

The History of Artificial Intelligence:
From Turing to Large Language Model Chatbots

Isabel Kilday
HISU-2900 Digital Revolution
April 14, 2024

The Initial Chatbox Response:

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

Gemini's response:

The journey from theoretical AI to chatbots we interact with today is a fascinating one, paved by brilliant minds and groundbreaking inventions. Here's a glimpse into that journey:

1. The Seeds of AI: Alan Turing (1950s)

- Alan Turing, a pioneer of computer science, laid the foundation with the Turing Test. This test proposed a way to gauge a machine's intelligence by its ability to mimic human conversation.
- He also conceptualized the Turing machine, a theoretical model of computation that's considered a cornerstone of computer science.

2. Early Building Blocks: Perceptrons and Coining the Term (1950s)

- In 1950, Alan Turing proposed the "Imitation Game" which later became known as the Turing Test.
- The invention of the perceptron by Frank Rosenblatt in 1958 marked a significant step. This early neural network, though simple, laid the groundwork for the more complex deep learning models used today.
- Around the same time, John McCarthy, another AI forerunner, coined the term "Artificial Intelligence" at a historic workshop in 1955.

3. Funding Fluctuations and Rise of Deep Learning (1960s-1980s)

- Despite early advancements, the field of AI faced challenges in the 1960s and 1970s due to limited funding and computational power.
- However, the groundwork was being laid for future breakthroughs. Research in areas like machine learning continued, and the concept of deep neural networks with multiple layers began to take shape.

4. AI Triumphs and Complex Learning (1990s-2000s)

- The 1990s saw a resurgence of interest in AI. IBM's Deep Blue defeating chess grandmaster Garry Kasparov in 1997 marked a significant milestone.
- This period also saw advancements in deep learning algorithms and the development of more powerful computers that could handle complex calculations required for training these models.

5. The Age of Chatbots and Large Language Models (2010s-Present)

- The 2010s witnessed a boom in AI applications, including the rise of chatbots. These chatbots, powered by natural language processing (NLP) techniques, could hold basic conversations with users.
- The development of large language models (LLMs) like me took AI to a new level. LLMs are trained on massive amounts of text data, allowing them to generate human-quality text, translate languages, write different kinds of creative content, and answer your questions in an informative way, just like I am doing now!

This is a simplified overview, of course. The field of AI continues to evolve rapidly, with new advancements happening all the time.

Essay Improvement Narrative:

Before beginning this project, I experimented frequently with Chat GPT-4, attempting to use it for a variety of purposes, with varying levels of success. For example, when I was studying abroad in Chile last year, I wanted to make chocolate chip cookies with my host dad. However, I could not find a recipe that was 1. In Spanish; 2. Using the metric system measurements; and 3. Substituting ‘panela’ (brown sugar loaf) for regular brown sugar. Thus, I decided to have Chat GPT-4 generate a recipe for me. Ultimately, my host family loved my AI-generated cookies, but they didn’t taste quite right to me. Whether or not this was due to country-specific ingredient differences or the limitations of artificial intelligence is impossible to say.

I have also used Chat GPT-4 to help me plan travel, apply to medical school, craft difficult texts, and perform translations between languages. This experimentation has allowed me to gain a comprehensive understanding of this LLM’s capabilities and limitations. More recently, I have begun using Gemini, as it has been integrated into Google’s search features, and Claude, since it was recommended in class. I was less familiar with these LLMs, so I was curious about implementing them for this project.

As Gemini states, the initial response to the essay response, was a “simplified overview.” I wanted the paper to be much more in-depth and much longer. From my experimentation with Large Language Models (LLMs), I knew that it can be quite difficult to get them to write longer narratives. Additionally, if you ask them to elaborate generally on the prompt, they will often repeat information already stated.

Given this knowledge, I knew I needed to break up the essay into smaller chunks. For a history essay, it made sense to do this by time period or historical figure. I also decided to have the LLM write the introductory paragraph and conclusion paragraph after completing the body sections, so that they would accurately reflect the content of the paper.

