From ELIZA to ChatGPT*

History of Human-Computer Conversation

V. Rajaraman

In 1950, Turing suggested a test to determine whether a computer can exhibit intelligent behaviour indistinguishable from that of a human. An aspect of the test was an imaginary conversation between a human and a computer. Since then, computer scientists have designed software to let humans interact with a computer in natural language. The first such program was ELIZA, designed in 1965. After ELIZA, numerous programs (christened Chatbots) have been written to enable humans to converse with computers in natural language. In this article, I describe several of them, including those that orally answer oral questions. I describe the latest program called ChatGPT, developed in late 2022, which has astounded everyone by its ability to answer questions posed in a natural language in any field and converse fluently with humans.



V. Rajaraman is at the Indian Institute of Science, Bengaluru. Several generations of scientists and engineers in India have learnt computer science using his lucidly written textbooks on programming and computer fundamentals. His current research interests are parallel computing and history of computing.

Keywords
Human-computer conversation, chatbot, ChatGPT, voice assistant.

Prologue

Alan Mathison Turing was born on 23 June 1912 in the UK. He obtained a Bachelor's degree from Cambridge University in 1934 and a PhD degree from Princeton University in 1938. He made three landmark contributions during his scientific career. In 1936, he published a paper titled "On Computable Numbers", in which he proposed a model of a universal computing machine, popularly known as "The Turing Machine", and proved that there exist numbers that cannot be computed by this machine. This demonstrated the theoretical limitations of computers. In 1940, he found a method of breaking the German Enigma code [1] and designed a machine to do it. This code-breaking contributed to the Allied

*Vol.28, No.6, DOI: https://doi.org/10.1007/s12045-023-1620-6



Figure 1. Alan Turing (Photo courtesy: Wikipedia).

victory in the second World War. The last was the paper "Computing Machinery and Intelligence" [2], which appeared in *Mind*, a little-known philosophical magazine, in October 1950. He starts the paper with the sentence, *"I propose to consider the question "Can machines think?" This should begin with the definitions of the meaning of the terms "machines" and "think"."* He argues that the common usage of these words is vague, and the question cannot be answered as posed. He replaces the question by postulating an imitation game. The game he proposed is between a computer that imitates a man, a woman, and an interrogator. A wall separates the interrogator from the computer and the woman. The interrogator types his questions using a teletypewriter, and it is routed to either the computer or the woman randomly. The answers are typed and relayed back to the interrogator. If after a reasonably long question-answering session, the interrogator is unable to decide which answers are from the computer and which are from the woman, then one may conclude that the woman and the computer are indistinguishable. In other words, one may conclude that the computer is intelligent. Turing believed that in fifty years computers will pass his test.

If after a reasonably long question-answering session, the interrogator is unable to decide which answers are from the computer and which are from the woman, then one may conclude that the woman and the computer are indistinguishable.

The Turing test has been extensively debated, and the current consensus is that it is an important milestone in the development of artificial intelligence (AI) but does not fully capture the nuances of human intelligence. The idea of a conversation between a human and a computer, however, captured the imagination of re-



Figure 2. Joseph Weizenbaum (Photo courtesy: Wikipedia).

searchers in AI and led to the development of several computer-based natural language question-answering systems.

ELIZA Program [3]

Joseph Weizenbaum, a professor at the Massachusetts Institute of Technology, wrote one of the earliest programs in 1966 to mimic a natural language conversation between a human and a computer. He designed a list-processing language called SLIP and embedded it in MAD (Michigan Algorithmic Decoder). He wrote a program using this language on the time-shared IBM 7094 computer at MIT. Weizenbaum called this program ELIZA after the flower seller Eliza Doolittle in Bernard Shaw's play *Pygmalion* (later adapted as the popular musical *My Fair Lady*) in which Professor Henry Higgins takes a bet that he will teach flower seller Eliza Doolittle to speak proper English in an upper-class accent. ELIZA program learns to converse gradually through several interactions, just like Shaw's character Eliza Doolittle. ELIZA was intended to mimic a non-directive psychotherapist—an approach in which a therapist tries to let a patient figure out his or her problem by framing questions based on a patient's response. Weizenbaum's primary aim was to program a computer to respond in natural language based on analysing a natural language input to it. A typical conversation with ELIZA is given in *Box* 1.

