Congressional Research Service
Informing the legislative debate since 1914

IN FOCUS

June 9, 2023

Generative Artificial Intelligence: Overview, Issues, and Questions for Congress

Generative artificial intelligence (GenAI) refers to AI systems, in particular those using machine learning (ML) and trained on large volumes of data, that are able to generate new content. In contrast, other AI systems may have a primary goal of classifying data, such as facial recognition image data, or making decisions, such as those used in autonomous vehicles. GenAI systems, when prompted (often by a user inputting text), can create various outputs, including text responses (e.g., OpenAI's ChatGPT and Google's Bard), images (e.g., Stability AI's Stable Diffusion and Midjourney's self-titled program), videos, computer code, or music.

The recent public release of many GenAI tools, and the race by companies to develop ever-more powerful models, have generated widespread discussion of their capabilities, potential concerns with their use, and debates about their governance and regulation. This CRS InFocus describes the development and uses of GenAI, concerns raised by the use of GenAI tools, and considerations for Congress. For additional considerations related to data privacy, see CRS Report R47569, *Generative Artificial Intelligence and Data Privacy: A Primer*, by Kristen E. Busch.

Background

AI can generally be thought of as computerized systems that work and react in ways commonly considered to require human intelligence, such as learning, solving problems, and achieving goals under uncertain and varying conditions, with varying levels of autonomy. AI is not one thing; AI systems can encompass a range of technologies, methodologies, and application areas, such as natural language processing, robotics, and facial recognition.

The AI technologies underpinning many GenAI tools are the result of decades of research. For example, recurrent neural networks (RNNs, a type of ML loosely modeled after the human brain that detect patterns in sequential data) underwent a period of much development and improvement in the 1980s-1990s. RNNs can generate text, but they are limited in retaining contextual information across large strings of words, are slow to train, and are not easily scaled up by increasing computational power or training data size.

More recent technical advances—notably the introduction of the Transformer architecture by Google researchers in 2017 and improvements in Generative Pre-Trained Transformer (GPT) models since around 2019—have contributed to dramatic improvement in GenAI performance. Transformer models process a sequence of whole sentences (rather than analyzing word by word), which helps them to "remember" past information. They use mathematical techniques called *attention* or *self-attention* to detect how data elements, even when far away sequentially, influence and depend on each other. These methods make GPT models faster to train, more efficient in understanding context, and highly scalable.

Other critical components to the recent GenAI advances are the availability of large amounts of data and the size of their language models. Large language models (LLMs) are AI systems that aim to model language, sometimes using millions or billions of parameters (i.e., numbers in the model that determine how inputs are converted to outputs). Repeatedly tweaking these parameters, using mathematical optimization techniques, and large amounts of data and computational power, increases model performance.

Notably, GenAI models work to match the style and appearance of the underlying training data. They also have been shown to have *capability overhang*, meaning hidden capabilities that their developers and users did not anticipate but that are emerging as the models grow larger.

Beneficial Uses and Concerns

The increase in size of recent GenAI systems (with hundreds of billions of parameters) has led to drastically improved capabilities over previous systems (with millions or a few billion parameters). For example, as released in 2020, OpenAI's GPT-3 could translate sentences from English to French with few to no training examples and outperformed previous models that were explicitly trained to solve that task. GPT-4, released in 2023, improved on various benchmarks, such as scores on simulated graduate and professional exams, and on traditional ML benchmarks.

OpenAI's ChatGPT, reportedly the fastest growing consumer application in history, uses GPT-3.5 or, for paid subscribers, GPT-4, which is the latest, most powerful version. ChatGPT is a conversational AI chatbot that can generate a variety of text outputs, such as emails, essays, and casual conversations. Microsoft has incorporated GPT-4 into its Bing search engine as a chatbot called Bing Chat, and numerous other chatbots also use the GPT-3 and GPT-4 models. In addition to GPT models, other LLM chatbots include Google's Pathways Language Model 2 (PaLM 2) in the Bard chatbot, and Hugging Face's BLOOM model in SambaNova's BLOOMChat.

