AI and the Pursuit of General Intelligence

A Research Paper submitted to the Department of Engineering and Society

Presented to the Faculty of the School of Engineering and Applied Science

University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree

Bachelor of Science, School of Engineering

Victor Xia

Spring 2023

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Advisor

Joshua Earle, Department of Engineering and Society

STS Research Paper

Introduction and Methods

The idea of an artificial intelligence, or machines capable of thought, has captured the imagination of science fiction writers and filmmakers for decades. Until recent years, that is all that artificial intelligence was in the public eye—a futuristic concept used in a dystopian or utopian context for novels. Today, a new wave of discourse on the relationship between artificial intelligence(AI) and humanity is on the rise. At the forefront of this discourse are AI models, like ChatGPT and DeepAI, that are "breaking ground in a much broader range of tasks", especially in the creative fields (Larsen & Narayan, 2023). Along with the rise of AI and machine learning models comes constant warnings by the researchers of the limitations of the models; artificial intelligence, as a field, is not yet capable of reproducing what is traditionally considered higher level intelligence. The limitations of the training data or expert knowledge used to create a model determine the extent to which the model can accurately represent and analyze information. The need to clarify that AI models are not intelligent shows the discrepancy between the traditional imagination of what an AI looks like versus the technological field that is being actively developed and researched.

This paper will explore how society has grown enamored with the concept of machine intelligence, in contrast to the current trajectory of the field towards creating assistive tools for human endeavors. The STS framework I will use to frame the paper will be the social construction of technology(SCOT), or how various actors in society influenced the development of AI as a technological field. Some actors that are of particular interest are science fiction writers, AI researchers, the government, and private corporations. To conduct this research, I will use literature synthesis and present my findings in chronological order, from past developments

to present. Given the sheer amount of history involved in the development of AI, however, I will be covering select parts of the history that made AI the way it is today. In this paper, I will discuss the concept of machine intelligence in literature, then pivot towards the progression of AI research throughout the 20th and 21st centuries. Finally, I will discuss the role of AI in modern society and tie together the paper with a discussion on realigning the public perception of AI with the current state of AI.

Machine Intelligence in Literature

Novelists and authors have utilized literature to explore the concept of intelligent machines, or machines capable of individual thought, from as early as the 19th century. One of the earliest pieces of literature discussing machine life and intelligence comes from Samuel Butler in 1872, in his novel *Erewhon.* Meant to be a piece set inside his fictional society with most mechanical inventions destroyed, Butler discusses the definition of consciousness, and how while "there is probably no known machine which is more than a prototype of future mechanical life", machines may have the capability to develop consciousness, just like how "the race of man has descended from things which had no consciousness at all" (Butler, 1872). This is one of the earlier pieces of literature to compare machines to humans and explore if machines can have human-like intelligence. To put into perspective, standard technologies today such as the telephone and the lightbulb were in their infancy in the 1870's, with computers decades away from being developed.

While Butler's book shows a long-standing fascination with the concept of intelligent machines, more recent popular films and novels like *2001: A Space Odyssey* and *Wall-E* portray a more mature vision of what artificial intelligence could look like. When people think about artificial intelligence, the first thing that typically comes to mind are the novels and films that

have prominent AI characters, ranging from Wall-E and HAL to the AI assistant J.A.R.V.I.S from the film *Iron Man*. In the film *2001: A Space Odyssey* (Kubrick, 1968), the artificial intelligence HAL is portrayed as the main antagonist. HAL endangers the crew of *Discovery*, the starship upon which the story takes place, due to a command to withhold information from the crewmates conflicting with his programming for "the accurate processing of information without distortion or concealment." The film portrays a dystopian view on artificial intelligence, highlighting the inhumanity of HAL's decisions where it decided to kill the crew to resolve the conflicts within its programming. As an artificial intelligence, HAL is also shown to have motivations and higher-level thoughts like other characters in the story, not just a device driven by a human mastermind.

