MLA-CCCC Joint Task Force on Writing and AI Working Paper: Overview of the Issues, Statement of Principles, and Recommendations [pages 5-6]

History, Nomenclature, and Key Concepts Theorizing computers as intelligent, and even sentient, has long been tied to their ability to produce writing (such as Christopher Strachey's LoveLetters in 1952), and interact with human beings in conversation (such as Joseph Weizenbaum's ELIZA in 1964). Members of the Modern Language Association and the College Conference on Composition and Communication have likewise been using and working with automatic text generation for decades, a corpus of research that has provided critical context for this group's work.

Generative artificial intelligence is often conflated with artificial general intelligence (AGI), which is the human-like, seemingly sentient AI that is still the stuff of science fiction. Generative AI, by contrast, refers to computer systems that can produce, or generate, various forms of traditionally human expression, in the form of digital content including language, images, video, and music. LLMs are a subset of generative AI used to deliver text-based formats like prose, poetry, or even programming code. The GPT grouping of LLMs currently enjoys the most public recognition, but there are others available; GPT itself stands for generative pre-trained transformer, with each piece of that nomenclature bearing a specific technical meaning.

For our purposes it is enough to say that in 2018 the nonprofit OpenAI developed the means to yoke generative pre-trained models to the so-called transformer architecture introduced by Google in the previous year, thus delivering dramatic increases in performance. The widely discussed ChatGPT is a specific application of GPT-3 (now GPT-4), released in late 2022 by OpenAI. It combines an easy-to-access browser interface with a chatbot style of interaction, whereby a user can enter a series of discursive prompts and engage with the outputs of the model in an ongoing dialogic stream.

LLMs work by using statistics and probability to predict what the next character (i.e., letter, punctuation mark, even a blank space) is likely to be in an ongoing sequence, thereby “spelling” words, phrases, and entire sentences and paragraphs. It is not unlike autocomplete, but more powerful. LLMs are trained on vast bodies of preexisting text (such as content from the Internet), which, to some extent, predetermine their output. All of the text a model generates is original in the sense that it represents combinations of letters and words that generally have no exact match in the training documents, yet the content is also unoriginal in that it is determined by patterns in its training data. The same language model may generate a variety of different sequences in response to the same input prompt. A model cannot reliably report on which sources in its training data contributed to any given output. All of this combines to make the output of LLMs qualitatively different from any other form of text, even texts that might have been computer generated according to some other method. It should be noted that given the assortment of software applications drawing on LLMs, presenting them through different user interfaces to offer various affordances, they are also unavoidable.

Although it is often tempting to speak in terms of what an LLM is “doing” or “intending” or even “thinking,” what we are witnessing is the production of word sequences that look like intentional human text through a process of statistical correlation. As the models are refined, expand their language corpora, and draw on greater computational power, their outputs mimic the writing of sentient humans more convincingly. LLMs do not, however, “think” in the way that we would define such an activity as it takes place in human cognition.

Source: https://aiandwriting.hcommons.org/working-paper-1/