Since I wanted to be creative with the structure of this paper, I decided to tell the story of the development of AI through some of the most important historical figures involved in its inception. This reflects the importance of collaboration across disciplines and building on the work of ones predecessors, which are important themes in this course’s book, *The Innovators*. A bonus of this style is that it naturally breaks up the structure of the paper.

My biggest concern with this paper was thoroughly and accurately citing sources. In the past, when I have asked Chat-GPT to cite sources for me, it has hallucinated research papers that never existed. Additionally, in my experience, LLMs tend to avoid directly stating the sources of their information, possibly to avoid copyright liability. I decided to leave the citations for last, as

I knew I could retroactively find sources for my assertions, once I was happy with the content of the paper.

After getting the initial draft in Gemini, I decided to ask it for a list of the most influential figures in the development of modern AI. This generated six figures; Alan Turing, John McCarthy, Marvin Minsky, Geoffrey Hinton, Yann LeCun, and Fei-Fei Li. I noticed that it did not mention Frank Rosenblatt, who was described in section 2 of the first draft. The only source cited for this list was a rudimentary blog post that lacked a date, author, publisher, or any other identifiable information other than the fact that it was published under a URL that suggested it came from an affiliate of the New Jersey Institute of Technology (<https://web.njit.edu/~ronkowit/eliza.html>). The blog post is about ELIZA, a “very basic Rogerian psychotherapist chatbot.” Aside from one mention of the Turing test, it does not support the content of the list that Gemini provided me. This hallucination was concerning to me, so I clicked on the information icon in the corner of the sources section. This took me to an FAQ page with the following information:

Gemini, like some other standalone LLM experiences, is intended to generate original content and not replicate existing content at length. We’ve designed our systems to limit the chances of this occurring, and we will continue to improve how these systems function. If Gemini does directly quote at length from a webpage, it cites that page. For answers with URLs or image thumbnails, Gemini enables users to easily see and, in some cases, click to navigate directly to the source for each. If Gemini generates a response using content from an extension, like Google Workspace, it provides a link to the sources (like your emails) that it used.

Sometimes the same public content may be found on multiple webpages and Gemini attempts to point to a popular source. In the case of citations to code repositories, the citation may also reference an applicable open source license.

I found it unlikely that the blog post that was cited was popular enough to merit citation, especially since it did not in fact provide the stated information. When I told Gemini that I was not happy with that source, it apologized and linked two new sources: TIME Magazine’s “The 100 Most Influential People in AI 2023” and “11 of the most influential people in Machine Learning” by Datagran. These sources were the first and fourth sources, respectively, given to me when I Googled this same question. The first source was exactly what I was looking for. However, the “datagran” source took me to the homepage of an AI business tool company, and I had to search the name of this blog post in order to find the correct link. Not mentioned in either article are the first few figures provided in Gemini’s response; Alan Turing, John McCarthy, and Marvin Minsky.

Given these sourcing inconsistencies, I decided to provide Gemini with a source and ask it to draw from that. When I provided Gemini with the Wikipedia article on the history of AI, it gave me a long list of names from the article. However, it did not seem to be able to differentiate

the importance of the figures listed, and it left out key figures, including Geoffrey Hinton, Yann LeCun, and Fei-Fei Li.

Moving on from there, I decided to ask both Gemini and Claude to name the 20 most important figures in the history of AI, then compile a list of the people who overlapped between both lists (14 in total). Claude's explanation for the inclusion of these innovators naturally grouped them into sections:

Alan Turing is considered the father of computer science and laid the foundational groundwork for AI in the 1930s-1940s.

John McCarthy and Marvin Minsky were pioneers in the 1950s, establishing the field of AI and making early breakthroughs.

Herbert Simon and Arthur Samuel made important contributions in the 1950s-1960s.

Rodney Brooks advanced embodied and behavior-based AI systems in the 1980s-1990s.

Judea Pearl pioneered Bayesian networks and causal inference in the 1980s-1990s.