ELIZA was intended to mimic a non-directive psychotherapist—an approach in which a therapist tries to let a patient figure out his or her problem by framing questions based on a patient's response.

Box 1. Conversation with Eliza[@]

Patient: Men are all alike.
Computer: In what way?
Patient: They're always bugging us about something or the other.
Computer: Can you think of a specific example?
Patient: Well, my boyfriend made me come here.
Computer: Your boyfriend made you come here?
Patient: He says I'm depressed much of the time.
Computer: I am sorry to hear you are depressed.
Patient: It is true I am unhappy.
Computer: Do you think coming here will help you not to be unhappy?
Patient: I need some help, that seems certain.
Computer: What would mean to you if you got some help?
Patient: Perhaps I could learn to get along with my mother.
Computer: Tell me more about your family.
Patient: My mother takes care of me.
Computer: Who else in your family takes care of you?
Patient: My father.
Computer: Your father?

———————- Conversation continues ———————

[@] V. Rajaraman, Turing test and after, *Resonance*, Vol.3, No.8, pp.50–59, August 1997.

ELIZA used the fact that a small number of keywords are used in most conversations. ELIZA stores sentence templates whenever a keyword appears. For example, a reference to a member of one's family, such as "my mother" or "my sister", results in the computer printing the response "Tell me more about your family". Often, a patient's statement is reversed as in the computer's response "Your boyfriend made you come here?" to a patient's statement, "Well, my boyfriend made me come here", in the above example. Several such simple rules to transform sentences made ELIZA seemingly an intelligent conversationalist and fooled many ordinary people into believing that another person was answering the questions. Weizenbaum wrote a book [4] based on his observa-

tions of people who used ELIZA. He was startled to see how people using ELIZA became emotionally involved with it. Even his secretary, who had seen him develop ELIZA, asked him to leave the room when she conversed with it!

Human-Computer Dialog Systems [5]

Weizenbaum's research kindled the interest of many researchers and they wrote natural language processing programs to mimic human-computer dialogs. Kenneth Colby of Stanford University wrote a program PARRY in 1972 that simulated a conversation of a paranoid schizophrenic with a doctor. A transcript of PARRY's responses and the responses from a human patient was given to five psychiatrists and they were asked to separate PARRY's responses from those of the real patient. The psychiatrists were able to separate the two responses only 48% of the time—almost like guessing. A demonstration of a conversation between ELIZA and PARRY was given at the International Computer Communication Conference (ICCC 72), which is interesting to read [6].

Kenneth Colby of Stanford University wrote a program PARRY in 1972 that simulated a conversation of a paranoid schizophrenic with a doctor.

JABBERWACKY, a program developed by the British programmer Rollo Carpenter in 1988, was the next important development in dialog systems. By then PCs had arrived and it was easier and cheaper to develop interactive dialog systems. It was the first system that combined natural language processing with machine learning. JABBERWACKY analysed a user's input, identified keywords, and created a database of input-output responses over a series of interactions. This was used by a machine learning system to improve its responses. JABBERWACKY became popular and is available on the Internet even now. Michael Mauldin coined the term chatterbot in 1994 for human-computer dialog systems. It was later abbreviated as *chatbot*. He wrote a dialog system named JULIA which ran on the Internet Relay Chat (IRC) system. JULIA's responses were humorous or clever and made JULIA popular. Chatbot entered dictionaries around 2000. In the Oxford Dictionary, it is defined as *A computer program designed to simulate conversation with human users, especially*

JABBERWACKY analysed a user's input, identified keywords, and created a database of input-output responses over a series of interactions. This was used by a machine learning system to improve its responses.

over the Internet.