In contrast to *2001: A Space Odyssey*, the Disney Pixar film *Wall-E* (Stanton et al., 2008) portrays the possibility of humanity in AI. Wall-E is the protagonist of the story, being the last remaining trash-compacting robot on Earth. As a robot, Wall-E is capable of higher-level thought and takes many initiatives that would no doubt deviate from his initial purpose as a trash compacting robot. Wall-E represents the evolution of machine life, where he is portrayed to have a very human-like personality while taking advantage of his machine identity through bodily repair and data transfer. In contrast, humanity in the film is portrayed to have grown over-reliant on AI, where humans have reduced authority over the running of the starships and have grown obese and complacent over their comfortable lives. *Wall-E* explores the idea that machines can have intelligence and a sense of humanity. As a film oriented towards children, it also helps instill an image in children's minds of what artificial intelligence may look like in the distant future.

A paper published by Cave and Dihal in 2019, "Hopes and Fears for Intelligent Machines in Fiction and Reality", sums up the common themes that the public associate with the idea of artificial intelligence. Cave and Dihal (2019) state that there are four dichotomies presented in fictional AI—immortality/inhumanity, ease/obsolescence, gratification/alienation, and dominance/uprising. In *Wall-E* (Stanton et al., 2008), we observe the immortality of AI, as seen through the age of Wall-E as well as the resurrection of its programming at the end of the film. In *2001: A Space Odyssey* (Kubrick, 1968), we see the inhumanity of AI, where HAL fails to recognize the importance of human lives over its programmed directive. The humans in *Wall-E* portray the ease, as well as obsolescence brought by AI; they are relieved from work duties, but also left with little purpose to their lives. These dichotomies form part of the public perception of AI, and shape our imagination when envisioning a future with artificial intelligence.

The Pursuit of Artificial General Intelligence

While the concept of machine intelligence has been explored in stories and literature since the 19th century, scientific research and discussion into machine intelligence did not kick off until the 1940s. It should be noted that AI research in the early 20th century was scattered and explored many different fields, each with a different viewpoint on how machine intelligence could look like. Various discoveries, including the binary nature of neuron pulses, the theory of computation, and the field of cybernetics, inspired the eventual formulation of the field of artificial intelligence research in 1956.

One of the major scientific discoveries that made the concept of artificial intelligence possible was the conception of a Turing machine by Alan Turing (1936), describing a device that could compute any computable function. His theoretical machine consisted of an infinitely long tape representing data storage, a read-write head to interact with the data contained in the tape,

and a finite set of instructions for writing to and moving the tape. The Turing machine proposed by Turing included many important insights and ideas,such as binary arithmetic, logical operations, and stored programs, that were later used in the development of early electronic computers. Some of the more significant advancements in digital computer technology as a result of Turing's ideas include the Electronic Numerical Integrator and Computer (ENIAC) and the Electronic Delay Storage Automatic Calculator(EDSAC) (Williamson, 2021). The EDSAC was considered one of the first computers to be Turing-complete, or capable of performing any computation that is algorithmically possible. With the advancements in computation and the creation of Turing-complete computers, researchers could freely formulate algorithms for an electronic mind, knowing that a device capable of running the algorithm will exist in the future.

Another significant scientific advancement that contributed to the start of artificial intelligence research was the introduction of cybernetics in the mid-20th century. In Norbert Wiener's 1948 book, *Cybernetics: Or Control and Communication in the Animal and the Machine*, he proposed that machines can be programmed to mimic human decision-making processes through the concepts of feedback and control. Through feedback and control theory, Wiener, as well as other influential cybernetics researchers, provided insight into how a system can be programmed to learn, adapt, and make decisions in a changing environment. The theoretical framework established by the field of cybernetics would then go on to inspire modern algorithms integral to artificial intelligence, such as neural networks and reinforcement learning.

Advancements in neurology inspired many of the core concepts in the field of artificial intelligence, and provide insight into why artificial intelligence as a field is structured the way it is now. One such discovery is the all-or-none response of nerve impulses, discovered by Edgar Adrian in 1932. The all-or-none response of neural impulses meant that the output of neurons

were binary in nature, a trait conducive to the binary computations able to be performed by a digital computer. In 1943, Warren McCulloch and Walter Pitts developed a simple mathematical model of a neuron, where the neuron takes in a weighted set of inputs from other neurons, computes the total weighted sum of the inputs, and outputs a binary value based on whether a threshold is exceeded or not. The system of neurons proposed by McCulloch and Pitts was one of the earliest neural networks, modeling the relationship between neurons in a network.

