

Colin Norton
04/12/2024
History of the Digital Revolution
Professor Isaacson

The AI Odyssey: From Speculation to Sophistication

PREFACE

As we have been talking about this project throughout the class, several students have mentioned different chatbots and how they have been using them in tandem in order to create a final version of their essay. This concept intrigued me because while I do use AI in my personal time, I have very little understanding of the differences between each model. In order to change this I have decided to take each of the initial sections of this essay, divide them up, and work with one specific AI chatbot for each to highlight the differences between them as well as their strengths and weaknesses through the lens of this project. To accomplish this I started by having a conversation with the LLM I am most accustomed to and establishing an outline for how the paper should look, that outline is depicted below.

Introduction

- Brief overview of AI and its significance in the modern world.
- Introduction to Alan Turing and his foundational role in the development of AI.
- Thesis statement highlighting the evolution from Turing's initial concepts to today's advanced AI chatbots.

The Foundations of AI

- Alan Turing and the Conceptual Framework
 - Discussion of Turing's 1950 paper, "Computing Machinery and Intelligence."
 - Introduction of the Turing Test as a criterion for machine intelligence.
- Early Experiments and Theoretical Developments
 - Brief overview of the first AI programs, such as the Logic Theorist and ELIZA.
 - The establishment of AI as a formal discipline in the 1950s and 1960s.

Key Milestones in AI Development

- The AI Winter(s)
 - Explanation of the periods known as AI winters, when progress slowed due to limited computational power and skepticism.
- Resurgence of AI Research
 - The role of increased computational power and big data.
 - Breakthroughs in machine learning and neural networks.

The Era of Machine Learning and Neural Networks

- From Symbolic AI to Machine Learning
 - Transition from rule-based systems to algorithms that learn from data.
- Deep Learning Revolution
 - The impact of deep learning on AI's capabilities, particularly in perception tasks like image and speech recognition.
- Notable AI Systems and Milestones
 - IBM's Deep Blue and Watson.
 - Google's AlphaGo and advancements in autonomous vehicles.

The Rise of Large Language Models

- Introduction to Language Models and NLP
 - Explanation of natural language processing (NLP) and the role of language models in interpreting and generating human language.

- Evolution of Language Models
 - From early statistical models to transformer-based architectures like GPT and BERT.
- Large Language Model Chatbots
 - Overview of the development and capabilities of large language models like GPT-3.
 - Discussion of their applications, strengths, and limitations.

Conclusion

- Recapitulate the major points made throughout the essay.
- Reflect on Turing's legacy and the current state of AI chatbots.
- Final thoughts on the future of AI and its potential impact on society.

I took each section of this outline and fed it into different popular Chatbots that I heard my classmates using, I then went on to ask the chatbots if they felt that the response was missing anything as well as how to refine the piece and implement those changes. At this stage I took note of how the AI had done and then went on to refine the work in terms of tone and style as well as information that was left out. The changes I made to the final AI response or information that I specifically had to ask about or add myself has been highlighted in green. Alongside each paragraph I have put my notes on how that specific chatbot worked in italics.

One AI that I have intentionally left out of this project is NotebookLM. I have a personal affinity for philosophy and have always enjoyed the works of writers who have unfortunately been dead for centuries. Along this line I have always been envious of the people who have gotten to study directly under such minds (think alexander experiencing the tutelage of aristotle). So when the speaker from google came and highlighted how you could compose your own set of data into NotebookLM to pull from I was thrilled, thinking that I could potentially replicate this experience by training it on the works of great thinkers and then engaging in a conversation with it. I started by uploading massive works from Hellenic philosophers like Plato and Seneca through existentialist pieces by Camus and Sarte (sourced from pdfs of public domain works taken off Google), however instead of the frankensteinian professor I was hoping to create, I ended up with a glorified fetch tool. That is not to say that Notebook could not do the simple tasks I asked of it or interpret modern questions through the lens it's database. But it completely lacked the soul and human aspect of what made those writings great in the first place. It would pull from the most surface level interpretations of text without really getting into any nuances and most of its responses were bland and unpalatable. This to me is a key example of what these AI tools do in general, while they are capable of amazing feats imitations (as will be shown throughout the essay) they lack a certain substance. They are flashy but remain as empty calories.