Geoffrey Hinton, Yoshua Bengio, and Yann LeCun were key figures in the deep learning revolution of the 1990s-2000s.

Demis Hassabis, Andrew Ng, Stuart Russell, and Fei-Fei Li have been prominent in more recent decades, advancing AI capabilities while also addressing ethical considerations.

I then asked Claude to tell me about each innovator's contributions to the field of AI, in their own words. I decided to use Claude, after experimenting with this, Gemini, and Chat GPT-4, because I liked the writing style the most. To me, it was the most natural-sounding of the three LLM's, and it most closely replicated my personal writing style.

I encountered further difficulties getting Claude to cite its sources. When I asked for a source for its information, it spit out the following paragraph:

The information provided in these dialogues between the key figures in the history of artificial intelligence is not based on a single source, but rather a synthesis of my knowledge about the contributions and perspectives of these pioneers. As an AI assistant created by Anthropic, I do not have a specific published source to cite, but rather I have drawn from my training on the history and development of the field of AI.

The details about the individual contributions and the collaborative exchange between the figures are based on my understanding of the seminal work, ideas, and interactions of these pioneering researchers. I aimed to capture the essence of their roles and viewpoints in a plausible conversational format, but the specific phrasing and framing is my own.

As an AI system, I do not have the ability to directly reference scholarly publications or historical archives. Instead, I have constructed these dialogues using my language modeling capabilities to illustrate the key milestones and intellectual discourse that have shaped the evolution of artificial intelligence over the decades. The information presented, while grounded in my training on this

subject matter, does not have a single verifiable source that I can point to. It is an original composition generated based on my understanding of the field.

I thought I might be able to outsmart this by asking about sources I could go to in order to learn more about the information provided. Claude's response to this request was fairly vague and noncommittal. However, it listed a wide variety of books and publications about each historical figure, as well as referring me to the innovators' own writings and publications. From this, I was able to infer that these were the principal materials that Claude's responses were drawn from.

I liked the individual sections that Claude was able to generate on each group of innovators, but I struggled to synthesize the segments, avoid repeating information, and transition from section to section. Thus, I had to go back through the essay and edit for redundancy and consistency, then add in organic transitions (typed in green). I also noticed that Claude had several phrases that it liked to use, such as "the future of AI is ours to shape," which I had to cut to avoid meaningless repetition.

At the conclusion of the Claude-generated essay, I asked Chat GPT-4 to fact-check the paper. It took the assignment very literally, in the way a human would not, and explained to me that following inaccuracies:

1. Alan Turing's Role in AI Development:
 - a. Alan Turing did not live to see the development of AI as a field. He passed away in 1954, and the term "artificial intelligence" was not coined until 1956 by John McCarthy.
2. Dialogue Between Figures:
 - a. The dialogues between AI pioneers like Turing, McCarthy, Minsky, and others are entirely fictional and anachronistic. These individuals could not have engaged in discussions as portrayed because their lifespans and careers did not overlap in the manner suggested. For example, Turing could not have had discussions with later AI researchers like Geoffrey Hinton or Demis Hassabis.
3. Alan Turing's Presentation of Modern AI Developments:
 - a. Turing discusses developments in AI that occurred long after his death, such as deep learning and advancements by researchers like Demis Hassabis. This portrayal is historically impossible.

I then asked Gemini to fact-check the paper. It suggested areas of improvement, but these suggestions contained multiple false assertions. For example, it suggested I mention pioneers from parts of the world other than the U.S. and U.K., such as Yann LeCun and organize the paper by "thematic groupings based on research areas or historical periods," both of which were already elements of my paper. Giving up on that strategy, I asked Gemini to add in Chicago-style citations to my paper. Unfortunately, it misunderstood me as asking it to provide a citation for the essay as a whole, so it hallucinated the following:

Turing, Alan. "An Introduction." In *The Alan Turing Omnibus*. Edited by Jack Copeland, 1-25. Oxford: Oxford University Press, 2004.