While McCulloch and Pitts proposed a mathematical model for the neuron, their model did not address the capability for neurons to learn and adapt over time. In 1949, Donald Hebb theorized in his book *The Organization of Behavior* that when two neurons are activated together, the strength of their connection is increased, also known as Hebbian learning (Hebb, 1949). Hebb's theory of synaptic plasticity inspired the further development of neural networks through altering the weights between neurons based on whether the network makes a correct or incorrect prediction. This process of altering neural weights based on the input data and desired output combines both Hebb's theory of synaptic plasticity and the field of cybernetics to create a feedback loop in the neural network. Through this basic framework for learning algorithms, it became possible for algorithms to learn and adapt from the accuracy of their outputs, a trait traditionally associated with higher-level intelligence.

With all the technological advancements making machine intelligence seem plausible, it was only natural for the scientific community to consider a question asked by Butler in 1972: can machines think? In 1950, Turing indirectly answered this question in his paper "Computing Machinery and Intelligence" by proposing the Turing test, an experiment to determine if an AI can be distinguished through anonymous chat (Turing, 1950). Given Person A, Person B, and an interrogator, if the interrogator cannot determine who the machine is after conversing with both

Person A and Person B via text, then the machine has passed the Turing test. The philosophy behind the test is that if a machine's responses are indistinguishable from that of a human, then it can be considered for all intents and purposes intelligent. While many criticize the test for focusing on natural language processing and communication, the test remains an important factor leading up to the establishment of artificial intelligence as a field, since it establishes a goal researchers can work towards when developing the capability of their AI models (Moor, 1976). The Turing test is also an example of how mid-20th century AI research focused on the concept of artificial higher intelligence, and how humans can determine if a machine has intelligence.

Up until 1956, research into the various areas of artificial intelligence was scattered, with different focuses by different researchers. The 1956 Dartmouth Conference aimed to bring together the researchers from different fields to define the future direction of this new field of research (McCarthy et al., 1955). The name that resulted from the conference to define the field of study was "artificial intelligence". Many consider this event as the start of the field of AI research. In the proposal for the conference, formulated by McCarthy and Minsky, they discussed subjects such as natural language processing, neural nets, and machine learning, topics in AI that continue to be researched and developed today.

Leading up to the 1956 Dartmouth conference that defined the field of AI research, researchers focused on recreating aspects of the human mind in the hopes of creating machine intelligence, otherwise known today as artificial general intelligence(AGI). Cybernetics aimed to recreate human capability to respond based on stimuli, while neural networks aimed to simulate the neural processes within the brain that allow humans to think and function. Discussions on machine intelligence, like Turing's paper on "Computing Machinery and Intelligence", focused on the metric by which intelligence can be ascribed to a machine. There is likely some degree of

correlation between the research focus of the mid-20th century AI researchers and the stories and concepts of machine intelligence present in literature at the time. To the researchers, they were progressing towards machine intelligence, a concept explored in literature and fictional stories. Given the initial progress into the field of AI research, researchers believed they would have "in... three to eight years... a machine with the general intelligence of an average human being" (Minsky, 1970).

Artificial Intelligence in the Late 20th Century

Entering the late 20th century, the focus of AI research gradually shifted away from the pursuit of artificial general intelligence, instead focusing on commercial applications of AI technology. The shift towards commercial applications of AI technology came as a result of funding issues for the field. As research into the technology continued into the late 20th century, little progress was made towards creating an AGI as initially promised during the early years of AI research. Institutions funding AI research at the time included the Japanese government, US government agencies like the Defense Advanced Research Projects Agency (DARPA), and universities like MIT, Carnegie Mellon University, and Stanford University. Due to the slow progress of the field towards the initially promised AGI, institutions grew disillusioned with the field, and funding declined in the late 20th century. As a result of the decline in funding, artificial intelligence grew narrower in scope compared to the vision of early AI researchers (National Research Council. 1999). In the late 20th century. funding for AI largely came from private companies who were more interested in the commercial applications of the technology, rather than the grand promise of an AGI (Newquist, 1994).

One of the first commercial applications resulting from AI research was the expert system, first developed by Edward Feigenbaum at Stanford University in the 1970s. These

systems differed from the neural network approach, instead being a program that responded to queries using a set of rules and heuristics carefully crafted by experts in the field of interest (Tan, 2017). This meant that expert systems were heavily dependent on the rulesets designed by experts, and could not learn or adapt from training data. Expert systems were widely used in corporations and governments in the 1980s and 1990s, in fields ranging from medical diagnosis to computer system configurations and traveler risk assessment. Ultimately, expert systems were limited by their pre-programmed nature, and were more akin to a correlation tool rather than a system capable of intelligent behavior.

