

Caroline Kurzweil

14 April 2024

Professor Isaacson

The Development of AI and its Controversies: An Essay Written with AI

Prompt:

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

Chat GPT 3.5 Answer:

The development of artificial intelligence (AI) has undergone significant stages, evolving from theoretical concepts to practical applications such as large language model chatbots. Here's a brief overview of key milestones in AI development:

Alan Turing (1936-1950s):

Alan Turing laid the theoretical groundwork for AI with his seminal paper "Computing Machinery and Intelligence" in 1950, where he proposed the Turing Test to determine a machine's ability to exhibit intelligent behavior indistinguishable from that of a human.

Early AI Research (1950s-1960s):

The term "artificial intelligence" was coined, and researchers began exploring symbolic reasoning and problem-solving through early AI programs.

The Logic Theorist (1956), developed by Allen Newell, J.C. Shaw, and Herbert A. Simon, demonstrated automated theorem proving.

Expert Systems (1970s-1980s):

AI research shifted towards expert systems, which used predefined rules and knowledge to solve specific problems.

MYCIN (1974), an expert system for medical diagnosis, and DENDRAL (1965), for chemical analysis, were notable examples.

AI Winter (1980s-1990s):

The AI field faced setbacks and reduced funding due to unmet expectations, leading to what is known as the "AI Winter."

Researchers shifted focus to more practical, narrow AI applications.

Machine Learning Resurgence (1990s-2000s):

Advances in machine learning, particularly neural networks, reignited interest in AI.

The development of backpropagation and the application of neural networks to various tasks led to improvements in pattern recognition and classification.

Big Data and Deep Learning (2010s):

The availability of vast amounts of data and increased computational power contributed to the rise of deep learning.

Breakthroughs in image and speech recognition, led by deep neural networks, demonstrated the effectiveness of these techniques.

Natural Language Processing (NLP) and Chatbots (2010s-2020s):

NLP techniques, combined with deep learning, enabled significant progress in understanding and generating human language.

Large language models like OpenAI's GPT (Generative Pre-trained Transformer) series, including GPT-3.5, showcased the ability to generate coherent and contextually relevant text.

Large Language Model Chatbots (2020s-ongoing):

Large language models like GPT-3 have been applied to various tasks, including chatbots that can engage in more natural and context-aware conversations.

These chatbots leverage pre-training on diverse datasets and fine-tuning for specific applications, demonstrating the potential of AI in understanding and generating human-like text.

The development of AI from Turing to large language model chatbots reflects a continuous evolution driven by theoretical advancements, algorithmic improvements, and increased computational capabilities. Ongoing research and innovation continue to shape the field, opening up new possibilities and applications for artificial intelligence.

Writing Phase

My main focus for this article was having the evolution of AI explained to me in business terms and also focusing on the controversies each development produced. Throughout this process, I had a lot of fun getting the different chatbots - Chat GPT 3.5, Gemini, and Claude - to disagree.

My first question to the AIs was what each thought the three most important contributions were to AI. Chat GPT answered first, telling me about the deep learning revolution, the development of large language models, and the practical applications of AI. Gemini answered with the Turing Test and Machine Learning (1930s-1950s), deep learning and increased processing power (1960s-1990s), and large language models (Present). Claude gave Alan Turing's theoretical foundations, the rise of machine learning and neural networks, and large language models and conversational AI as answers. Since the previous answers were similar, I delved deeper by asking each chatbot the specific groups or people that were the most influential. Chat GPT mentioned Alan Turing, Geoffrey Hinton, Yoshua Bengio, and Yann

LeCun, then OpenAI, Google Brain, and DeepMind. Gemini gave Alan Turing (1912-1954), Artificial Intelligence Research Groups (1950s-Present), and Geoffrey Hinton, Yoshua Bengio, and Yann LeCun (1980s-Present). Claude mentioned Alan Turing, the researchers at the Dartmouth Conference (1956), and the teams behind breakthroughs in machine learning and neural networks.