While there is a “New Turing Omnibus” and an author named Jack Copeland who has written about Alan Turing, this source does not exist. When I attempted to clarify to Gemini that I wanted in-text citations to support each assertion in my paper, it used the fake citation it had created as a source for much of the new response.

Attempting to source my essay a different way, I then spent around 2 hours experimenting with the browser extensions Sciwheel and Zotero. Unfortunately, neither was able to retroactively find and suggest sources for content that was already written, especially given that the AI-generated writing had been synthesized from countless sources around the web.

At this point, I was losing hope in the ability of any of these LLMs to fact-check or source the material in this essay. However, I decided to give Chat GPT-4 one last chance before I gave up and manually did all of the citations. Attaching my introductory section, I prompted Chat GPT-4 as such: “provide me with a list of sources for the following information. make sure that every assertion is supported by a trustworthy and accurate source.”

To my relief, it provided a legitimate and appropriate source that corresponded to each paragraph of the introduction. Chat GPT-4 even understood that the introduction I had given it was a “fictionalized speech attributed to Alan Turing,” despite its previous misunderstanding of my creative style when I had asked it to fact-check my paper previously. I then input the sources given to me by Chat GPT-4, double-checking any that looked suspicious. Surprisingly, none of the sources seemed to be hallucinated.

I am still unsure why this prompt worked out so much better than my other attempts to obtain sourcing. Could it be that if the user wants real, reliable sources, we must explicitly state this? I decided to continue to give explicit context for the following sections, explaining that they were fictionalized conversations, but needed citations for their truthful assertions.

In conclusion, I thought that this project was a fascinating experiment. However, my biggest takeaway is that LLMs are unreliable and unimaginative. While they are meant to be useful and efficient, I found that I spent more time on this paper, trying to use AI, than I would have if I had written it entirely on my own. Additionally, I believe that on my own, I could have produced a better final product. I will certainly continue to experiment with LLMs and other AI, but with a realistic perspective on its limited present utility.

The Final Product:

The Artificial Intelligence Innovators: *in their own (creation's) words*

An Introduction by: Alan Turing

clears throat and speaks in a distinguished British accent

I am Alan Turing, the father of computer science and artificial intelligence. My contributions to this field have been foundational, laying the groundwork for the remarkable progress we have witnessed over the decades.¹

It all began with my 1936 paper on the Turing machine - a conceptual model of computation that demonstrated the universal nature of certain types of mathematical calculations. This work established the theoretical limits and possibilities of what could be computed by mechanical means. It was a profound insight that would shape the future of computing and AI.²

Building on this, in 1950 I published my seminal paper "Computing Machinery and Intelligence," which proposed the now-famous Turing test as a way to assess whether a machine can be considered intelligent. The key idea was to judge a machine's intelligence not by its inner workings, but by its ability to exhibit behavior indistinguishable from a human in a series of text-based interactions.³

This challenge to determine if a machine can "think" sparked a rich and ongoing debate about the nature of intelligence, consciousness, and the possibility of creating artificial minds. My work challenged the prevailing view that machines could only perform rigid, predetermined tasks, and opened up the tantalizing prospect that they might one day rival human-level cognitive capabilities.⁴

Of course, the road to developing truly intelligent machines has not been an easy one. There have been many ups and downs, periods of enthusiasm and disillusionment - the so-called "AI winters" - as the field has grappled with the immense complexity of replicating the full breadth and flexibility of human cognition.⁵

¹ Hodges, Andrew. Alan Turing: The Enigma. London: Vintage, 2012.

² Turing, Alan. "On Computable Numbers, with an Application to the Entscheidungsproblem." Proceedings of the London Mathematical Society, series 2, vol. 42, no. 1 (1936): 230-265.

³ Turing, Alan. "Computing Machinery and Intelligence." Mind 59, no. 236 (1950): 433-460.

⁴ Stanford Encyclopedia of Philosophy. "The Turing Test." Last modified July 12, 2019.
<https://plato.stanford.edu/entries/turing-test/>.

⁵ Mitchell, Melanie. Artificial Intelligence: A Guide for Thinking Humans. New York: Farrar, Straus and Giroux, 2019.