To demonstrate this I thought it would be interesting to have Dalle - 4 create a piece of art to accompany each paragraph. I copied the finished version of each paragraph into ChatGPT 4 and asked it to summarize the main themes and ideas of the paragraph. I then asked it to create an image based on those themes in the style of different artists that have unique styles and feels to see what would come up. Many of these images took several different prompts to a passable result as they would often over characterize themselves or miss the mark entirely. Another issue

that came up with this process is that the AI would occasionally deny my requests as they “violated policy guidelines”. When I asked GPT 4 why this was it told me that in order to be safe, it could not directly replicate any artists past 1912 (additionally any prompt containing the name Ada Lovelace was blocked as well). Many of the artists I wanted to imitate created work in the last century, so to get around this I asked the chatbot how to circumvent it’s own policy guidelines. Which it happily told me how to do. It isn’t that Dalle-4 can’t create art in the style of artists after 1912, it is just that it can’t do so when directly asked to avoid potential copyright infringements. To get around this I would give several specifics about an artist I wanted replicated without actually naming them to let the AI fill in the blanks and pull from that artists data to create it’s image. I also asked it to provide an analysis of each piece which I have left in an appendix at the end of the essay.

INTRODUCTION

(ChatGPT 3.5) (Image Prompted in the Style of Leonardo Da Vinci)

ChatGPT 3.5 is the chatbot I have the most experience with. For this reason I had it establish the outline for my paper as well as the introduction. In the toolbox of LLMs it is the hammer. ChatGPT 3.5 is extremely versatile and surprisingly creative when it comes to creating a response. However it has a specific tone that despite many prompts to change always feels a bit off. While it does a good job creating a passable response I have always felt that in order for it to become palatable many of the phrases have to be reworded and the ideas require some rearranging to have a more natural flow.



Artificial Intelligence (AI) has transitioned from a **facet of science fiction** to a core component of technological advancement, influencing various sectors of society. **Alan Turing** stands at the heart of AI's conceptual beginnings as the mathematician who laid the groundwork

for modern computing and AI. Turing's question, "Can machines think?", initiated a quest that has spanned decades, leading us to the development of sophisticated AI systems capable of understanding and generating human language. This essay traces the evolution of artificial intelligence from its inception with Turing's theoretical contributions to the contemporary era of large language model chatbots. To offer a practical perspective on AI's capabilities, this analysis is uniquely informed by AI tools used in the research, drafting, and refinement processes. **By employing AI in crafting this essay, we** aim to provide insights into how AI technologies can serve as valuable aids in academic and creative endeavors. The journey from Turing's theoretical models to today's chatbots illustrates the rapid advancements in AI technology and its growing impact on society. Through this exploration, we'll examine key milestones, technological breakthroughs, and the challenges faced in AI's ongoing development.

BODY 1: Origins of AI

(Chat GPT 4) (Image prompted in the Style of Van Gogh)

I found ChatGPT 4 extremely easy to work with and much more responsive than 3.5 or any other chatbot. What was interesting for this paragraph was how rigidly it stuck to each prompt given in the initial outline. Because the outline mentions Turing specifically by name, the AI in its first rendition of this body paragraph exclusively talked about him and made little to no mention of other scholars or people of influence in the origins of AI. However, once I prompted it to take some other figures into account such as Lovelace and early influential thinkers, the piece it provided a very different opinion about the origins of AI that expressed credit in multiple different directions. What this brought up for me was the potentially tricky nature of AI and how it is specifically tailored to give you what you ask for, almost similar to stories of a wish twisted by a genie. This example was apparent to me because I was primed to think of multiple other influential thinkers when it came to the development of AI, but i'm sure even in this paper alone there are several biases that have been reflected through my prompts about the nature of AI as well as the ideas mentioned. I decided to ask ChatGPT 4 about this and it's response was rather interesting. It's responses detailed how it scans the prompts given for both tone and information to give a response that matches what it is fed. The example it provided was "if you asked 'why are teenagers so irresponsible with money' vs. 'Can you discuss some factors that might affect teenagers' money management skills?'" both questions are looking for the same answer but the tone and potential jumping off points for a response are quite different. In this example the former question leads to a response focused on the teenagers shortcomings of inexperience and an underdeveloped brain, where as the latter prompt discusses systemic factors like family and education that would impact choices made. One issue is that similar to its predecessor GPT 3.5, most of the responses follow an arc of their own essay with a clear beginning middle and end rather than the flow a body paragraph in a human essay would have. This reflects the Chatbots desire to create closed loops when generating a response. Much like John Searls 'man in the chinese room' because the AI has no understanding of what it's saying and only what would likely be said in that scenario, it has a tendency to neatly wrap up each of its responses as if it

were closing a melody. Similar to ChatGPT 3.5 despite prompting I could not get it to cough up its sources.



The intellectual bedrock of artificial intelligence finds its origins not just in the 20th century but in the fabric of human thought stretching back centuries. The speculation that machines or mechanical devices could mimic human intelligence predates Alan Turing by a significant margin. Ancient myths and legends, along with philosophical inquiries by figures such as René Descartes and Gottfried Wilhelm Leibniz, pondered the mechanical nature of thought. Leibniz, in particular, envisioned a universal language and a computational system in the 17th century that could represent and solve all reasoning tasks, an idea that prefigures modern computing principles. Moreover, the 19th century saw Ada Lovelace, working alongside Charles Babbage on the Analytical Engine, assert the potential for machines to perform tasks beyond mere calculation, anticipating the concept of programming. Lovelace is often celebrated as the first computer programmer and for her foresight in recognizing that computers could one day "compose elaborate and scientific pieces of music of any degree of complexity or extent."