Another commercial application resulting from AI research was the development of fuzzy logic systems, a tool used in both control systems and artificial intelligence. Developed in the 1960s by Lotfi Zadeh, fuzzy logic served to mimic human reasoning by assigning partial truth values to data, a representation of the uncertainty and imprecision in human data interpretation (Zadeh, 1965). Similar to expert systems, fuzzy logic systems operated using a ruleset crafted by experts, with the goal of the system functioning as if it was being operated by an expert. Fuzzy logic systems were most notably applied to household appliances like washing machines and temperature regulation systems (Singh et al., 2013). My technical project utilized fuzzy logic controllers as a tool to regulate and maintain a desired set temperature for a portable temperature-controlled container.

The shift of AI funding towards the commercial sector has resulted in the gradual shift in the field of AI in the late 20th century away from general intelligence. The public perception of AI was still that of machines with human-like intelligence and reasoning, while organizations were instead funding AI development towards decision-making and business tools. Applications like expert systems and fuzzy logic systems were a far cry from what was originally promised

from artificial intelligence research, and they did not seem to get any closer to the ideal of a general intelligence. Many AI researchers called their research by terms other than AI, such as informatics or knowledge-based systems, which made it seem like progress in the field had come to a standstill. The term "artificial intelligence" was tainted by the failures of the earlier generation to deliver on their promises of an AGI, and this led to a quiet shift in focus in AI research, largely shielded from public perception (McCorduck, 2004).

Rise of Neural Networks and Big Data

Going into the 21st century, one of the key features that fueled advancements in AI technology was the advancement of computation. The neural networks first developed in the mid-20th century were contingent on training data and computational power to allow the models to learn and adapt to the data. Both data storage and computational power came at a premium during the mid-20th century, but advancements in storage and computation made many of the algorithms and neural networks more feasible to implement and train.

One of the major developments of the 21st century that helped spur the growth of AI was the backpropagation algorithm. Backpropagation was a technique used to adjust the weights of a neural network, with the goal of aligning its output results with the expected output results. In other words, backpropagation made it possible for neural network AIs to leverage the structure of a neural network and train itself to produce more accurate results. First proposed in 1974 by Werbos, backpropagation was limited in development by the computational power of computers in the late 20th century (Werbos, 1974). Through advancements in computation, backpropagation became one of the most important techniques for allowing neural networks to learn and adapt from complex data. Through the use of neural networks and backpropagation, organizations could create AI that did not require expert knowledge or occasional maintenance, a flaw

observed in AIs like expert systems and fuzzy logic systems. Unlike the progression of AI in the late 20th century, realizing the capability of learning in machines seemed like a step towards a true artificial general intelligence.

Another major development of the 21st century for AI was the rise of big data. Along with the advancements in data storage and computation technologies came an influx of data from digital technologies and devices. The development of digital technologies has also resulted notably in the digitization of most of the data in the world, with an estimated 94% of all information storage being digital by 2007 (Hilbert & López, 2011). With backpropagation techniques and neural networks, artificial intelligence technology had great potential in being able to leverage big data to generate insights. In particular, AI had the capability to learn trends and patterns from hugely complex datasets like natural language and images, a task too complex for traditional statistical analysis techniques previously used to process and generate insights from data. If a concept had a digital representation, then an AI could learn how to perform that concept with enough training data and tuning. In a way, the rise of big data represented the alignment of the world's information to be in a form that AI could understand and learn from.

The effect that big data had on the development of AI and machine learning models can be seen from many applications and services that people interact with on a daily basis, like social media and voice assistants. One key application of big data and AI is in personalization, or the tailoring of content on a service to be aligned closer to the interests of the user. Social media is one of the most prominent areas of personalization, with apps like Instagram and TikTok utilizing the data generated by user interactions with the app to predict content the user would likely engage with (Invisibly, 2022). Personalization is the concept of making a service align more closely to one's own interests and wants, something that is typically done through the user

deciding what they themselves want to see, or professionals recommending content the user should see. The role of AI in personalization is to understand the wants and interests of the user, then predict what they would like to see. As an application of AI, personalization has commercial potential in increasing user interest in a service or application. That, however, is its limitation; the usage of big data and AI for personalization has only commercial applications, with little advancement towards an AGI or envisioned representations of intelligent machines.