I was unsatisfied with their answers, then realized I should be asking which breakthroughs were most important to their specific chatbot model. By asking each chatbot a general important breakthrough or person, they most likely searched the internet and found common answers. I hoped that by asking what was most important to their specific algorithms, they would give me different answers. Chat GPT and Gemini felt that the development of deep learning techniques, particularly the backpropagation algorithm, contributed most to their creation. They then discussed how this breakthrough, pioneered by Geoffrey Hinton, has fundamentally changed how neural networks are trained and optimized. When I asked Chat GPT why it thought deep learning was most important, it said that GPT's unique scale, generative capabilities, and pre-training methodology make deep learning particularly crucial for its creation and effectiveness compared to other chatbots. Claude gave Anthropic as its biased answer, but this sort of bias is exactly what I was looking for in each chatbot. Anthropic is the company that created Claude, so clearly it would say that their contributions to AI development were the most important. It stated that Anthropic has been at the forefront of exploring ways to imbue AI systems like myself with robust ethical principles, values, and behavioral boundaries.

Since this is a paper on AI, I wondered who better to ask for questions than the AI models themselves. When asked what topic would most likely receive disagreeing answers from multiple chatbots, it suggested asking each model its opinion on controversial AI topics. Chat

GPT then gave me a list of ten controversial topics brought about by the creation of AI, which I then asked each model to answer. At first, each of the three models defined each topic with a neutral view. Unsatisfied, I asked each model to be more controversial, allowing me to feel their clear position on the topics. Claude was the only chatbot that refused to get controversial with its answers, telling me that it was designed to have conversations based on ethical principles and that getting controversial would go against its fundamental purpose. Chat GPT and Gemini, however, had no qualms about getting more controversial, and even took sides that went against the support of AI. In terms of ethical considerations, Chat GPT said that industry self-regulation coupled with ethical guidelines would foster responsible AI development more than strict regulations. Gemini took a more controversial route, simply stating that profit trumps people in our current system. Each chatbot took on the same tone with job displacement as Chat GPT said that society should adapt to change instead of resisting it and Gemini said that AI isn't replacing workers, just replaceable workers. The rest of their answers held similarly worded answers on the following topics: bias and fairness, AI safety, existential risk, privacy, algorithmic transparency, distribution of benefits, human-AI collaboration, and long-term implications. Chat GPT gave thoughtful answers, giving a clear stance but backing it up with ways to fix the issue such as focusing on minimizing bias and ensuring fairness rather than eliminating bias in algorithms due to the inherent biases in data. Gemini, though more entertaining to read, gave highly controversial answers that seemed to be against AI which I found interesting given that it gave answers based on an AI algorithm. However, it did reinstate my faith in AI since it had no bias towards its very creation.

Satisfied with my collected information on the important people and breakthroughs in AI based on each chatbot's opinion, I decided it was time to structure my paper. My first prompt

leading Chat GPT 3.5 towards a structured essay was having it write me a paper about the evolution of artificial intelligence and, after each foundational discovery, include the controversy it presented after each breakthrough. I liked the paper I received since it was broken up into paragraphs of breakthroughs chronologically and was followed by a paragraph about the controversy it introduced. However, I wanted it to be more interesting to read. I had the chatbot rewrite the same paper as though it were a children's story, but it added too much fluff to the story and took out a lot of the controversial discussion. I then asked Chat GPT what style it could write the paper in to make it more unique. It suggested writing the paper as a quest in which the goal was to make AI as developed as possible but each controversy that arose was like a creature the adventurers must face. I went with this idea but found that the writing became very convoluted. Since this topic is very technical with a lot of terms that are still unfamiliar to me, I figured I should stick to a writing style that would not drown out the technical side of this essay. I did not want to go with just a beginner's guide to AI, so I decided to have it tell me the same information in terms I have studied. So I told Chat GPT to tell me about the evolution of AI and the controversies it created - so I asked it to keep all the previous information we had discussed - as though it were talking to a business school student since I am in the A.B. Freeman School of Business. The analogies it gave in terms of business speak were more related to everyday life than specific business terms, but I enjoyed how it made the topics easier to understand.

Just when I thought I was satisfied with the final product and had it add a bibliography, I realized that Chat GPT could not keep up with my add-ons and had not mentioned Alan Turing or any other prominent figures a single time in the paper. This led me to insert a new final prompt and realize I would have to add in some of my information to incorporate other aspects I had asked it to include earlier.