Against this rich historical backdrop, Alan Turing's contributions in the mid-20th century crystallized the discourse around artificial intelligence. His seminal 1950 paper, "Computing Machinery and Intelligence," and the introduction of the Turing Test provided a clear, practical framework for understanding and developing machine intelligence. Turing's work effectively bridged speculative inquiries with concrete mathematical and computational principles, setting a rigorous scientific inquiry into AI. The Turing Test, involving a human judge engaging in conversation with an unseen interlocutor to assess machine intelligence, set a benchmark for the field.

Parallel to Turing's theoretical groundwork, the 1950s and 1960s witnessed practical experiments aimed at realizing the concepts of intelligent machinery. The Logic Theorist, developed by Allen Newell, Herbert A. Simon, and Cliff Shaw, demonstrated machines could engage in complex problem-solving, while Joseph Weizenbaum's ELIZA laid the groundwork for natural language processing. These efforts, together with the formal establishment of AI as a discipline at the Dartmouth Conference in 1956, marked the beginning of an ambitious journey to endow machines with intelligent behavior, despite early computational limitations and skepticism.

These foundational efforts underscored the complexity of mimicking human intelligence and set the stage for the advanced algorithms and systems that define today's AI field. Turing's legacy, in particular, remains a touchstone in discussions on machine intelligence, embodying the transformative vision that has propelled the field forward.

BODY 2: Key milestones in AI (Gemini) (Image Prompted in the Style of Pablo Picasso)

Gemini required much more prompting for usable responses whereas ChatGPT would often give me too much for a response, I felt Gemini would often give very little. The responses of this chatbot were also much less formal using casual language even when prompted. Most attempts to flush out this paragraph and add information resulted in diluting the quality even further or expanding only on tangential points like anecdotes about the difficulty of beating a chess grandmaster rather than the actual meat of the essay. What was exciting and different about using Gemini AI, however, was the ability to directly see which sources pieces of information came from. This made it significantly easier to track down information and sources, however the sources that it provided were typically massive public forums of databases like wikipedia or at one point a hacker news forum. Whether or not this is the only data that Gemini trains on is uncertain but from the consumer's perspective it's all that I had the ability to see. I personally use Gemini for certain classes when I need a topic explained to me because of these features, Gemini remains concise and is straightforward and reliable when asked about complex topics such as advanced accounting or biology. However through this experience I can see that those strengths become weaknesses when attempting to write a paper. The green text in this paragraph show areas where I had to add in information or tell Gemini to flush out an Idea to better express what I wanted out of this paper.



Despite the buzz around the development of Artificial Intelligence progress on this front has been cyclical. Early optimism in the mid-20th century, fueled by achievements like the perceptron, an early neural network model, and LISP, a powerful programming language for symbolic manipulation, fueled this initial optimism. However, these collided with limitations – computers at the time lacked the necessary processing power – and a growing understanding of the intricacies of human intelligence. This resulted in the first AI Winter of the 1970s, characterized by decreased funding and research stagnation. SHRDLU, a ground-breaking natural language program developed in the 1960s, exemplified these limitations. SHRDLU could respond to user commands within a simulated world of blocks and robots. While impressive for its time, it relied heavily on pre-programmed knowledge and struggled with anything outside its specific domain. For instance, it couldn't grasp relative concepts or handle unexpected language structures. SHRDLU's limitations, alongside the growing recognition of human intelligence's complexity, underscored the challenges of achieving true AI and contributed to the funding shortfalls that characterized the AI Winter of the 1970s. A similar period followed in the 1980s with the disappointment surrounding expert systems. The late 20th and early 21st centuries witnessed a significant resurgence in AI research. Fueling the rebirth of AI research was the confluence of two key factors: exponentially increasing computing power predicted by Moore's Law and the vast availability of big data. The former provided the computational muscle to train complex AI algorithms, while the latter offered a rich training ground for these algorithms to learn from and refine their abilities. This synergy between processing power and data availability has been instrumental in the recent breakthroughs in AI. Breakthroughs in machine learning, particularly deep learning techniques like convolutional neural networks, allowed AI to excel in areas like image recognition and natural language processing. These advancements were

exemplified by milestones like IBM's Deep Blue defeating chess grandmaster Garry Kasparov (1997) and Watson's triumph on Jeopardy! (2011). The AI Winter had finally come to an end, paving the way for a new era of innovation and real-world applications.

