. One of the defining characteristics of deep learning is its ability to automatically discover and extract features from raw data, eliminating the need for manual feature engineering. This capability is particularly advantageous in handling unstructured data, such as images, text, and audio, where traditional algorithms often struggle to capture relevant information. The deep learning revolution began to gain momentum in the 2010s, driven by advancements in computational resources, algorithmic innovations, and the availability of large-scale datasets. Notable historical figures such as Yann LeCun, Geoff Hinton, and Yoshua Bengio were instrumental in pioneering key concepts and architectures in deep learning. Yann LeCun's work on Convolutional Neural Networks (CNNs) revolutionized image recognition tasks, enabling hierarchical feature extraction and pattern recognition (LeCun et al., 1998). Geoff Hinton's contributions to backpropagation algorithms and deep belief networks laid the foundation for training deep neural networks efficiently (Hinton et al., 2006). Yoshua Bengio's research on unsupervised learning and neural network architectures contributed significantly to advancing deep learning methodologies (Bengio et al., 2007). In addition to CNNs, Recurrent Neural Networks (RNNs) and their variants, such as Long Short-Term Memory (LSTM) networks, revolutionized sequence modeling and natural language processing. Alex Graves and colleagues further refined LSTM networks, making them effective for tasks like language modeling and speech recognition (Graves et al., 2009). The importance of deep learning for AI development lies in its capability to learn complex representations of data, model intricate relationships, and generalize patterns across diverse domains. Unlike traditional machine learning models that rely on handcrafted features and shallow architectures, deep neural networks can automatically extract hierarchical features from raw input data, leading to superior performance in tasks like image classification, speech synthesis, and recommendation systems. Furthermore, deep learning techniques have democratized AI research and applications, making sophisticated algorithms and models accessible to a broader community of researchers, developers, and practitioners. Open-source libraries and frameworks like TensorFlow, PyTorch, and Keras have contributed to the widespread adoption of deep learning across various industries and domains. The integration of deep learning with NLP has led to significant advancements in language understanding, generation, and dialogue systems. Models like the Transformer architecture, exemplified by the BERT (Bidirectional Encoder Representations from Transformers) model, have achieved remarkable performance in tasks such as language translation, sentiment analysis, and question answering, pushing the boundaries of AI capabilities in linguistic tasks. Ethical considerations in deep learning, including fairness, transparency, and bias mitigation, are critical aspects that require ongoing research and development efforts. Addressing these challenges is essential to ensuring responsible AI deployment and fostering trust in AI systems among users and stakeholders. The deep learning revolution represents a paradigm shift in AI, empowering machines to learn complex representations from data and perform tasks with human-like capabilities. By advancing deep learning techniques, AI systems become more versatile, adaptable, and capable of tackling real-world challenges across diverse domains.

The emergence of Large Language Models (LLMs) built upon the foundations laid by deep learning and natural language processing (NLP). Large Language Models (LLMs) are deep learning-based models trained on massive amounts of text data to understand and generate human-like language. The concept gained prominence with models such as OpenAI's GPT (Generative Pre-trained Transformer) series, including GPT-2 and GPT-3, which demonstrated remarkable capabilities in text generation, conversation, and language understanding (Radford et al., 2019). These models leverage transformer architectures and pre-training techniques to learn rich contextual representations of language, enabling them to perform tasks ranging from translation and summarization to dialogue generation and question answering. The development of LLMs was propelled by advances in computing power, data availability, and algorithmic innovations. Ilya Sutskever at OpenAI, as well as Samy Bengio and Quoc Le at Google, played pivotal roles in advancing transformer architectures and training methodologies, contributing to the scalability and efficiency of large-scale language models (Sutskever et al., 2014). Chatbots, on the other hand, are AI systems designed to engage in natural language conversations with users, simulating human-like interaction. Early chatbot technologies date back to the 1960s with programs like ELIZA, which utilized pattern matching and scripted responses to simulate conversation (Weizenbaum, 1966). However, modern chatbots leverage deep learning, NLP, and LLMs to achieve higher levels of conversational sophistication and contextual understanding. The importance of LLMs and chatbots in AI development lies in their ability to bridge the gap between AI systems and human users, enabling seamless communication, information retrieval, and task automation. LLMs enhance language understanding and generation capabilities, enabling chatbots to comprehend complex queries, provide relevant responses, and engage in meaningful dialogue across diverse domains. Moreover, LLM-powered chatbots find applications in customer service, virtual assistants, education, healthcare, and entertainment. Companies like Google, Microsoft, and Facebook have integrated chatbot technology into their platforms, offering users personalized experiences, real-time assistance, and interactive interfaces. Large Language Models (LLMs) and chatbots represent a significant milestone in AI evolution, leveraging deep learning, NLP, and transformative architectures to enable human-like language understanding and communication. The ongoing advancements in LLM technologies and chatbot capabilities are reshaping the way humans interact with AI systems, paving the way for more intelligent, conversational, and empathetic AI experiences.