In 1991, Creative Labs released a chatbot named DR. SBAITSO (Sound Blaster AI Text to Speech Operator) which mimicked a psychologist and used a very simple text-based dialog system. It was capable of giving spoken answers but the system which was released gave answers as text. The use of speech-based dialogs became popular later in interactive systems designed to answer customer queries. Richard Wallace of Lehigh University developed the next notable system named A.L.I.C.E. (Artificial Linguistic Internet Computer Entity) in 1995 for use on the Internet. It was a rule-based system that used a natural language processor with a database of 41000 keywords. One of its interesting features was its use of linguistic deflection in answering queries. For example, if you asked the question: "Why are apples red?", it responded, "Does everything have to have an explanation?" ALICE was rewritten and improved, and the latest version uses AI markup language for specifying conversation rules. It was a remarkable system in its time and won the Loebner Prize (a prize given to a most human-like dialog system) in 2000, 2001, and 2004. The perfume retailer, CyberScent used the technology of ALICE for online customer support. Spike Jonze, a movie writer and director, cites ALICE as the inspiration for his highly acclaimed movie *Her* in which the hero falls in love with a chatbot named Samantha.

One of A.L.I.C.E.'s interesting features was its use of linguistic deflection in answering queries. For example, if you asked the question: "Why are apples red?", it responded, "Does everything have to have an explanation?"

Voice Assistants

Bell Labs and IBM, along with academic researchers in universities, pioneered research on speech recognition and synthesis beginning in the 1960s. Very powerful integrated circuits were available by 2001, which led to the commercial availability of powerful speech recognition and synthesis systems. Chatbot developers used these extensively.

ActiveBuddy Inc. designed one of the first voice-based intelligent chatbot named *SmarterChild* in 2001. It had access to a large database and used this to answer queries such as the results

SmarterChild had access to a large database and used this to answer queries such as the results of football matches, stock prices, weather, arrival and departure time of airlines, etc.

of football matches, stock prices, weather, arrival and departure time of airlines, etc. It was available on America On-Line (AOL) Instant Messenger and Windows Live Messenger. It inspired the design of a large number of interactive voice-based customer relations systems.

Siri was one of the first voice assistants integrated with smartphones and devices using Apple's iOS. Apple Computers released it in 2010. SRI International collaborated with the Federal Polytechnic of Lausanne and Nuance Communications and developed the technology for voice recognition and response which was used by Siri. The system was trained by an advanced neural network developed by Nuance. Siri is quite sophisticated and can obey commands such as: Hey Siri call my mom, remind me tomorrow to send greetings to Tom, and what will be the weather like tomorrow? It answers questions such as "What is the height of Mount Everest?" using the information available on the Internet. Siri finds it difficult to recognize foreign accents. It stores a huge number of conversations (anonymously) ostensibly to train the system worrying many privacy advocates.

Siri was one of the first voice assistants integrated with smartphones and devices using Apple's iOS. Apple Computers released it in 2010.

IBM developed a notable voice assistant named *Watson* in 2011. IBM Watson beat the champions Brad Rutter and Ken Jennings in a television game called "Jeopardy!" on 16 February 2011. In this game, a player guesses the question for which the answer is given. Microsoft released the voice assistant *Cortana* in 2014. Cortana worked with Windows OS, Android, and iOS. It was integrated with the Bing search engine. It answered questions, scheduled appointments, and could converse in several languages, including Chinese. It was discontinued in 2021.

IBM Watson beat the champions Brad Rutter and Ken Jennings in a television game called "Jeopardy!" on 16 February 2011. In this game, a player guesses the question for which the answer is given.

Amazon released an intelligent voice-based chatbot named *Alexa* in 2014. Amazon bought a Polish speech synthesiser company Nona in 2013 and designed Alexa using its technology. Alexa uses AI and natural language processing for speech recognition and synthesis. Amazon built a smart speaker (a combination of a microphone and a speaker which could be connected by WiFi or Bluetooth to a computer) named *Echo*, which is used as the main device for voice-based interaction with a smartphone/computer

Alexa is woken up using the command (also called a *hot word*) "*Hey Alexa*". It could be used to answer real-time queries on the weather, stock prices, and sports scores. It can also be commanded to play music, set alarms, stream podcasts, read audiobooks, and control smart devices.

running on Android, iOS, Linux, or Windows. Alexa is woken up using the command (also called a *hot word*) "*Hey Alexa*". It could be used to answer real-time queries on the weather, stock prices, and sports scores. It can also be commanded to play music, set alarms, stream podcasts, read audiobooks, and control smart devices. It understands several languages, including English, French, German, Italian, Spanish, Portuguese, Japanese, Hindi, and Arabic. It is reported that Amazon employs around 10,000 in the Alexa group and it cost approximately 200 million USD to develop the system. As of 2019, 100 million Alexa-enabled smart speakers have been sold.