Artificial intelligence research and applications are often perceived as neutral and unbiased, but that simply does not hold true when the data they are trained upon are biased. Big data brought about great advancements to the capability of AI, but also introduced biases and issues that are commonly associated with statistics and big data. One such example of the bias of AI algorithms comes in the form of confirmation bias for racist alt-right movements. Jessie Daniels cites AI algorithms that power the internet for their ability to "deliver search results for those who seek confirmation for racist notions and connect newcomers to like-minded racists" (Daniels, 2018). We think of AI to be objective and morally right, but morals are often a matter of personal opinion. The power of AI to foster the growth of the violent alt-right movement responsible for the 2017 Charlottesville "Unite the Right" rally shows the dangers of unregulated AI in the spreading of ideologies and potentially false information.

After decades of AI research staying under the radar of the public, one particular application of big data and AI seems to rekindle the idea of an artificial general intelligence: generative AI. The potential of generative AI lies in its capability to create responses to any requests asked by the user. One such program that showcases the potential of generative AI is ChatGPT. ChatGPT, released in November 2022, is an AI model trained using the vast amount of public data and information aggregated on the internet, and is capable of performing a wide

range of requested tasks, from generating code to explaining concepts and writing papers (Larsen & Narayan, 2023). In other words, ChatGPT manages to perform at a human-like level in creative tasks that previous AI models struggle with. As a chatbot in the digital age where text communication is commonplace, interacting with ChatGPT feels like interacting with another human over the internet, given how capable the AI is at interpreting anything the user says. Models like ChatGPT bring a new wave of interest towards what it means for a machine to be intelligent, and a rekindled interest towards the Turing test that Alan Turing proposed back in 1950. The rise of generative AIs like ChatGPT have inspired extra research into the concept of intelligence; on February 17th, 2023, Sejnowski released the paper *Large Language Models and the Reverse Turing Test*, exploring the intelligence of models like ChatGPT (Sejnowski, 2023). With the development of generative AI, the field of AI has taken a step towards the concept of intelligent machines, as first popularized back in literature and media.

Conclusion

The start of artificial intelligence, or machine intelligence, was rooted in the idea of a machine capable of displaying human levels of intelligence. Inspired by literary works exploring machine intelligence, scientists had hoped to replicate the thought processes of the human mind using the digital 0's and 1's that make up a traditional electronic computational device. In the early days of AI research, with sufficient funding and lack of commercial pressure, both researchers and the public had aligned public perceptions of what AI should look like: the intelligent machines that seemed no different from a human in behavior.

With the shift of AI research into commercial applications and programs designed to learn and predict outputs from a tailored set of data, a discrepancy formed between the early vision of AI, and the actual state of AI research. Since the public perception of artificial

intelligence has remained largely in line with the early researchers, people envision a machine with higher-level intelligence when they think of AI. This public perception is reinforced by films like *Wall-E* and *2001: A Space Odyssey*, both showing the idea of an artificial general intelligence with human-like behavior. Instead, much of what AI is can be seen in devices and applications used by people everyday. Social media personalization, Netflix show recommendations, fridge temperature control, and other features are all possible through AI technologies, as algorithms to improve the functionality of a device or service. Until the scientific community comes to a consensus on whether a particular AI model has intelligence, it can only be assumed that AI will remain a supportive tool for applications, rather than the intelligent robots envisioned back in Butler's *Erewhon*.

Rather than focus on whether AI will become artificial general intelligence, it is more important to understand the role of AI in society today. With the advent of generative AIs like ChatGPT, we should understand the role and risks of big data in determining the capabilities of an AI model. As a rapidly developing field whose capability is powered by the availability of big data, understanding the role of data in AI can help the public make more informed decisions towards the relationship between humans and AI. The dream of artificial general intelligence is not dead yet, but it is not likely to happen soon. By educating the public about the current state of AI, researchers and the public alike can once again realign to advance AI technology, whether that leads the field further into tools for human use, or back towards the dream of a true artificial intelligence.