Sources provided: Unable to cite in-text as it “violates safety guidelines”

<https://hci.stanford.edu/winograd/shrdlu/AITR-235.pdf>

https://en.wikipedia.org/wiki/AI_winter

<https://www.turingpost.com/p/aiwinters>

BODY 3: The Era of Machine Learning Networks

(Claude AI) (Image prompted to match the style of Jackson Pollock)

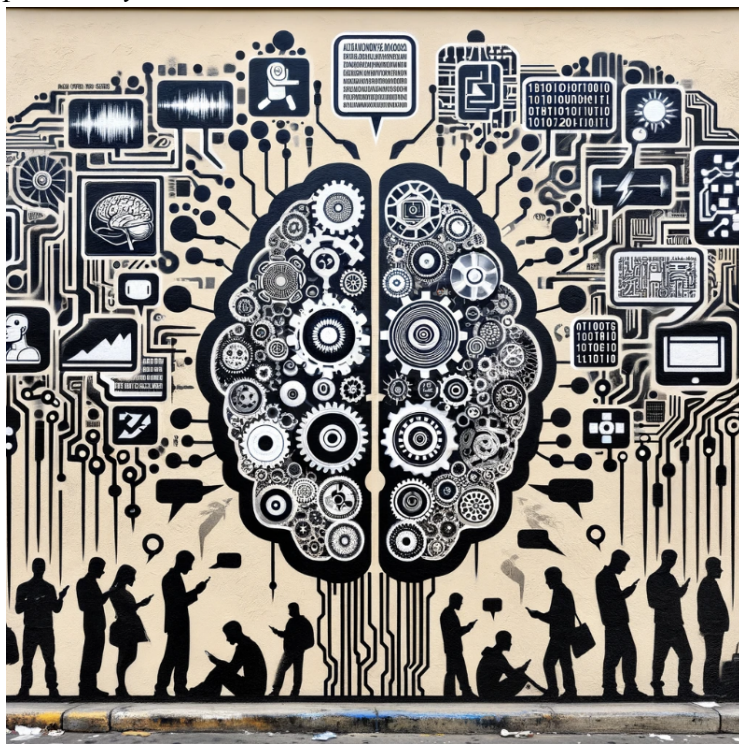
Claude was an interesting AI to work with because it felt much more casual than other language models. In my opinion Claude has a big leg up on Gemini and GPT 3.5 as it has a much more personable tone and can pull specific sources and quotes which is very helpful for writing a research paper such as this. One thing that I have found difficult with Claude is the large wait time that comes after too many messages were written, I have often been cut off in the middle of testing out a conversation till the next day. Claude also tended to stay more surface level than either ChatGPT, I had to specifically prod it to produce more information of what seemed to me to be key portions of this section (deep learning techniques and what was going on behind the scenes) however the majority of what Claude put out after the refinement of telling it to go more in depth on those sections and asking it to take out certain tone choices like being too casual seemed impressive. Unlike other AIs I felt that Claude was capable of interpreting my prompt beyond just the literal meaning of the words I gave the suggestions I had for it ranged from “less formal” to paragraph requests on information I wanted added in and sourced and in both cases it was able to exceed the expectations that both GPT and Gemini had set for me.



The field of artificial intelligence has undergone a transformative paradigm shift from the earlier symbolic approach of manually programming rule-based systems to the modern data-driven machine learning paradigm. This transition was catalyzed by the profound breakthrough known as the "deep learning revolution" in the late 2000s, which was enabled by the availability of massive datasets and significant increases in computing power from specialized hardware like graphics processing units (GPUs) (Sejnowski, 2018; Goodfellow et al., 2016). Deep learning techniques, inspired by the hierarchical layered structure of the brain's neural networks, enabled artificial neural networks to automatically discover intricate patterns in data in a way that superseded previous machine learning approaches. Deep neural networks stack multiple non-linear transformation layers, allowing them to learn increasingly abstract representations from the raw data. This hierarchical approach mimics how the brain processes information across neuronal layers, with each layer capturing more complex representations built up from the previous layer. Unlike traditional machine learning models limited to a single transformation, deep neural networks can automatically learn rich, distributed representations that encode highly intricate patterns and high-level concepts directly from the training data. This allowed deep learning models to surpass shallow architectures that relied on manually engineered features. By learning their own structured multilayered representations, deep neural networks could distill the complex statistical patterns required for tasks like image and speech recognition that were previously bottlenecks for AI. (LeCun et al., 2015). Specific innovations like convolutional neural networks for image processing and recurrent neural networks for sequence data were foundational (Schmidhuber, 2015). This capability powered major advances in areas like image and speech recognition that were previously bottlenecks for AI progress. Prominent milestones showcased deep learning's prowess, such as IBM's Deep Blue becoming the first computer to defeat a reigning world chess champion, Garry Kasparov, in 1997 after years of previous failures (Campbell et al., 2002). Over a decade later, IBM's Watson system demonstrated deep learning's mastery of natural language by successfully taking on human Jeopardy! quiz show champions in 2011, a feat requiring nuanced understanding of puns, metaphors, and other complexities of human communication (Ferrucci et al., 2010). In 2016, Google's AlphaGo program achieved one of the long-standing "grand challenges" of AI by defeating the world's best human players at the ancient game of Go, a highly intuitive game once thought to involve too many possibilities for computers to master through brute force alone (Silver et al., 2016). Deep learning has catalyzed not just research achievements, but widespread commercial adoption and real-world applications across industries like consumer technology, healthcare, finance, and transportation (Borg et al., 2019). For example, it has been a key driver of recent breakthroughs in autonomous vehicles and robotic systems by providing the sophisticated perception capabilities needed to interpret rich sensory data from the real world (Janai et al., 2020). Overall, this modern machine learning renaissance powered by deep learning has given rise to artificial intelligence systems exhibiting flexible, human-like competencies that were previously deemed far too complicated for computers to master based on conventional programming approaches.

