

# How Generative AI Is Changing Creative Work

by Thomas H. Davenport and Nitin Mittal

November 14, 2022



Andriy Onufriyenko/Getty Images

**Summary.** Generative AI models for businesses threaten to upend the world of content creation, with substantial impacts on marketing, software, design, entertainment, and interpersonal communications. These models are able to produce text and images: blog... [more](#)

Large language and image AI models, sometimes called generative AI or foundation models, have created a new set of opportunities for businesses and professionals that perform content creation. *Some of these opportunities include:*

- 1. Automated content generation: Large language and image AI models can be used to automatically generate content, such as articles, blog posts, or social media posts. This can be a valuable time-saving tool for businesses and professionals who create content on a regular basis.*
- 2. Improved content quality: AI-generated content can be of higher quality than content created by humans, due to the fact that AI models are able to learn from a large amount of data and identify patterns that humans may not be able to see. This can result in more accurate and informative content.*
- 3. Increased content variety: AI models can generate a variety of content types, including text, images, and video. This can help businesses and professionals to create more diverse and interesting content that appeals to a wider range of people.*
- 4. Personalized content: AI models can generate personalized content based on the preferences of individual users. This can help businesses and professionals to create content that is more likely to be of interest to their target audience, and therefore more likely to be read or shared.*

How adept is this technology at mimicking human efforts at creative work? Well, for an example, the italicized text above was written by GPT-3, a “large language model” (LLM) created by OpenAI, in response to the first sentence, which we wrote. GPT-3’s text reflects the strengths and weaknesses of most AI-generated content. First, it is sensitive to the prompts fed into it; we tried several alternative prompts before settling on that sentence. Second, the system writes reasonably well; there are no

grammatical mistakes, and the word choice is appropriate. Third, it would benefit from editing; we would not normally begin an article like this one with a numbered list, for example. Finally, it came up with ideas that we didn't think of. The last point about personalized content, for example, is not one we would have considered.

Overall, it provides a good illustration of the potential value of these AI models for businesses. They threaten to upend the world of content creation, with substantial impacts on marketing, software, design, entertainment, and interpersonal communications. This is not the “artificial general intelligence” that humans have long dreamed of and feared, but it may look that way to casual observers.

## **What Is Generative AI?**

Generative AI can already do a lot. It's able to produce text and images, spanning blog posts, program code, poetry, and artwork (and even winning competitions, controversially). The software uses complex machine learning models to predict the next word based on previous word sequences, or the next image based on words describing previous images. LLMs began at Google Brain in 2017, where they were initially used for translation of words while preserving context. Since then, large language and text-to-image models have proliferated at leading tech firms including Google (BERT and LaMDA), Facebook (OPT-175B, BlenderBot), and OpenAI, a nonprofit in which Microsoft is the dominant investor (GPT-3 for text, DALL-E2 for images, and Whisper for speech). Online communities such as Midjourney (which helped win the art competition), and open-source providers like HuggingFace, have also created generative models.

These models have largely been confined to major tech companies because training them requires massive amounts of data and computing power. GPT-3, for example, was initially trained on 45 terabytes of data and employs 175 billion

parameters or coefficients to make its predictions; a single training run for GPT-3 cost \$12 million. Wu Dao 2.0, a Chinese model, has 1.75 trillion parameters. Most companies don't have the data center capabilities or cloud computing budgets to train their own models of this type from scratch.

But once a generative model is trained, it can be “fine-tuned” for a particular content domain with much less data. This has led to specialized models of BERT — for biomedical content (BioBERT), legal content (Legal-BERT), and French text (CamemBERT) — and GPT-3 for a wide variety of specific purposes. NVIDIA's BioNeMo is a framework for training, building and deploying large language models at supercomputing scale for generative chemistry, proteomics, and DNA/RNA. OpenAI has found that as few as 100 specific examples of domain-specific data can substantially improve the accuracy and relevance of GPT-3's outputs.

To use generative AI effectively, you still need human involvement at both the beginning and the end of the process.

To start with, a human must enter a prompt into a generative model in order to have it create content. Generally speaking, creative prompts yield creative outputs. “Prompt engineer” is likely to become an established profession, at least until the next generation of even smarter AI emerges. The field has already led to an 82-page book of DALL-E 2 image prompts, and a prompt marketplace in which for a small fee one can buy other users' prompts. Most users of these systems will need to try several different prompts before achieving the desired outcome.