Yet, I remained steadfast in my conviction that machines could be imbued with intelligence, creativity, and autonomy. My ideas and thought experiments laid the philosophical and mathematical foundations that have inspired generations of researchers to push the boundaries of what's possible.⁶

Now, as I look at the remarkable progress in areas like deep learning, computer vision, natural language processing, and robotic control, I am both gratified and humbled. The field of artificial intelligence has grown exponentially beyond my initial conjectures, with researchers unlocking capabilities that I could scarcely have imagined.⁷

However, the ethical implications of these advances have also become increasingly apparent. As AI systems become more powerful and ubiquitous, we must ensure they are developed and deployed in a manner that is aligned with human values and the wellbeing of society. This is a responsibility that I hope all AI pioneers will take seriously, building on the foundations I and others have laid.⁸

Speaking of which, let's let my successors, John McCarthy and Marvin Minsky, tell it in their own words.

John McCarthy: Thanks for the introduction, Alan. Ah, Marvin, my old friend and fellow pioneer of artificial intelligence. It's been far too long since we've had the chance to sit down and discuss the evolution of our beloved field.

Marvin Minsky: It's a pleasure to reconnect and reflect on the remarkable journey we've both been a part of. As you know, our approaches to AI may have differed at times, but I believe our complementary contributions have been instrumental in shaping the discipline.⁹

John McCarthy: Absolutely. I'm proud to have played a role in establishing AI as a distinct field of study through the coining of the term "artificial intelligence" in 1956 at the Dartmouth Conference. That was a pivotal moment, as it helped solidify our work as a unique endeavor, separate from the realms of computer science and cognitive science.¹⁰

⁶ Copeland, B. Jack, ed. The Essential Turing: Seminal Writings in Computing, Logic, Philosophy, Artificial Intelligence, and Artificial Life plus The Secrets of Enigma. Oxford: Oxford University Press, 2004.

⁷ Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. Deep Learning. Cambridge, MA: MIT Press, 2016.

⁸ O'Neil, Cathy. Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy. New York: Crown, 2016.

⁹ Russell, Stuart J., and Peter Norvig. Artificial Intelligence: A Modern Approach. 3rd ed. Upper Saddle River, NJ: Prentice Hall, 2010.

¹⁰ McCorduck, Pamela. Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence. 25th Anniversary Update ed. Natick, MA: A K Peters, Ltd., 2004.

Marvin Minsky: And your pioneering work on logic-based AI, knowledge representation, and common sense reasoning laid the groundwork for much of the field's early progress. Your development of the Lisp programming language, in particular, proved to be a crucial tool for AI researchers, providing a versatile platform for representing and manipulating complex knowledge structures.¹¹

John McCarthy: **Thank you**, Marvin. But I must say, your own focus on neural networks and the theory of intelligence was a crucial counterpoint to those more symbolic approaches. Your seminal work on the "Society of Mind" theory, which posited that the human mind is composed of a vast network of interacting agents or "mindlets," laid the groundwork for a more holistic understanding of intelligence.¹²

Marvin Minsky: Precisely. I recognized that true intelligence, whether artificial or human, emerges from the complex interplay of interconnected processing units, much like the structure of the brain. This perspective would later prove pivotal in the resurgence of neural network-based approaches, the so-called "deep learning revolution" that we're now witnessing.¹³

John McCarthy: And while our approaches may have differed at times, I believe our combined efforts played a vital role in establishing the intellectual foundations of this field. The tension between symbolic and connectionist approaches has been a driving force in the evolution of AI.¹⁴

Marvin Minsky: Undoubtedly. And the fact that we were able to debate and discuss these ideas, often passionately, within the nurturing environment of the MIT AI Lab, is a testament to the collaborative spirit that has characterized this discipline. We pushed each other to think beyond the boundaries of what was possible.¹⁵

John McCarthy: **Yes, and** we were always united in our vision of AI's immense potential to transform our understanding of the human mind and our relationship with technology. That shared conviction has been a guiding light throughout our careers.