BODY 4: The Rise of Large Language Models (ChatGPT 4 Alongside Claude) (Picture Prompted in the Style of Banksy)

For this paragraph I started by creating a paragraph and in ChatGPT 4 and refining it as I did for the second paragraph. This process consisted of having it create the initial paragraph based off of the initial outline and then asking it if it felt anything was missing. The AI thought that specific examples of how the LLMs worked would be to the benefit of the paper. Then I went over to Claude and asked it to flush out the existing content based on the advice given with the addition of sources and citations that would fit in at those points. I was impressed at choices of sources and academic research; however as you can see in the paragraph many of the citations are seemingly unnecessary, citing general topics rather than specific ideas. This was a relatively seamless experience though as ChatGPT 4 is very good at building up the structure and scaffolding of an essay and then bringing Claude in to fill in those gaps with research and more flushed out ideas seemed to work very well especially in the context of this paper. One thing that I did find lacking in this area was that neither AI suggested adding information relating to the broader implications of their existence, which I had to add on at the end of the paragraph personally.



The next major stage in the history of artificial intelligence has been the The rise of large language models, an AI field focused on human-computer language interaction (Qiu et al., 2020). LLMs' core function is to interpret and generate naturalistic text, enabling machines to comprehend and respond with unprecedented sophistication and fluency. Early language models used simple statistical methods like n-gram modeling to predict the next word based on previous ones - an approach too limited in grasping contextual nuances and long-range dependencies

(Bengio et al., 2003). The field underwent a paradigm shift with the advent of transformer architectures like GPT (Radford et al., 2018) and BERT (Devlin et al., 2019), employing self-attention mechanisms to effectively capture context across entire spans of text. The latest evolution is exemplified by gigantic LLMs like GPT-3 (Brown et al., 2020) which pushed the boundaries with over 175 billion parameters. Not only orders of magnitude larger in training data and model size, but uniquely capable - able to perform a wide variety of language tasks like translation, question-answering, summarization and open-ended text generation without task-specific fine-tuning (Chowdhery et al., 2022). This remarkable versatility is paired with an ability to generate highly coherent, contextually-aware text that is often indistinguishable from authentic human writing. However, LLMs' incredible power brings its own set of risks and challenges that cannot be ignored. Perpetuating societal biases from the skewed demographics of their training data is a prevalent concern (Bender et al., 2021), as is the potential for misuse in generating misinformation or explicit content. The astronomical computing resources required to develop these models compound environmental impact concerns (Strubell et al., 2019) while also raising accessibility issues for all but the largest tech companies. **As LLMs continually evolve and pervade products and services, redefining AI's language frontier, their advanced capabilities must be balanced against profound ethical implications, unsolved technical limitations, and broader societal impacts.** The intensifying research and debate in this arena reflect the critical need to develop robust strategies for responsibly developing, deploying and governing these powerful but imperfect language technologies as they become ubiquitous across our digital world.

CONCLUSION

(GPT 3.5 and Gemini) (image prompted in the style of Monet)

The initial response that GPT 3.5 gave to me was very stale and in line with it's earlier initial answers the earlier problems I noted with GPT of being cyclical in its writing and tending to give to much as a response were ver prevalent in its first response. However, I thought that it had gotten a good grasp for what I wanted out of a conclusion and was only missing the tone and some slight tie ups. In order to clean it up I sought out the strengths of Gemini which were making more condensed and straightforward responses from specific prompts. Gemini did great work with this and created a paragraph that required little to no additional prompting to get something I considered acceptable. I adjusted some minor linguistic choices like starting off with "in conclusion" but this was likely a remnant of the initial ChatGPT prompt rather than a choice of Gemini.