Despite the impressive abilities of GenAI, its growing prominence is raising concerns. For example, the tendency of GenAI to make things up, sometimes referred to as *hallucinating*, might result in the tools generating and amplifying misinformation or being used to create and spread disinformation. OpenAI notes that even a powerful model like GPT-4 "is not fully reliable" and "great care should be taken when using language model outputs, particularly in high-stakes contexts, with the exact protocol (such as human review, grounding with additional context, or avoiding high-stakes uses altogether) matching the needs of a specific use-case." Additionally, because the models are generally trained on large amounts of data scraped from the internet, they can incorporate, reflect, and potentially amplify biases in such data.

Particular GenAI use cases, such as in education and scientific research, have also raised questions about ethical and transparent use, whether restrictions should be implemented and whether they are effective, and the accuracy of detection tools.

As GenAI use grows, analysts have begun considering how it might affect jobs and productivity. For example, will these tools complement workers' skills in existing jobs and create new jobs? Will GenAI automate some jobs, displacing workers? While these have been long-standing concerns for automation and AI technologies, the speed, capability, and widespread use of the latest GenAI models have heightened them.

LLMs have been characterized as *foundation models* (also called *general-purpose AI*), meaning models trained on broad data that can be adapted to a wide range of downstream tasks. (In contrast, many other AI systems are built, trained, and used for particular purposes.) As described by the Stanford University Institute for Human-Centered AI, foundation models may be built upon or integrated into multiple AI systems across a variety of domains, with the potential for both benefits (e.g., concentrating efforts to reduce bias and improve robustness) and drawbacks (e.g., security failures or inequities that flow out to downstream applications, amplifying their harms).

Federal AI Laws and GenAI Legislation

Numerous bills focused on AI, or including AI-focused provisions, have been enacted in prior Congresses. For example, the National Artificial Intelligence Initiative Act of 2020 (Division E of P.L. 116-283) codified the establishment of a national AI initiative and associated federal offices and committees. In the 118th Congress, at least 50 AI-related bills have been introduced.

Specifically regarding GenAI, the Identifying Outputs of Generative Adversarial Networks (IOGAN) Act (P.L. 116-258) directed federal support of research on generative adversarial networks. In the 118th Congress, at least four bills have been introduced that specifically include GenAI, including those pertaining to transparency and accountability of GenAI use in political advertisements, disclosure of GenAI outputs, and federal oversight of digital platforms.

Potential Questions for Congress

Like the private sector, federal agencies and Congress have begun testing GenAI uses, including for its potential to help with office tasks such as creating and summarizing content, writing speeches, and drafting bills. At the same time, Congress has begun considering whether and how to implement guardrails for GenAI technologies. As that work continues, along with further development and growing use of GenAI tools, Congress might consider a range of questions and potential actions.

- **Bias and ethics of use.** How are the private and public sectors using federal guidance documents and frameworks for AI evaluation and risk management to address bias and manage risk in GenAI systems? Should the deployment of GenAI models in high-risk scenarios (e.g., mental health therapy or generating forensic sketches) be restricted?
- **Testing and transparency.** The biggest models deployed today, such as GPT-4 and PaLM 2, are closed source, proprietary models. While many companies state that they conduct internal testing and are evaluating options for external validation and testing, Congress might consider whether and how to support or require independent testing and reporting of results.
- **Economic and workforce impacts.** Researchers in industry, the private sector, and academia have begun analyzing GenAI's potential widespread effects on labor. The National Academy of Sciences is currently updating a study on automation and the U.S. workforce. In addition to assessing those findings, Congress might consider the appropriate federal role in supporting U.S. workforce reskilling or upskilling in response to shifting job tasks caused by the implementation of GenAI. It might also consider whether and how to increase AI expertise in the government's own workforce.
- **Research and competition.** Estimates put training costs for GenAI models like GPT-3, with 175 billion parameters, at over $4.6 million. Some analysts have argued that cost, use of proprietary data, and access to vast computing power will create a divide between those who can train the most cutting-edge LLMs (e.g., large technology firms) and those who cannot (e.g., nonprofits, startups, universities). Congress might consider ways to support access to data, training, and computing resources, such as through codifying recommendations in the final report of the National AI Research Resource Task Force.
- **Oversight and regulation.** How might Congress regulate GenAI technologies while supporting innovation and international competitiveness? Do federal regulatory agencies have the authorities and resources to adequately oversee and regulate GenAI tools to minimize risks while supporting benefits? If not, what additional authorities are needed? How might federal oversight and regulation for GenAI be distinct from that for AI technologies more broadly?

Laurie A. Harris, Analyst in Science and Technology Policy