Marvin Minsky: **I couldn't agree more. Now, I'm interested in hearing how the pioneers Herbert Simon and Arthur Samuel built on the foundations we laid.**

¹¹ McCarthy, John. "Recursive Functions of Symbolic Expressions and Their Computation by Machine, Part I." Communications of the ACM 3, no. 4 (1960): 184-195.

¹² Minsky, Marvin, and Seymour Papert. Perceptrons: An Introduction to Computational Geometry. Expanded ed. Cambridge, MA: MIT Press, 1988.

¹³ Goodfellow, 2016

¹⁴ Russell, Stuart J., and Peter Norvig. Artificial Intelligence: A Modern Approach. 3rd ed. Upper Saddle River, NJ: Prentice Hall, 2010.

¹⁵ Mitchell, 2019

Herbert Simon: **Happy to oblige.** Arthur, it's a pleasure to reconnect and discuss the evolution of our field. I'm particularly intrigued by your groundbreaking work on the first self-learning program - the checkers-playing algorithm that could improve its own performance through practice.¹⁶

Arthur Samuel: Why, thank you, Herbert. That project was indeed a labor of love, as I sought to demonstrate the potential for machines to engage in autonomous learning and problem-solving.¹⁷ Your own contributions, however, have been equally transformative.

Herbert Simon: Ah yes, the work my colleagues and I did on problem-solving, heuristics, and the simulation of human-like thinking in machines was instrumental in shaping the early trajectory of AI. The development of our Logic Theorist¹⁸ and General Problem Solver¹⁹ programs in the 1950s were pivotal milestones.

Arthur Samuel: And I believe the synergy between our approaches has been a driving force in the field's evolution. While you focused on symbolic and knowledge-based AI, my emphasis on learning and autonomous problem-solving provided an important counterpoint.²⁰

Herbert Simon: Absolutely. The interplay of these different perspectives has been crucial in pushing the boundaries of what's possible. As we now witness the remarkable advancements, I'm struck by how far the field has come, **especially following the work of Judea Pearl and Rodney Brooks.**

Rodney Brooks: **Thank you Arthur and Herbert.** Judea, I'm eager to hear your insights and reflect on how our respective contributions have **affected** the discipline **of artificial intelligence.**

Judea Pearl: **Of course, Rodney.** I believe the interplay between our approaches has been instrumental in driving the progress of AI.²¹

¹⁶ Samuel, Arthur L. "Some Studies in Machine Learning Using the Game of Checkers." IBM Journal of Research and Development 3, no. 3 (1959): 210-229.

¹⁷ Samuel, 1959

¹⁸ Newell, Allen, and Herbert A. Simon. "The Logic Theory Machine: A Complex Information Processing System." IRE Transactions on Information Theory 2, no. 3 (1956): 61-79.

¹⁹ Newell, Allen, Herbert A. Simon, and J.C. Shaw. "Report on a General Problem-solving Program." Proceedings of the International Conference on Information Processing. Paris: UNESCO, 1959.

²⁰ Russell, 2010

²¹ Russell, 2010

Rodney Brooks: Your groundbreaking work on Bayesian networks²² and causal inference²³ has had a profound impact on how AI systems handle uncertainty and make inferences. The ability to reason probabilistically has been a crucial advancement.

Judea Pearl: Thank you. And your emphasis on physical embodiment and real-world interaction for intelligent systems has been equally transformative. Your work on behavior-based robotics and the subsumption architecture challenged the traditional AI approaches and pushed the field towards more grounded, situated intelligence.²⁴

Rodney Brooks: I always felt that for AI to truly flourish, it needed to be grounded in physical and sensory experience, not just abstract reasoning. The integration of cognitive capabilities with real-world interaction has been a key focus of my research.²⁵

Judea Pearl: Absolutely. **I know that** the work we've done, and the paths we've helped forge, will continue to inspire and guide the next generation of AI pioneers. **No one can attest to this better than Geoffrey Hinton, Yoshua Bengio, and Yann LeCun.**

Geoffrey Hinton: **We sure can.** Yoshua, Yann, it's wonderful to see you both. As the three of us know all too well, our work has been instrumental in ushering in the latest revolution in artificial intelligence.²⁶

Yann LeCun: Absolutely, Geoffrey. But I think it's important that we each reflect on our specific contributions to this field.