Looking ahead, the evolution of large language model chatbots opens intriguing possibilities and ethical dilemmas. These advanced chatbots, powered by deep learning and natural language understanding, could redefine how we interact with AI and each other. In the realm of creativity, imagine chatbots collaborating with writers and artists to co-author novels or create visual art. These collaborations could blend human creativity with AI's vast knowledge and computational abilities, resulting in unique and innovative works that push the boundaries of traditional art forms. In healthcare, large language model chatbots could revolutionize patient care by providing personalized medical advice, facilitating telemedicine consultations, and assisting in medical research. However, ensuring patient privacy, data security, and ethical use of AI in healthcare remain critical challenges that require careful consideration and regulation. Looking even further ahead, AI's development could lead to AI-augmented humor, where chatbots generate jokes and memes based on vast datasets of humor styles and cultural references. While this adds a lighthearted element, it also raises questions about AI's ability to understand and navigate nuanced aspects of human communication, including humor and emotions. Ethically, as AI becomes more sophisticated, the need for transparency, accountability, and ethical guidelines becomes paramount. Developments in explainable AI (XAI) can help users understand how AI chatbots make decisions and ensure that these systems operate ethically and responsibly. Further concerns arise considering the ability of AI to surpass humans on an intellectual level. From Deep Blue defeating Garry Kasparov in chess in 1997, to AlphaGo defeating Lee Sedol in Go in 2016, AI has continually proved itself capable, with time, in various ways (through sheer computational force or neural networks), to grow beyond humans, which at first glance seems dangerous. However, we should endeavor to have control over AI, rather than limiting it or disavowing its use. AI is a useful tool, as seen in the process of drafting this paper. In conclusion, the future of large language model chatbots holds immense potential for creativity, innovation, and improved human-machine interactions. By addressing ethical challenges, fostering collaboration between humans and AI, and maintaining a human-centered approach, we can harness the full benefits of AI while mitigating risks and ensuring a positive impact on society.

Bibliography

Bengio, Yoshua, et al. "Greedy Layer-Wise Training of Deep Networks." Advances in Neural Information Processing Systems 19, 2007, pp. 153-160.

Charniak, Eugene, et al. "Statistical Parsing with a Context-Free Grammar and Word Statistics." Proceedings of the Fourteenth National Conference on Artificial Intelligence (AAAI-97), 1993, pp. 598-603.

Collins, Michael, et al. "Three Generative, Lexicalized Models for Statistical Parsing." Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics, 2002, pp. 16-23.

Graves, Alex, et al. "A Novel Connectionist System for Unconstrained Handwriting Recognition." IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 31, no. 5, 2009, pp. 855-868.

LeCun, Yann, et al. "Handwritten Digit Recognition with a Back-Propagation Network." Proceedings of the 1989 IEEE Conference on Neural Information Processing Systems (NIPS), 1998, pp. 396-404.

Newell, Allen, Herbert A. Simon, and J.C. Shaw. "The Logic Theorist: A Case Study in Heuristic Problem Solving." Proceedings of the Western Joint Computer Conference, 1956, pp. 230-240.

Pang, Bo, and Lillian Lee. "A Sentimental Education: Sentiment Analysis Using Subjectivity Summarization Based on Minimum Cuts." Proceedings of the Association for Computational Linguistics, 2008, pp. 271-278.

Radford, Alec, et al. "Language Models are Unsupervised Multitask Learners." OpenAI Blog, 2019.

Sutskever, Ilya, et al. "Sequence to Sequence Learning with Neural Networks." Advances in Neural Information Processing Systems, 2014, pp. 3104-3112.