It is reported that Amazon routinely records user interactions (ostensibly to train the system) without users' consent and uses it for marketing. This is vehemently denied by Amazon.

Google released *Google Assistant*, a chatbot similar to Alexa, in 2016 for Android-based smartphones. It uses AI extensively and accepts both voice and text commands. It is also adapted to work with several other operating systems. It answers questions, plays games, schedules appointments, adjusts device settings, and carries out many other daily tasks. In a recent update, Google announced that it identifies objects using the phone's camera, supports e-commerce, and can be used to send money. Google Assistant may also be used with a smart speaker, just like Alexa. There are concerns about the privacy of Google Assistant users.

Starting early 2016, customer complaints and queries are being answered by chatbots. Human interaction has almost disappeared, to the chagrin of several customers, including me.

Chatbots developed by companies for customer interaction proliferated with advances in AI, natural language processing, and the advent of social media platforms. Starting early 2016, customer complaints and queries are being answered by chatbots. Human interaction has almost disappeared, to the chagrin of several customers, including me. Interactions with chatbots can often be exasperating. It is estimated that over a million chatbots covering a wide area including marketing, health care, entertainment, education, and many other fields are now active. The global chatbot market was valued at USD 2.9 billion in 2020.

ChatGPT

The availability of almost inexhaustible computing power from clouds such as Microsoft's Azure, and huge amounts of multimedia data from the Internet, rekindled research in deep neural networks in the first decade of this century. Rapid advances were made in natural language recognition and learning systems. OpenAI (funded by several companies, including several billion dollars from Microsoft) started working with a group of gifted AI researchers in 2015. It released a revolutionary text-based chatbot named ChatGPT 3.5 to the public in November 2022. It caused a furore in the world of AI and astonished its users with its prowess. It uses a large language model and a deep neural network to generate text. It has been trained on around 570 GB of diverse text data including web pages and books (data available until 2021) and provides coherent human-like natural language answers to any arbitrary question one may ask. (I got most of the data about ChatGPT by asking it). It learns based on feedback from users. To illustrate its power, I reproduce below some of my interactions with ChatGPT 3.5 during April 2023.

The availability of almost inexhaustible computing power from clouds such as Microsoft's Azure, and huge amounts of multimedia data from the Internet, rekindled research in deep neural networks in the first decade of this century.

Query: *How many languages can be used with ChatGPT?*

Answer: ChatGPT can theoretically support any language for which there is sufficient training data available. However, the degree to which it can understand and generate text in a given language depends on the quality and quantity of the data available for that language. Currently, the most widely used and supported languages for ChatGPT are English, Spanish, French, German, Italian, Portuguese, Dutch, Japanese, Chinese, and Korean. However, with the availability of more data in other languages, it is possible for ChatGPT to support additional languages in the future.

ChatGPT can theoretically support any language for which there is sufficient training data available. However, the degree to which it can understand and generate text in a given language depends on the quality and quantity of the data available for that language.

Query: *What are the shortcomings of ChatGPT?*

1. Lack of real-world knowledge: ChatGPT relies on large amounts of training data to generate its responses, which means that it has a limited understanding of the real world outside of the context of

its training data. This can lead to errors and misunderstandings when responding to questions or comments that fall outside of its training data.

ChatGPT is essentially a sophisticated pattern recognition system, and it does not have the ability to reason or think critically in the way that humans do.

2. Inability to reason or think critically: ChatGPT is essentially a sophisticated pattern recognition system, and it does not have the ability to reason or think critically in the way that humans do. This means that it may struggle with certain types of complex or nuanced conversations, or may give answers that are technically correct but lack context or nuance.