Editing Phase

As mentioned during the writing phase, I went through many versions of this essay before settling on one. The main issues with prior versions were their convoluted verbiage and the lack of information it provided. The multiple chatbots' inability to write me a creative but well-flowing essay without making it confusing proved to me that AI cannot think like a human because of its lack of creativity. Furthermore, even though I did come to a final product that satisfied me, that is not to say the essay was without its faults.

The first check I did was to verify all the sources used for my bibliography. At first, the sources all seemed very reputable and even included prestigious works through Ivy League universities. It even cited the original article on deep learning written by those who pioneered its existence. Problems arose when I got farther down the list of citations. The citation for Radford's paper on language models as unsupervised multitask learners is a good source but was never cited in the paper. I think it should have been around the section on large language model chatbots, either for the explanation or controversy. I am guessing this is based on the paper itself and the fact that the paragraph for this section had no citations, so I added it myself.

During my final read-through, I found lots of the information very vague. The title had no mention of controversies which confused me because that was my primary focus. For controversies, each paragraph was one to two sentences so I could not understand the full point being made. For example, neural networks were compared to a "black box" but did not explain why. With green text, I had to go through the paper and add more explanations based on Chat GPT and articles in the bibliography. I also was unsatisfied with some of the words used, finding their definition a little too far from what they were looking for. I think that when I asked them to

write a technical paper, Chat GPT prioritized higher-level words over the accuracy of the definition.

I also wanted the paper to describe AI to me as though I am a business student, which I am, and one version did this very well. Once I asked it to add in sources, however, the fun analogies went away. This required me to add them back in manually. I also had to add back the controversial opinions Chat GPT originally gave me for each topic.

After all this, the paper had only hit 1,000 words. I then put the essay back into Chat GPT and requested that it double the word count without adding more sources. This request, however, kept returning essays of the same length. No matter how many times I asked it to double the word count, Chat GPT would only add 50-100 words. I then had to enter in the original prompt, ask it to re explain itself to a business student, then enter those into the paper manually. I kept those inserts in black because they were still AI-generated. Still below the word count, I decided to have Chat GPT explain AI winter to me in business terms and had it discover some controversies it brought about. This point was very interesting to me because when I asked it to give me sources, it refused. I found this intriguing because Chat GPT had given me all the previous citations so I asked why it gave me a bibliography earlier and not now. Chat GPT then apologized to me and provided me with an accurate bibliography for information on AI winter. After this point, I had hit the word count and was satisfied both with what I had learned about AI and what we had produced.

The Final Essay

The Evolution of Artificial Intelligence **for Business Students**: From Turing to Large Language Model Chatbots **and Their Controversies**

Artificial Intelligence (AI) has undergone a remarkable evolution since its inception, traversing through various milestones from Alan Turing's conceptualization of intelligent machines to the **emergence** of large language model chatbots. As a business student, understanding this journey is crucial, as AI continues to reshape industries, redefine business processes, and revolutionize customer interactions. **AI has been compared to that of a curious child in the way that it observes, imitates, and experiments until receiving satisfaction from its users (Domingos, 2015).** This essay explores the development of AI, highlighting key controversies at each stage, and describes its implications for business.

Turing's Test and the Birth of AI

Imagine Alan Turing as the visionary architect designing the blueprint for a groundbreaking skyscraper. Just as Turing proposed the Turing Test as a criterion for determining a machine's intelligence, the architect lays the foundation for the skyscraper's construction. The blueprint represents the starting point—a meticulously crafted plan that sparks the imagination and sets the stage for future developments in the world of architecture. Similarly, Turing's work laid the groundwork for AI research, igniting discussions and debates about the potential of intelligent machines to revolutionize various industries, **including business**. Just as the architect's blueprint guides the construction process, Turing's ideas guided researchers as they embarked on the journey to create machines capable of human-like intelligence. The design for AI started in the 1950s with his famous Turing Test. He proposed that if a machine could engage in conversation indistinguishable from a human, it could be considered intelligent (Turing, 1950). **This sparked the beginning of AI research and ignited discussions about its potential in various industries, including business.**

Controversy:

The Turing Test sparked controversy over its adequacy as a measure of genuine intelligence. Critics argued that passing the test did not equate to a true understanding of consciousness, raising questions about the essence of AI (Harnad, 2000). This test poses further issues today as researchers urge those in the AI field to recognize the dangers of creating applications that believably mimic humans (Bender & Gebru, 2021). Furthermore, the Turing Test led to debates about the nature of intelligence and consciousness, raising philosophical and ethical concerns about the implications of AI for humanity's understanding of itself and its relationship with technology.