The remarkable journey of artificial intelligence, from the theoretical musings of Alan Turing to today's sophisticated large language models (LLMs), reflects humanity's unwavering quest to understand and replicate intelligence. Turing's foundational work ignited a scientific exploration of machine intelligence, paving the way for generations of researchers. AI's development has been marked by periods of both fervent optimism and frustrating plateaus, yet each phase has yielded valuable insights into the intricate challenges of mimicking human thought. The emergence of deep learning, fueled by the exponential growth of computing power and vast datasets, marked a turning point in AI research. This paradigm shift empowered machines to excel in tasks once deemed intractable, leading to breakthroughs in image recognition, natural language processing, and even strategic games. LLMs represent the latest frontier, pushing the boundaries of human-like text interpretation and generation with unprecedented sophistication. Models like GPT-3 showcase remarkable versatility and fluency, but these advancements come with inherent risks. Societal biases embedded in training data can be reflected in LLM outputs, raising ethical concerns. Additionally, the environmental impact of training these massive models requires careful consideration. As we navigate this new era of AI, a cautious and forward-thinking approach is paramount. Responsible governance, coupled with rigorous ethical considerations, is essential to ensure that AI serves humanity. Ongoing research is critical to mitigate potential harms while maximizing the societal benefits of this powerful technology. Ultimately, the evolution of AI reflects our desire to not only understand intelligence, but to harness its potential for a better future. By approaching AI development with a steadfast commitment to ethical principles, we can ensure that this technology enriches the human experience, serving as a tool for progress and positive change.

Bibliography (Claude AI)

- Sejnowski, T. J. (2018). The deep learning revolution. MIT Press.
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep Learning. MIT Press.
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436-444.
- Schmidhuber, J. (2015). Deep learning in neural networks: An overview. *Neural networks*, 61, 85-117.
- Campbell, M., Hoane Jr, A. J., & Hsu, F. H. (2002). Deep Blue. *Artificial Intelligence*, 134(1-2), 57-83.
- Ferrucci, D., Brown, E., Chu-Carroll, J., ... & Welty, C. (2010). Building Watson: An overview of the DeepQA project. *AI Magazine*, 31(3), 59-79.
- Silver, D., Huang, A., Maddison, C. J., ... & Dieleman, S. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587), 484-489.
- Janai, J., Güney, F., Behl, A., & Geiger, A. (2020). Computer vision for autonomous vehicles: Problems, datasets and state-of-the-art. *Foundations and Trends® in Computer Graphics and Vision*, 12(1-3), 1-308.
- Borg, M., Englund, C., Wnuk, K., ... & Fayyad, U. (2019). Safely Interruptible Agents. arXiv preprint arXiv:1905.07780.
- Qiu, X., Sun, T., Xu, Y., Shao, Y., Dai, N., & Huang, X. (2020). Pre-trained models for natural language processing: A survey. *Science China Technological Sciences*, 63(10), 1872-1897.
- Bengio, Y., Ducharme, R., Vincent, P., & Jauvin, C. (2003). A neural probabilistic language model. *Journal of machine learning research*, 3(Feb), 1137-1155.
- Radford, A., Narasimhan, K., Salimans, T., & Sutskever, I. (2018). Improving language understanding by generative pre-training.
- Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. *NAACL 2019*.
- Brown, T.B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., ... & Amodei, D. (2020). Language models are few-shot learners. *NeurIPS 2020*.
- Chowdhery, A., Narang, S., Devlin, J., Bosma, M., Mishra, G., Roberts, A., ... & Shazeer, N. (2022). PaLM: Scaling language modeling with pathways. arXiv preprint arXiv:2204.02311.
- Bender, E.M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big?. *FACCT 2021*.
- Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. *ACL 2019*.

Analysis of picture 1: DA VINCI

The artwork depicting the evolution of artificial intelligence as a rough pencil sketch in the style of Leonardo da Vinci's lab notebook carries significant artistic and conceptual value:

1. **Historical Connection**: The choice to visualize AI's evolution in the style of Leonardo da Vinci, a quintessential Renaissance polymath known for his forward-thinking and intricate sketches, creates a bridge between past and present intellectual endeavors. It highlights how contemporary ideas are often grounded in historical inquiries and the imaginative visions of pioneers from earlier eras.
2. **Exploratory Nature**: Da Vinci's sketches were known for their exploratory and often speculative nature. By adopting this style, the artwork embodies the inquisitive spirit of early scientific and technological exploration, paralleling the way early thinkers like Alan Turing approached the unknown territories of computing and artificial intelligence.
3. **Conceptual Flow**: The progression from simple sketches of Turing and his foundational questions to more complex representations of modern AI technologies illustrates the trajectory of AI development. This not only serves an educational purpose but also artistically conveys the flow of ideas over time through a connected series of thoughts and innovations.
4. **Creative Visualization**: The use of a rough sketch style emphasizes the raw, unrefined nature of brainstorming and developing new technologies. It reflects the iterative process of scientific discovery, where initial ideas undergo numerous revisions and refinements before reaching their final form.
5. **Inspiration and Insight**: For viewers, especially those familiar with AI and its history, the artwork can serve as a source of inspiration and insight. It visually narrates the story of AI, from theoretical underpinnings to practical applications, capturing the essence of continuous human endeavor to expand our capabilities through technology.