References

Adrian E.D. (1932) *Mechanisms of nervous action: electrical studies of the neurone*, University of Pennsylvania Press, Philadelphia

Butler, Samuel. (1872). *Erewhon, or, Over the range.* Trübner & Co.. Retrieved from https://doi.org/10.5479/sil.1036699.39088016476525

Cave, S., & Dihal, K. S. (2019). Hopes and Fears for Intelligent Machines in Fiction and Reality. *Nature Machine Intelligence*, 1 (2), 74-78. https://doi.org/10.1038/s42256-019-0020-9

Daniels, J. (2018). The Algorithmic Rise of the "Alt-Right." *Contexts*, *17*(1), 60–65. https://doi.org/10.1177/1536504218766547

Hebb, D. O. (1949). *The organization of behavior; a neuropsychological theory.* Wiley.

Hilbert, M., & López, P. (2011). The World's Technological Capacity to Store, Communicate, and Compute Information. *Science*, *332*(6025), 60 –65. doi:10.1126/science.1200970

Invisibly. (2022, March 29). Personalized Feeds: What Users & Brands Need to Know in 2022. *Invisibly Blog*. Retrieved from https://www.invisibly.com/learn-blog/personalized-feeds/

Kubrick, S. (Director). (1968). *2001: A Space Odyssey* [Film]. Stanley Kubrick Productions.

Larsen, B. & Narayan, J. (2023, January 9). Generative AI: a game-changer that society and industry need to be ready for. *World Economic Forum*. Retrieved from https://www.weforum.org/agenda/2023/01/davos23-generative-ai-a-game-changer-industries-and-society-code-developers/

McCarthy, J., Minsky, M. L., Rochester, N. & Shannon, C. E. (1955). A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence. http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html

McCorduck, Pamela (2004). *Machines Who Think* (2nd ed.). A. K. Peters.

McCulloch, W.S., Pitts, W. A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics* 5, 115–133 (1943). https://doi.org/10.1007/BF02478259

Moor, J.H. An analysis of the turing test. *Philos Stud* 30, 249–257 (1976). https://doi.org/10.1007/BF00372497

National Research Council. (1999). *Funding a Revolution: Government Support for Computing Research.* Washington, DC: The National Academies Press. https://doi.org/10.17226/6323.

Newquist, HP. (1994). *The Brain Makers: Genius, Ego, And Greed in the Quest For Machines That Think*, New York: Macmillan/SAMS

O'Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy* (1st ed., Vol. 4.1). Penguin Books.

Sejnowski, T. J. (2023). Large language models and the reverse turing test. *Neural Computation*, 35(3), 309-342. https://doi.org/10.1162/neco_a_01563

Shahab, D. (2019, November 10). Traditional Statistics vs. Artificial Intelligence and Machine Learning. *Journal of Petroleum Technology*. Retrieved from https://jpt.spe.org/traditional-statistics-vs-artificial-intelligence-and-machine-learning

Harpreet Singh, Madan M. Gupta, Thomas Meitzler, Zeng-Guang Hou, Kum Kum Garg, Ashu M. G. Solo, Lotfi A. Zadeh. (2013). *Real-Life Applications of Fuzzy Logic*. Advances in Fuzzy Systems, vol. 2013, https://doi.org/10.1155/2013/581879

Stanton, A., Morris, J., Reardon, J., Burtt, B., Knight, E., & Garlin, J. (2008). *Wall-E*. United States; Walt Disney Studios Motion Pictures.

Tan, H. (2017, September). A brief history and technical review of the expert system research. *Materials Science and Engineering Conference Series, 242*, 012111. doi:10.1088/1757-899X/242/1/012111

Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59, 433–460. https://doi.org/10.1093/mind/LIX.236.433

Turing, A. (1936). On Computable Numbers, with an Application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society* 42 (1):230-265.

Werbos, P. (1974). *Beyond regression: New tools for prediction and analysis in the behavioral sciences. Thesis (Ph. D). Appl. Math. Harvard University.*

Wiener, N. (1961). *Cybernetics, or control and communication in the animal and the machine* (2nd ed.). John Wiley & Sons, Inc.; Boston Review. https://doi.org/10.1037/13140-000

Williamson, T. (2021). History of computers: A brief timeline. *Live Science*. Retrieved from https://www.livescience.com/20718-computer-history.html

Zadeh, L.A. (1965). *Fuzzy Sets.* Information Control, 8, 338-353. http://dx.doi.org/10.1016/S0019-9958(65)90241-X