



Updated April 2, 2025

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

Generative artificial intelligence (GenAI) refers to AI models, in particular those that use machine learning (ML) and are trained on large volumes of data, that are able to generate new content. In contrast, other AI models may have a primary goal of classifying data, such as facial recognition image data, or making decisions, such as those used in automated vehicles. GenAI, when prompted (often by a user inputting text), can create various outputs, including text, images, videos, computer code, or music.

The public release of many GenAI tools, and the race by companies to develop ever-more powerful AI models, have generated widespread discussion of their capabilities, potential concerns with their use, and debates about their governance and regulation. This CRS In Focus describes the development and uses of GenAI, concerns raised by the use of GenAI tools, and considerations for Congress.

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 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 detects patterns in sequential data, underwent much development and improvement in the 1980s-1990s. RNNs can generate text, but they have limited ability to retain 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. 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 recent GenAI advances have been 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 have also demonstrated *emergent* abilities, meaning capabilities that their developers and users did not anticipate but that are emerging as the models grow larger.

LLMs have been characterized as *foundation models* (also called *general-purpose AI*), meaning models trained on large and diverse datasets that can be adapted to a wide range of downstream tasks. As described by the Stanford University Institute for Human-Centered AI, foundation models may be built on or integrated into multiple AI systems across various domains (e.g., text-based GPT models that can perform arithmetic and computer programming tasks, which were outside the scope of their original training). This capability has 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 to downstream applications).

Capabilities and Advances

The increase in size of recent GenAI models (with hundreds of billions or trillions of parameters) has led to improved capabilities over previous systems (with millions or a few billion parameters). According to the *AI Index 2024 Annual Report*, “LLMs have surpassed human performance on traditional English-language benchmarks,” and the report argues that the “rapid advancement has led to the need for more comprehensive benchmarks.”

Initial GenAI tools tended to excel at single input and output types—such as text to text for chatbots or text to image for image generators. More *multimodal* GenAI models are available now, meaning they can process and integrate information from multiple types of data or modalities simultaneously. For example, Google’s Gemini models can use text-image-audio-video inputs and provide text-image outputs.

Beginning in late 2024, companies introduced what have been termed *reasoning models*—models that use a *chain-of-thought* technique in an attempt to “refine their thinking process, try different strategies, and recognize their mistakes” (e.g., OpenAI’s o1 and o3-mini models, Anthropic’s Claude 3.7 Sonnet, and DeepSeek’s R1 model), though reasoning models still frequently make mistakes. Along with the development of large-scale models,

innovations in model design and training have led to improvements in the speed and capabilities of GenAI tools.

Concerns and Potential Risks

Despite the impressive abilities of GenAI, its rapid growth has raised concerns and discussions about managing potential risks. For example, the tendency of GenAI to produce incorrect or misleading results, sometimes referred to as *confabulation* or *hallucinating*, might lead to the tools generating and amplifying misinformation or being used to create and spread disinformation. For example, in January 2024, researchers found that legal mistakes with LLMs were “pervasive,” and in December 2024, expert testimony was thrown out after a judge discovered fake information produced by GenAI. 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. OpenAI has noted that even powerful models like GPT-4 have limitations, being “not fully reliable,” and that “great care should be taken when using language model outputs, particularly in high-stakes contexts.” Concerns about model inputs and limitations have led to calls for greater transparency about model training and validation.

Safety and security risks are ongoing concerns with GenAI tools. For example, *jailbreaks* can trick an AI system into ignoring built-in safety rules. *Prompt injection attacks* disguise malicious inputs as non-malicious prompts, manipulating a system into such actions as leaking sensitive data. Also, depending on a company’s privacy policies, U.S. user data may be stored on servers in foreign countries as has been flagged with the Chinese company DeepSeek.

As GenAI use grows, analysts have been considering how it might affect jobs and productivity. For example, will these tools complement workers’ skills in existing jobs, create new jobs, or automate some jobs, displacing workers? What impact might GenAI have on economic productivity? While these have been long-standing questions for automation and AI technologies, the speed, capability, and adoption of GenAI has heightened labor-related concerns.

Federal AI Laws and GenAI Legislation

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. 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. Some Members of Congress have introduced bills that included GenAI, such as those pertaining to the development of technical standards, deepfakes and AI-generated voice messages, and transparency and accountability of GenAI use.

Considerations for Congress

Federal agencies and Congress have been exploring GenAI uses, including for office tasks such as creating and summarizing content, writing speeches, and drafting bills. At the same time, several bills have been introduced in the 119th Congress regarding GenAI, building on various bills introduced in the 118th Congress to implement guardrails

for GenAI technologies in the private sector. As GenAI development continues and its use grows, Congress might consider questions and potential actions such as these:

- **Bias and ethics of use.** The private and public sectors may be using federal guidance and frameworks for AI to address bias and manage risks for GenAI. Congress might consider whether additional sector-specific guidance would be helpful and whether the deployment of GenAI models in high-risk scenarios (e.g., mental health therapy or generating forensic sketches) requires restrictions.
- **Testing and transparency.** Many of the biggest models deployed today are closed-source and proprietary. Companies state that they conduct internal testing and are evaluating options for external validation and testing. Congress might consider the adequacy of industry self-assessment and whether and how to support or require independent testing and reporting of results.
- **Economic and workforce impacts.** Researchers in industry and academia have begun analyzing GenAI’s potential widespread effects on labor. In November 2024, the National Academy of Sciences released an updated study on AI and the future of work. 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, such as using statutory authority to establish a federal AI scholarship-for-service program (per P.L. 117-167).
- **Research and competition.** Estimates have put training costs for large GenAI models in the millions to tens of millions of dollars. 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, start-ups, universities). Congress might consider whether to support access to data, training, and computing resources, such as through codifying an effort like the National AI Research Resource. Further, the release of advanced, open-source models such as DeepSeek’s R1 have raised concerns about national security and international competition. Congress might consider whether to support U.S. leadership in AI innovation, such as through technical standards development and coordination with allied countries.
- **Oversight and regulation.** In considering possible regulation of GenAI technologies, Congress might weigh potential impacts on innovation and international competitiveness. Do federal regulatory agencies have the authorities and resources to adequately oversee GenAI tools to minimize risks and support 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 Harris, Analyst in Science and Technology Policy

Disclaimer

This document was prepared by the Congressional Research Service (CRS). CRS serves as nonpartisan shared staff to congressional committees and Members of Congress. It operates solely at the behest of and under the direction of Congress. Information in a CRS Report should not be relied upon for purposes other than public understanding of information that has been provided by CRS to Members of Congress in connection with CRS's institutional role. CRS Reports, as a work of the United States Government, are not subject to copyright protection in the United States. Any CRS Report may be reproduced and distributed in its entirety without permission from CRS. However, as a CRS Report may include copyrighted images or material from a third party, you may need to obtain the permission of the copyright holder if you wish to copy or otherwise use copyrighted material.