The Road to GPT

In November 2022, OpenAI debuted ChatGPT, an AI chatbot with an extensive knowledge base and deductive abilities capable of producing natural language responses almost indistinguishable from those written by humans. By January 2023, it had more than 100 million users—an explosive arrival that followed years of research and experimentation. Here we look at a few key milestones along the way. For more on the role of generative AI in scientific publishing, see this month's cover story (p. 26).

1948
A seminal paper by Claude Shannon lays the mathematical groundwork for modern communication and **information theory**.

1950
Alan Turing presents the **Turing test**, a measure of a machine's ability to exhibit intelligent behavior indistinguishable from that of a human.

1956
Noam Chomsky describes **phrase structure grammar**, a syntax-based approach that informs some early work on natural language processing.

1966
Joseph Weizenbaum develops the first chatbot, **ELIZA**. It simulates conversation by using a simple algorithm to generate text responses to questions.

1982
John Hopfield develops the **Hopfield network**, a recurrent neural network (RNN) that can learn and remember patterns.

1997
Sepp Hochreiter and Jürgen Schmidhuber introduce the concept of **long short-term memory**, a type of RNN capable of learning long-term dependencies.

2003
Yoshua Bengio and team develop a **feed-forward neural network** to model language, which predicts the next word when given a sequence of words.

2014
Ian Goodfellow develops **generative adversarial networks**, which pit two neural networks against each other to generate increasingly realistic content.

2015
Dzmitry Bahdanau introduces the **attention model**, which focuses only on the words that best help formulate an output, instead of remembering the entire input.

2017
Ashish Vaswani and a team at Google develop **transformer architecture**. It is based entirely on attention mechanisms, which require less training time than RNNs.

2018
Google implements transformers into its large language model (LLM), **BERT**. It can automatically learn relationships between words to predict meaning.

2018
Alec Radford shows how an LLM, through **generative pre-training** on a large and diverse set of data, can acquire knowledge and process dependencies.

2018
OpenAI releases **GPT**, the first generation of its generative pre-trained transformer LLMs. Trained on extensive data, it can generate novel human-like content.

November 2022
OpenAI debuts **ChatGPT**, a generative-AI chatbot built on top of version **GPT-3.5** of its LLM. Its responses are almost indistinguishable from those written by humans.

Since ChatGPT's remarkable debut, progress on generative AI has continued rapidly. For example, in March of this year, Google released its generative-AI chatbot **Bard**, and OpenAI unveiled **GPT-4**, a multimodal version of its LLM that allows image input as well as text.