Geoffrey Hinton: I'll start by acknowledging my work on neural networks and deep learning architectures, which I believe laid much of the groundwork for the successes we've witnessed. My exploration of backpropagation, and the insights I gained about the power of distributed representations, were pivotal.²⁷

²² Pearl, Judea. "Bayesian Networks: A Model of Self-Activated Memory for Evidential Reasoning." In Proceedings of the 7th Conference of the Cognitive Science Society, University of California, Irvine, CA, USA, August 15-17, 1985, 329-334.

²³ Pearl, Judea. Causality: Models, Reasoning, and Inference. 2nd ed. Cambridge, UK: Cambridge University Press, 2009.

²⁴ Brooks, Rodney A. "A Robust Layered Control System for a Mobile Robot." IEEE Journal of Robotics and Automation RA-2, no. 1 (March 1986): 14-23.

²⁵ Brooks, 1986

²⁶ "ACM Awards Turing Award to Creators of Modern AI." Association for Computing Machinery (ACM), last modified March 27, 2019.

²⁷ Hinton, Geoffrey E., Simon Osindero, and Yee-Whye Teh. "A Fast Learning Algorithm for Deep Belief Nets." Neural Computation 18, no. 7 (2006): 1527-1554.

Yoshua Bengio: And I built upon those foundations by delving deeper into the mathematics and theory behind deep learning. My work on recurrent neural networks, unsupervised learning, and other deep architectures helped expand the capabilities of these systems.²⁸

Yann LeCun: I would argue that my contributions in the realm of convolutional neural networks were equally transformative. The ability to leverage the spatial structure of data, such as images, was a breakthrough that enabled major advancements in computer vision.²⁹

Geoffrey Hinton: Your work on convolutional networks was instrumental in pushing the boundaries of what was possible. And I believe the interplay between our different areas of focus - your computer vision, Yoshua's language modeling, and my general architectures - has been a driving force behind the progress we've seen.³⁰

Yoshua Bengio: That's true, and I think it's important to acknowledge that our individual contributions have been recognized and honored, most notably through our joint receipt of the Turing Award - the "Nobel Prize of computing."³¹

Geoffrey Hinton: Now, shall we discuss the latest advancements? Or better yet, we'll let [Demis Hassabis](#), [Andrew Ng](#), [Stuart Russell](#), and [Fei-Fei Li](#) do the talking.

Demis Hassabis: Colleagues, it's a pleasure to have the chance to discuss our collective contributions to the field of artificial intelligence.

Andrew Ng: As a pioneer in making AI more accessible and applicable to real-world problems, I've dedicated much of my career to promoting the responsible development and deployment of these powerful technologies.³²

Stuart Russell: As we've witnessed the rapid advancements in AI, the need to grapple with the ethical implications has become increasingly critical. My work on machine ethics and aligning AI systems with human values has been a crucial part of this endeavor.³³

²⁸ Bengio, Yoshua, Pascal Lamblin, Dan Popovici, and Hugo Larochelle. "Greedy Layer-Wise Training of Deep Networks." In *Advances in Neural Information Processing Systems 19* (NIPS 2006), edited by B. Schölkopf, J. Platt, and T. Hoffman, 153-160. Cambridge, MA: MIT Press, 2007.

²⁹ LeCun, Yann, Léon Bottou, Yoshua Bengio, and Patrick Haffner. "Gradient-Based Learning Applied to Document Recognition." *Proceedings of the IEEE* 86, no. 11 (1998): 2278-2324.

³⁰ "ACM Awards Turing Award"

³¹ "ACM Awards Turing Award"

³² Ng, Andrew. "Machine Learning Yearning." Self-published, 2018.