3. Potential for bias or inappropriate responses: Because ChatGPT relies on large amounts of training data, it can potentially be influenced by biases or inappropriate content in that data. For example, if the training data contains sexist or racist language or ideas, ChatGPT may unintentionally reproduce those biases in its responses.

4. Difficulty with long-term context: While ChatGPT is capable of understanding context within the scope of a single conversation, it can struggle to maintain long-term context across multiple conversations or interactions with the same user. This can make it difficult for ChatGPT to provide personalized or consistent responses over time.

5. Lack of emotional intelligence: While ChatGPT is capable of recognizing and responding to certain emotional cues in user input, it lacks the ability to truly understand or empathize with human emotions in the way that humans can. This can make it difficult for ChatGPT to provide appropriate emotional support or guidance in certain situations.

Overall, while ChatGPT is an impressive achievement in the field of natural language processing and conversational AI, it is important to recognize its limitations and shortcomings in order to use it effectively and responsibly.

Overall, while ChatGPT is an impressive achievement in the field of natural language processing and conversational AI, it is important to recognize its limitations and shortcomings in order to use it effectively and responsibly.

Query: *Is it correct to say that ChatGPT has passed the Turing test?*

Answer: No, it is not correct to say that ChatGPT has passed the Turing test. The Turing test is a subjective and controversial measure of artificial intelligence that involves convincing a human judge that a machine is indistinguishable from a human in a conversation. While ChatGPT is a highly advanced natural language processing system that is capable of generating responses that are sometimes difficult to distinguish from those of a human, it has not been formally tested in the context of a Turing test. Moreover, passing the Turing test is a difficult and highly debated accomplishment in the field of artificial intelligence, and there is ongoing debate among experts as to whether the Turing test is a valid or useful measure of true AI. While ChatGPT is an impressive achievement in the field of natural language processing and conversational AI, it is not accurate to say that it has definitively passed the Turing test.

Passing the Turing test is a difficult and highly debated accomplishment in the field of artificial intelligence, and there is ongoing debate among experts as to whether the Turing test is a valid or useful measure of true AI.

I have explained how ChatGPT works in *Box* 2. Besides answering questions, ChatGPT could write poetry, write programs given a description of requirements, and write essays when a topic is given. ChatGPT has been a runaway success since its launch. It had 100 million users by January 2023 and 1.6 billion visits by March 2023. It has been upgraded to ChatGPT 4, a larger language model, and a paid service (see *Box* 3). It has been reported that ChatGPT 4 was placed in the top 90th percentile in the International Math and Biology Olympiads and the US Bar examination.

Box 2. How Does ChatGPT Work?

The letters GPT in ChatGPT stand for Generative Pretrained Transformer. Generative indicates that it Generates words—predicting what the next word should be based on a sequence of previous words in a sentence. Pre-trained because the neural network it uses is trained repeatedly on massive amounts of data until its performance is satisfactory as judged by the trainers. The word Transformer indicates the type of neural network used by it. Transformer is a new type of deep neural network* that was developed by Vaswani and his colleagues at Google in 2017 called a Large Language Model (LLM). It was developed specifically for natural language processing and language translation. Its uniqueness is its capability of figuring out the syntax and semantics of natural language, identifying dependency among words in a sentence to enable it to generate coherent, meaningful text as output. Another important feature of the Transformer network is that it can be trained in parallel. This allows it to use the parallel architecture of GPGPUs (General Purpose Graphics Processing Units), which are the primary hardware used to train neural networks.

ChatGPT's Transformer neural net is trained using a database of around 570 GB. The database is taken from a publicly available repository known as Common Crawl, which stores data from the Internet and a vast repository of books, and articles. The Internet data covers the period from 2016 to 2021.

The secret of the success of the Large Language Model (LLM) is in the architecture of the neural network that consists of a large number of layers and sub-layers. It requires an enormous effort in training it—billions of trials and errors. The training uses a sequence of words (the longer the sequence, the better), and the network figures out the relationships between words in a sentence and generates the next word reasonably accurately. The huge database on which it was trained enabled the neural network to learn the relationships between words and phrases and generate contextually relevant next word. The training method is called unsupervised as data is not labelled by humans and the neural net "discovers" the next word in a sentence automatically and checks it with the database for correctness.