Expert Systems: Mimicking Human Expertise

Think of expert systems as a state-of-the-art library with a team of expert librarians. Each librarian specializes in a specific field—be it history, science, or literature—and possesses a wealth of knowledge and expertise. Visitors to the library can seek guidance from these expert librarians, who efficiently organize information and provide valuable insights. Similarly, expert systems utilize rule-based approaches to mimic the decision-making abilities of human experts in specific domains. Just as expert librarians facilitate access to knowledge and provide valuable guidance to library visitors, expert systems offer businesses the ability to leverage expertise in various fields and make informed decisions. Like a well-organized library, expert systems streamline processes, enhance efficiency, and empower businesses to solve complex problems effectively. During the 1960s and 1970s, expert systems emerged, aiming to replicate the decision-making abilities of human experts in specific domains (Buchanan & Shortliffe, 1984). These systems utilized rule-based approaches to solve problems and make recommendations,

heralding a new era of AI applications. Business leaders grappled with the trade-offs between efficiency gains and the limitations of rigid decision-making frameworks.

Controversy:

Critics voiced concerns about the limited scope and rigid nature of expert systems. They argued that these systems lacked adaptability and learning capabilities, constraining their effectiveness in dynamic environments. They follow strict rules, which can make their recommendations hard to understand. This lack of clarity makes it tough to trust them, especially in important areas like healthcare or finance (Dreyfus & Dreyfus, 1988). Additionally, the reliance on expert systems raised fears about job displacement and the devaluation of human expertise, sparking debates about the ethical implications of AI for employment and socio-economic inequality.

AI Winter: A Halt to the Development of AI

Occurring during the 1970s and 1980s, enthusiasm and investment in artificial intelligence technology sharply declined after a period of rapid growth and high expectations. Just like seasons change in nature, the AI winter represents a cooling-off period in the AI industry.

During an AI winter, businesses and investors become disillusioned with AI technology due to various factors. These factors could include overhyped promises that fail to materialize, technical limitations that hinder progress, or economic downturns that lead to decreased funding for research and development - much like what can happen to a business ("AI Winter", 2018). For a business student, it's important to understand that AI winters can have significant implications for companies involved in AI research, development, and implementation. Investments may dry up, projects may be shelved, and talented professionals may leave the field, impacting both the short-term and long-term prospects of AI-driven initiatives. However, it's also essential to recognize that AI winters are not permanent. Just as winter gives way to spring, periods of

stagnation in the AI industry are often followed by renewed interest and innovation. Businesses that can weather the AI winter by maintaining strategic focus, adapting to changing market conditions, and continuing to invest in research and development stand to benefit when the industry eventually rebounds.

Controversy:

One notable controversy that arose specifically because of AI winter is the debate over the viability and future of artificial intelligence research and funding. During AI winters, funding for AI research projects may be cut, leading to a slowdown in innovation and a brain drain as researchers leave the field in search of more promising opportunities. This lack of investment can perpetuate a cycle of decline, further dampening enthusiasm for AI technology and exacerbating the challenges faced by the industry. Moreover, the perception of AI as a failed or overhyped technology during AI winters can damage the reputation of the field and deter potential investors and collaborators (“AI Winter”, 2018). This controversy highlights the importance of sustained support for AI research and development, even during periods of uncertainty and skepticism.

Neural Networks Resurgence: Unleashing Computational Power

Consider neural networks as a vast network of interconnected roads in a bustling city. Each road represents a pathway for information to travel, and the intersections serve as decision points where routes are determined. Like the intricate network of roads, neural networks consist of interconnected nodes that process and transmit information. These networks leverage the power of parallel processing to perform complex computations and learn from data. Just as the city's infrastructure facilitates the flow of traffic and enables efficient navigation from one point to another, neural networks harness computational power to analyze data, extract patterns, and

make predictions. Like a city's transportation network, neural networks facilitate the efficient flow of information, enabling businesses to gain valuable insights and make data-driven decisions. In the 1980s and 1990s, researchers like Geoffrey Hinton played a pivotal role in advancing neural network algorithms, paving the way for the modern era of deep learning (Rumelhart et al., 1986). These networks, inspired by the structure of the human brain, demonstrated remarkable capabilities in pattern recognition and machine learning tasks.