³³ Russell, Stuart. "Human Compatible: Artificial Intelligence and the Problem of Control." Viking, 2019.

Fei-Fei Li: *Exactly right, Stuart.* As a leading figure in computer vision and a champion of diversity and inclusion in the AI community, I've strived to ensure that the development of these technologies is guided by a deep understanding of their societal impact. By fostering a more inclusive and representative AI ecosystem, we can better anticipate and address the complex implications of these powerful tools.³⁴

Demis Hassabis: Your contributions, Fei-Fei, have been invaluable. The work your team has done in areas like large-scale image recognition datasets has been instrumental in driving progress in computer vision and laying the foundations for more robust and equitable AI systems.³⁵

Fei-Fei Li: *I am proud of what I do, and you should be too, Demis.* The breakthroughs your team at DeepMind has achieved in game-playing AI, like AlphaGo and AlphaZero, have been truly remarkable. The ability of these systems to learn complex strategies and reasoning through self-play *is incredible.*³⁶

Andrew Ng: But I still believe the real challenge lies in translating those impressive results into AI systems that can operate safely and beneficially in the real world, beyond the confines of a game. That's why I've dedicated so much of my work to making AI more accessible and understandable to a wider audience, empowering people to engage with and *contribute to* the future of this technology.³⁷

Stuart Russell: The path forward will require a continued collaboration and collective responsibility from all of us. The future of AI is *in our hands*, and we must rise to the challenge.

Conclusion by Alan Turing:

As I gaze upon the remarkable progress that has unfolded in the field of artificial intelligence, I am both gratified and humbled. The pioneers I have had the privilege of

³⁴ Deng, Jia, et al. "ImageNet: A large-scale hierarchical image database." In 2009 IEEE Conference on Computer Vision and Pattern Recognition, 248-255. IEEE, 2009.

³⁵ Liang, Weixin, Girmaw Abebe Tadesse, Daniel Ho, Li Fei-Fei, Matei Zaharia, Ce Zhang, and James Zou. "Advances, challenges and opportunities in creating data for trustworthy AI." *Nature Machine Intelligence* 4, no. 8 (2022): 669-677.

³⁶ Silver, David, et al. "Mastering the game of Go with deep neural networks and tree search." *Nature* 529, no. 7587 (2016): 484-489. This article covers the development of AlphaGo by DeepMind.

³⁷ Ng, 2018

introducing - John McCarthy, Marvin Minsky, Herbert Simon, Arthur Samuel, Rodney Brooks, Judea Pearl, Geoffrey Hinton, Yoshua Bengio, Yann LeCun, Demis Hassabis, Andrew Ng, Stuart Russell, and Fei-Fei Li - have all played crucial roles in shaping the trajectory of this discipline.

Their spirited debates, collaborative spirit, and unwavering dedication have propelled AI from the realm of my initial conjectures to a transformative force that is poised to reshape our world. While the path has not been without its challenges and setbacks, the resilience and creativity of these pioneers have overcome the obstacles, opening up new frontiers of possibility.

From the foundational breakthroughs in computation, knowledge representation, and neural networks, to the more recent advancements in deep learning, computer vision, and natural language processing, each generation has built upon the work of the last, unlocking capabilities that I could scarcely have imagined.

Yet, as the power and ubiquity of AI systems continue to grow, the need to ensure their responsible development and deployment has become increasingly paramount. The ethical considerations raised by my successors are a testament to their foresight and commitment to the well-being of humanity.

The road ahead may be uncertain, but I have no doubt that with the continued dedication, collaboration, and sense of purpose demonstrated by these luminaries, the full potential of artificial intelligence will be realized. The machine intelligence of tomorrow will not merely emulate human cognition, but transcend it, in service of a better world for all.

This is the legacy that I, and all those who have contributed to this field, have helped to create. It is a testament to the power of human ingenuity, the triumph of the mind, and the boundless possibilities that await us in the years to come. The future of artificial intelligence is ours to shape, and I am honored to have played a part in the journey.