ChatGPT uses several databases of human-to-human dialogs to train its neural network to enable it to give human-like replies to queries. The dialog database includes Persona-Chat (160,000 dialogs between two persons), Cornell Movie Dialog Corpus, Ubuntu Dialog Corpus (answers to technical questions), and a corpus of daily dialogs [9]. After initial training, ChatGPT's designers used supervised learning and fine-tuned the system, filtering inappropriate outputs. An important feature of ChatGPT is that it stores a user's earlier conversations with it and replies appropriately. The to-and-fro between ChatGPT and a user follows a natural course and gives a user the impression that he/she is talking to another person.

Contd.



Box 2. ***Contd.***

It is observed that sometimes ChatGPT 3.5 can hallucinate, that is, give incorrect answers which are not in its training database. Another drawback of ChatGPT is that it cannot always explain how it arrived at an answer for a given question.

The heart of the system is its trained LLM neural net. It costs a lot to train it. It is estimated that OpenAI spent around 4.5 million USD to train ChatGPT 3.5 and 1.3 Gwatt-hour of power was consumed during the training. Training the more powerful successor ChatGPT 4 is estimated to have cost nearly 100 million USD.

I give below a conversation between me and ChatGPT which is self-explanatory.

Q: *Can I assume that your response is not entirely based on search of a huge database but generated from the knowledge embedded in the Transformer neural net?*

ChatGPT: Yes, that is correct. My responses are not solely based on a search of a pre-existing database. While I do rely on a large corpus of text to train my language model, the knowledge and understanding of language is embedded within the parameters and weights of the Transformer neural network.

When presented with a question, I use a combination of pattern recognition and machine learning algorithms to generate a response based on the information and patterns learned during the training process. This allows me to generate responses that are not only relevant to the specific question but are also informative and accurate, even if the answer is not explicitly stated in the database.

So while my responses are based on my understanding of human language and the patterns and information learned from my training data, they are not a simple regurgitation of pre-existing information. Instead, my responses are generated based on a complex interplay of language understanding, pattern recognition, and machine learning algorithms.

Q: *Is it true that ChatGPT cannot explain how it got an answer to a question?*

ChatGPT: As an AI language model, ChatGPT generates answers based on patterns it has learned from vast amounts of data during its training process. While it can provide explanations for some answers, it is true that for some questions, ChatGPT may not be able to provide a clear and logical explanation for how it arrived at its answer.

Contd.

Box 2. ***Contd.***

This is because the inner workings of deep learning models like ChatGPT are complex and difficult to interpret. ChatGPT is a neural network with numerous layers, and the specific calculations and weights involved in producing an answer may be difficult to trace back and explain in human terms.

However, recent advances in the field of explainable AI aim to make machine learning models more transparent and interpretable. There are also some techniques like attention mechanisms that can help identify which parts of the input were most important in the model's decision-making process, providing some degree of explanation. Nonetheless, the extent to which ChatGPT can provide an explanation will depend on the specific question and context in which it is asked.

---

*D. Sharma, Deep learning without tears, *Resonance*, Vol.25, No.1, pp.15–32, Jan 2020.

Box 3. ChatGPT 4's Capabilities

ChatGPT 4 was released by OpenAI in March 2023. I give below its capabilities compared to ChatGPT 3.5.

1. It can accept images as input data and output text unlike ChatGPT 3.5, which can accept only textual input data.
2. It is trained on Terabytes of image data in addition to enormous text data exceeding 600 GB (the exact size is not published by OpenAI).
3. It can accept input data of up to 25,000 words as opposed to about 3125 words of input data that ChatGPT 3.5 can accept.
4. It is capable of giving better answers to questions in mathematics and science compared to ChatGPT 3.5 as it is trained on a large quantity of appropriate data. The database used is not disclosed by OpenAI.
5. It can write stories given a plot idea, write poetry, write programs given the requirements in natural language, and correct programs.
6. The answers given by it to questions are more cogent and correct compared to ChatGPT 3.5. It gives fewer incorrect answers.