Controversy:

The resurgence of neural networks sparked controversy due to their opaque "black box" nature **because it's hard to understand how they make decisions**. Critics raised concerns about the interpretability of neural network decisions, hindering their acceptance in certain applications (Domingos, 2015). Additionally, they can accidentally pick up biases from the data they're trained on, leading to unfair outcomes. So, while neural networks have lots of potential, we need to be careful about how we use them to make sure they're fair and reliable for everyone.

However, eliminating bias in AI algorithms completely may be impractical due to the inherent biases in data and human decision-making. Instead, efforts should focus on minimizing bias and ensuring fairness through continuous monitoring and adjustment. Moreover, the rapid advancement of neural networks raised ethical concerns about their potential misuse for purposes such as surveillance, manipulation, and control, prompting calls for regulatory oversight and responsible AI governance to safeguard against abuses of power and protect individual rights and freedoms.

Deep Learning and Big Data: A Marriage of Algorithms and Data

Picture deep learning algorithms as master chefs in a bustling kitchen, equipped with an extensive pantry of ingredients sourced from around the world. With their culinary expertise and creativity, they can create exquisite dishes by combining the right ingredients in precise proportions. Similarly, deep learning algorithms leverage vast amounts of data—often referred to as big data—to uncover valuable insights and deliver innovative solutions to businesses. Just as the master chefs experiment with different ingredients and cooking techniques to create culinary masterpieces, deep learning algorithms analyze data from diverse sources, identify patterns, and generate actionable insights. Like a well-stocked pantry, big data provides a rich source of information that fuels the creative process, enabling businesses to uncover hidden opportunities, optimize processes, and drive innovation. The 2000s witnessed the convergence of deep learning algorithms and the abundance of big data, catalyzing breakthroughs in AI research (LeCun et al., 2015). Deep learning techniques, coupled with massive datasets, revolutionized tasks such as image recognition, natural language processing, and autonomous driving. **The development of deep learning techniques - particularly the process of fine-tuning neural networks by adjusting connections played a pivotal role in the creation of chatbots like GPT and Gemini. Deep learning techniques also empowered businesses to enhance customer experiences, optimize supply chains, and personalize marketing campaigns.**

Controversy:

Ethical dilemmas surfaced regarding the use of big data for training AI systems. Issues of privacy infringement, algorithmic bias, and potential societal implications fueled debates surrounding the responsible **use of AI technologies (Mittelstadt et al., 2016). While privacy concerns are valid, the potential benefits of AI-powered surveillance, such as crime prevention and public safety, should not be overlooked. Striking a balance between privacy and security is**

essential in modern society. Some opinions even state that privacy is a thing of the past—AI surveillance will only become more pervasive. However, businesses must prioritize ethical considerations and regulatory compliance when leveraging big data and AI technologies to ensure responsible and transparent practices. Moreover, the proliferation of deep learning algorithms raised concerns about the concentration of power and wealth in the hands of tech giants, exacerbating existing inequalities and monopolistic tendencies in the digital economy.

Large Language Model Chatbots: Conversational AI at Scale

Imagine large language model chatbots as skilled diplomats representing a country on the world stage. Like diplomats, these chatbots possess excellent communication skills and cultural awareness, allowing them to engage in meaningful conversations with people from diverse backgrounds. They carefully craft their messages to convey information effectively and build rapport with their audience. Similarly, large language model chatbots leverage advanced natural language processing techniques to understand and generate human-like text, enabling them to engage in conversational interactions with users. Just as diplomats facilitate dialogue and bridge cultural divides, large language model chatbots facilitate communication between businesses and their customers, fostering stronger relationships and driving engagement. Like skilled diplomats navigating diplomatic negotiations, these chatbots navigate complex conversations, address customer inquiries, and provide personalized assistance, enhancing the overall customer experience and driving business success. In the 2010s, large language model chatbots emerged as the peak of AI development, exemplified by models like GPT. These models, trained on vast amounts of text data, demonstrated remarkable proficiency in generating human-like text and engaging in conversational interactions (Radford et al., 2019). These chatbots, exemplified by

models like GPT, offer businesses the ability to provide personalized and natural language interactions at scale.