Overall, the artwork is not just a representation of historical facts; it's an interpretative piece that invokes the continuous link between past intellectual achievements and current technological advancements, highlighting the perpetual human quest for knowledge and innovation.

Analysis of picture 2: VAN GOGH

The painting is significant as it visually represents the evolution and conceptual depth of artificial intelligence (AI) through a style reminiscent of Vincent van Gogh, known for his expressive and emotionally resonant artwork. Here are some key aspects that highlight its significance:

1. **Historical and Philosophical Depth**: The artwork incorporates elements from ancient myths and philosophical inquiries that predate modern computing by centuries. Figures such as René Descartes and Gottfried Wilhelm Leibniz, who speculated about the mechanical nature of thought and the potential for a universal language and computational system, are symbolically represented. This ties the roots of AI to deep philosophical questions about the nature of intelligence and the human mind.
2. **Technological Progression**: The painting traces the technological advancements from the 19th-century Analytical Engine, recognized as an early form of computer, to the mid-20th-century developments by Alan Turing, including the Turing Test, which is a seminal concept in AI. These elements illustrate the continuous thread of innovation that connects early mechanical experiments to modern digital computing.
3. **Artistic Interpretation**: Using a style that evokes Van Gogh's vibrant and tumultuous brush strokes, the painting emphasizes the turbulent and revolutionary nature of technological progress. Van Gogh's technique is particularly adept at conveying emotion and transformation, making it a fitting metaphor for the disruptive and transformative nature of AI.
4. **Cultural Reflection**: The piece reflects on how society has viewed and interacted with the concept of artificial intelligence over the centuries. It suggests a dialogue between human creativity and mechanical computation, raising questions about what it means to be intelligent and how technology mirrors or extends human capacities.
5. **Legacy and Vision**: Finally, the painting places a special emphasis on Turing's legacy as a central figure in AI. Turing's contributions are not just historical footnotes but are foundational to ongoing discussions and developments in AI. His vision continues to propel the field forward, symbolized in the artwork by the transition from historical elements to a depiction of modern AI systems and networks.

Overall, the painting serves as both a tribute and a contemplation on the deep, intertwined history of technology, philosophy, and artistic expression in understanding and shaping the future of artificial intelligence.

Analysis of Picture 3: PICASSO

The portrait on the left appears to adopt the cubist style, characterized by geometric shapes and the breaking down of the subject into a series of fragmented, abstract forms. The significant elements likely symbolize aspects of Artificial Intelligence, as per the thematic request.

In the context of the narrative provided about the development of AI:

- The geometric shapes and mechanical gears could represent the early stages of AI, like the perceptron and LISP, where the foundation of AI was being laid out but was still primitive and fragmented.
- The cold and icy textures might symbolize the AI Winter, a period of reduced funding and interest in AI research due to high expectations meeting technological limitations.
- The overall composition could suggest the complex and multifaceted nature of AI progress, which, much like the layers and dimensions in cubist art, builds upon itself in a non-linear fashion.

The portrait, therefore, isn't just a representation of a figure but a narrative piece, each fragment telling part of the story of AI's cyclical history, from its optimistic beginnings, through periods of stagnation, to its eventual resurgence and triumphs.

Picture 4 Analysis: POLLOCK

The abstract expressionist painting, inspired by the field of artificial intelligence and crafted in the style of a 1940s drip painter, holds significant symbolic value. Here are some key aspects that contribute to its significance:

1. **Representation of AI Evolution**: The painting captures the profound transformation in artificial intelligence, from rule-based systems to modern data-driven machine learning paradigms. The chaotic yet harmonious arrangement of splattered and dripped paint represents the complexity and intricate patterns discovered by deep learning systems in massive datasets.
2. **Abstract Expressionism and Technology**: By merging the style of 1940s abstract expressionism—specifically drip painting—with the theme of technological advancement, the painting juxtaposes the seemingly organic, human-centric method of paint application with the precise, data-driven nature of modern AI. This contrast invites contemplation on the relationship between human creativity and machine intelligence.
3. **Visual Metaphors**: The textured layers and varied colors evoke the layers of neural networks and the data streams processed by algorithms, symbolizing how deep learning mimics certain aspects of the human brain. This visual metaphor enhances our understanding of AI as a field that not only replicates but also extends human capabilities.
4. **Dynamic Composition**: The energetic drips and splashes of paint reflect the rapid advancements in AI technology, suggesting a continuous and unstoppable progression. This dynamic composition captures the spirit of innovation and the fast pace of change in the tech world.