Epilogue



Figure 3. Samples of pictures generated by DALL-E. A girl standing in a desert wearing a Texas hat and the painting of a man in a boat style of Cezanne. (Pictures courtesy of OpenAI).

Human-computer interaction in natural language has come a long way from 1965 to 2023. After the release of ChatGPT, many competitors have emerged [7, 8]. Google released its version of a large language model chatbot, called *Bard* in February 2023, and is now available in India. It is based on Google's Language Model for Dialog Applications (LaMDA). Microsoft had invested around 10 Billion USD in Open AI, and in return, it has used ChatGPT technology in *Bing AI*, a conversational search engine. Google Deepmind, an AI company, has developed *Chinchilla*, which is reported to outperform ChatGPT as it uses a larger transformer model. It is not yet released.

Google released its version of a large language model chatbot, called *Bard* in February 2023.

OpenAI has also released a system called DALL-E which generates pictures from a natural language description. Given a text such as "a girl standing in a desert wearing a Texas hat", it generates four images that match the description. A request to "paint a man in a boat in the style of Cezanne" generates four pictures matching the request (see *Figure* 3).

OpenAI released a system called DALL-E which generates pictures from a natural language description.

The advent of ChatGPT and similar generative large language model AI tools has made many groups anxious. Teachers are worried that students may use these AI tools to write essays, write poetry, write small programs, and do their homework. Many professions which depend on routine writing skills would be impacted. Lawyer's assistants and advertisement copywriters feel

threatened. In my opinion, these tools will assist many to improve their productivity and will not make them redundant. The nature of work would change and quality would improve. There are also concerns about the accuracy and bias in the responses by some current chatbots. The technology is very young—less than a year old. It is expected that over time chatbots would improve. Some countries are already drafting regulations on the development and release of LLM based neural networks (popularly called AI) fearing their social impact.

Acknowledgment

I thank Prof. P.C.P. Bhatt, Dr. N.Dayasindhu, Prof. R. Govindarajan, and Prof. S.Ramani for reading this article and for their comments that improved it.

Suggested Reading

[1] Imperial War Museum, UK, How Alan Turing Cracked the Enigma Code.https://www.iwm.org.uk/history/how-alan-turing-cracked-the-enigma-code#:~:text=With%20the%20help%20of%20captured,of%20all%20German%20naval%20signals.

[2] A. M. Turing, Computing machinery and intelligence, *Mind*, 59, pp.433–60, October 1950. https://academic.oup.com/mind/article/LIX/236/433/986238.

[3] J. Weizenbaum, ELIZA—A computer program for the study of natural language communication between man and machine, *Com. ACM*, Vol.9, No.1, pp.36-45, Jan 1966. https://dl.acm.org/doi/10.1145/365153.365168.

[4] J. Weizenbaum, *Computer Power and Human Reasoning*, W. H. Freeman, San Francisco, 1976.

[5] E. Adamopoulou, and L. Moussiades, Chatbots: History, technology, and applications, *Machine Learning and Applications*, Vol.2, December 2020, Elsevier. https://www.sciencedirect.com/science/article/pii/S2666827020300062#b107.

[6] V. Cerf, PARRY encounters ELIZA, Based on a demonstration at ICCC 1972. https://datatracker.ietf.org/doc/html/rfc439.

[7] ChatGPT Competitors Emerging in 2023, https://gigster.com/blog/5-chatgpt-competitors-emerging-in-2023/.

[8] Caroline M. Nyce, Nine Chatbots you can play with right now, The Atlantic, March 28, 2023. https://www.theatlantic.com/technology/archive/2023/03/chatgpt-generative-ai-chatbots-bing-google-bard/673533/

[9] D.Gewirtz, How does ChatGPT Work? ZDNET, March 10,2023. https://www.zdnet.com/article/how-does-chatgpt-work/

[10] J Grudin, ChatGPT and Chat History: Challenges for the New Wave, IEEE Computer, Vol.56, No.5, pp.94–100, May 2023.

Address for Correspondence
Prof. V. Rajaraman
Supercomputer Education &
Research Centre
Indian Institute of Science
Bengaluru 560 012, India.
Email: rajaram@iisc.ac.in