Controversy:

The advent of large language model chatbots introduced controversies surrounding misinformation, fake content generation, and ethical considerations. Concerns about the potential misuse of AI-generated text and its impact on society underlined the need for responsible AI governance. Concerns have also been raised about the danger of applications that allow chatbots to mimic humans due to the additional risks they raise by generating believable content (Bender & Gebru, 2021). But Chat GPT shared that the existential risk posed by AI is overstated. AI researchers and policymakers should focus on more immediate ethical and societal concerns rather than speculative scenarios of AI surpassing human intelligence. Moreover, the rise of large language model chatbots raised questions about the authenticity of online interactions and the erosion of trust in digital communication channels, highlighting the importance of transparency, accountability, and ethical standards in the development and deployment of AI technologies. Overall, the evolution of AI from Turing to large language model chatbots has been marked by significant technological advancements and ethical dilemmas. As business students, it is imperative to comprehend these developments, as AI continues to shape the landscape of industries and pose complex challenges for businesses. Some say human-AI collaboration is a myth and that AI will eventually outperform humans in all tasks, making collaboration unnecessary. However, the fear of AI replacing humans in the workforce is exaggerated. Rather than viewing AI as a threat, humans should embrace its potential to augment our capabilities and enhance productivity in various fields. By navigating the controversies surrounding AI responsibly, businesses can harness its transformative potential while mitigating risks and

upholding ethical standards. The long-term implications of the AI revolution are uncertain and may not be as significant as some predict. Rather than speculating about distant future scenarios, society should focus on addressing immediate challenges and maximizing the benefits of AI for business in the present. As AI continues to evolve, business students need to remain engaged with the latest developments and debates surrounding this technology. By understanding the historical context and controversies surrounding AI, students can develop a nuanced perspective that enables them to navigate the complex ethical, social, and economic issues associated with its adoption. Moreover, business leaders must prioritize responsible AI governance and invest in strategies that promote transparency, fairness, and accountability in the development and deployment of AI systems. This includes actively addressing issues such as algorithmic bias, data privacy, and the impact of automation on employment and society.

In conclusion, the evolution of AI has been characterized by a series of milestones and controversies that have shaped its trajectory and impact on business and society. From the theoretical concepts of Alan Turing to the practical applications of large language model chatbots, AI has transformed the way businesses operate, interact with customers, and make decisions. As AI technologies continue to advance, business students must engage with the opportunities and challenges they present, ensuring that AI is harnessed in a manner that promotes human well-being, economic prosperity, and social progress. By embracing a holistic approach to AI education and governance, businesses can leverage the full potential of this transformative technology while addressing its ethical and societal implications responsibly and sustainably.

Bibliography

"AI Winter: Causes and Mitigation Strategies." IEEE Spectrum, 2018.

Bender, E. M., & Gebru, T. (2021). On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? Retrieved from <https://arxiv.org/abs/2105.14017>

Buchanan, B. G., & Shortliffe, E. H. (1984). Rule-based expert systems: The MYCIN experiments of the Stanford Heuristic Programming Project. Addison-Wesley.

Domingos, P. (2015). The master algorithm: How the quest for the ultimate learning machine will remake our world. Basic Books.

Dreyfus, H. L., & Dreyfus, S. E. (1988). Mind over machine: The power of human intuition and expertise in the era of the computer. Free Press.

Harnad, S. (2000). Turing indistinguishability and the epistemology of hypotheticals: A reply to Dietrich. *Psychology*, 11(060).

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436-444.

Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2), 2053951716679679.

Radford, A., Narasimhan, K., Salimans, T., & Sutskever, I. (2019). Language Models are Unsupervised Multitask Learners. Retrieved from https://cdn.openai.com/better-language-models/language_models_are_unsupervised_multitask_learners.pdf

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, 323(6088), 533-536.

Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59(236), 433-460.