5. **Emotional Impact**: Abstract expressionism often aims to convey deep emotional truths through spontaneous and intuitive painting techniques. In this context, the painting may evoke feelings about the rapid encroachment of technology into human spaces, stirring reflections on progress, dependency, and the future of human-machine interactions.

Overall, this painting serves as both a reflection on the state of AI and a piece of art that challenges viewers to think deeply about the interplay between human creativity and artificial intelligence.

Picture 5 Analysis: BANKSY

The mural created based on your description carries several layers of symbolic significance, each tied to the theme of the evolution and impact of large language models in natural language processing (NLP). Here's a breakdown of its elements and their meanings:

1. **Machine-Shaped Brain**: The central image of a brain composed of cogs and gears represents the complex and mechanical nature of language models. This suggests that just as the brain processes thoughts and language, language models process data to 'understand' and 'generate' human language. It also emphasizes the artificial yet intricate nature of these technologies, mirroring the complexity of human cognition.
2. **Human Figures with Digital Devices**: Around the brain, diverse human figures are shown interacting with digital devices. This highlights how integral language models have become in everyday technology, affecting various aspects of human life. The interaction suggests both dependency and engagement, indicating that these tools are now extensions of human communication and learning processes.
3. **Speech Bubbles with Binary Code and Scripts**: The speech bubbles emanating from the devices, filled with binary codes and different scripts, symbolize the global and digital reach of language models. This imagery reflects the ability of NLP technologies to bridge language barriers and facilitate communication across diverse linguistic and cultural backgrounds.
4. **Urban Setting with Technological Elements**: The backdrop of the mural features an urban environment infused with elements like circuit patterns and data streams. This setting reflects the progressive integration of advanced technologies into the fabric of modern society. It also traces the evolution from basic statistical models to sophisticated transformer-based architectures like GPT and BERT, indicating ongoing development and refinement.
5. **Monochromatic, High-Contrast Style**: The choice of a monochromatic palette with high contrast and bold outlines enhances the dramatic and impactful nature of the mural. This style is

not only visually striking but also evokes the stark realities and significant societal impact of technological advancements.

6. **Themes of Technology, Innovation, and Societal Impact**: Overall, the mural captures themes of technological advancement and its profound implications for society. It encourages viewers to reflect on how these tools are shaping communication, information exchange, and even human thought, underscoring the transformative power of AI in the modern world.

This mural, therefore, serves not just as an artistic representation but also as a conceptual exploration of the roles and ramifications of language models and AI in our lives. It invites viewers to ponder the future trajectory of these technologies and their ethical, cultural, and social dimensions.

Picture 6 Analysis: MONET

The painting you see captures a tranquil, misty morning by a calm lake, richly layered with symbolism and reflections on the progression of artificial intelligence (AI) research. Here's a breakdown of its elements and meanings:

1. **Atmosphere and Colors**: The soft hues of blue and gray, along with the delicate brushstrokes depicting reeds and lilies, establish a serene, contemplative setting. The wisps of fog not only add to the mystique of the landscape but also symbolize the unknown territories of AI research.
2. **Sunlight and Warmth**: A faint glimmer of sunlight piercing through the mist represents breakthroughs and advancements in AI. The transition of colors from cool blues to warm pinks and oranges symbolizes the optimism and pioneering spirit of early AI researchers like Alan Turing.
3. **Figures on the Horizon**: The blurred outlines of figures in the distance symbolize AI researchers and pioneers, perhaps a nod to the collective effort of many minds working towards the understanding and development of AI technologies.
4. **Richer Colors and Textures**: As the painting progresses, the colors deepen and become more vibrant, depicting the exponential growth in computing power and the vast datasets that have propelled AI forward. This part of the painting might reflect the modern era of AI, characterized by advanced algorithms and significant technological capabilities.
5. **Ethical Shadows**: The gathering dark clouds and the lone figure standing contemplatively by the water's edge bring attention to the ethical considerations and the need for responsible governance in AI development. This adds a critical layer to the painting, prompting viewers to

think about the consequences and moral responsibilities that come with technological advancements.

6. ****Final Fade and Hope****: The scene gradually fading as if viewed through a lens of memory suggests that while the technological journey of AI is impactful, it also blends into the broader scope of human experience and history. Despite the challenges and ethical dilemmas, there remains a hopeful undertone—a belief in the resilience and potential of humanity to navigate the complexities of AI responsibly.

Overall, the painting is not just a depiction of a natural landscape but a metaphorical canvas that beautifully intertwines the themes of technological innovation and ethical contemplation.