Then, once a model generates content, it will need to be evaluated and edited carefully by a human. Alternative prompt outputs may be combined into a single document. Image generation may require substantial manipulation. Jason Allen, who won the Colorado “digitally manipulated photography” contest with help from Midjourney, told a reporter that he spent more than 80 hours

making more than 900 versions of the art, and fine-tuned his prompts over and over. He then improved the outcome with Adobe Photoshop, increased the image quality and sharpness with another AI tool, and printed three pieces on canvas.

Generative AI models are incredibly diverse. They can take in such content as images, longer text formats, emails, social media content, voice recordings, program code, and structured data. They can output new content, translations, answers to questions, sentiment analysis, summaries, and even videos. These universal content machines have many potential applications in business, several of which we describe below.

## **Marketing Applications**

These generative models are potentially valuable across a number of business functions, but marketing applications are perhaps the most common. Jasper, for example, a marketing-focused version of GPT-3, can produce blogs, social media posts, web copy, sales emails, ads, and other types of customer-facing content. It maintains that it frequently tests its outputs with A/B testing and that its content is optimized for search engine placement. Jasper also fine tunes GPT-3 models with their customers' best outputs, which Jasper's executives say has led to substantial improvements. Most of Jasper's customers are individuals and small businesses, but some groups within larger companies also make use of its capabilities. At the cloud computing company VMWare, for example, writers use Jasper as they generate original content for marketing, from email to product campaigns to social media copy. Rosa Lear, director of product-led growth, said that Jasper helped the company ramp up our content strategy, and the writers now have time to do better research, ideation, and strategy.

Kris Ruby, the owner of public relations and social media agency Ruby Media Group, is now using both text and image generation from generative models. She says that they are effective at

maximizing search engine optimization (SEO), and in PR, for personalized pitches to writers. These new tools, she believes, open up a new frontier in copyright challenges, and she helps to create AI policies for her clients. When she uses the tools, she says, “The AI is 10%, I am 90%” because there is so much prompting, editing, and iteration involved. She feels that these tools make one’s writing better and more complete for search engine discovery, and that image generation tools may replace the market for stock photos and lead to a renaissance of creative work.

DALL-E 2 and other image generation tools are already being used for advertising. Heinz, for example, used an image of a ketchup bottle with a label similar to Heinz’s to argue that “This is what ‘ketchup’ looks like to AI.” Of course, it meant only that the model was trained on a relatively large number of Heinz ketchup bottle photos. Nestle used an AI-enhanced version of a Vermeer painting to help sell one of its yogurt brands. Stitch Fix, the clothing company that already uses AI to recommend specific clothing to customers, is experimenting with DALL-E 2 to create visualizations of clothing based on requested customer preferences for color, fabric, and style. Mattel is using the technology to generate images for toy design and marketing.

## **Code Generation Applications**

GPT-3 in particular has also proven to be an effective, if not perfect, generator of computer program code. Given a description of a “snippet” or small program function, GPT-3’s Codex program — specifically trained for code generation — can produce code in a variety of different languages. Microsoft’s Github also has a version of GPT-3 for code generation called CoPilot. The newest versions of Codex can now identify bugs and fix mistakes in its own code — and even explain what the code does — at least some of the time. The expressed goal of Microsoft is not to eliminate

human programmers, but to make tools like Codex or CoPilot “pair programmers” with humans to improve their speed and effectiveness.

The consensus on LLM-based code generation is that it works well for such snippets, although the integration of them into a larger program and the integration of the program into a particular technical environment still require human programming capabilities. Deloitte has experimented extensively with Codex over the past several months, and has found it to increase productivity for experienced developers and to create some programming capabilities for those with no experience.

In a six-week pilot at Deloitte with 55 developers for 6 weeks, a majority of users rated the resulting code’s accuracy at 65% or better, with a majority of the code coming from Codex. Overall, the Deloitte experiment found a 20% improvement in code development speed for relevant projects. Deloitte has also used Codex to translate code from one language to another. The firm’s conclusion was that it would still need professional developers for the foreseeable future, but the increased productivity might necessitate fewer of them. As with other types of generative AI tools, they found the better the prompt, the better the output code.

## **Conversational Applications**

LLMs are increasingly being used at the core of conversational AI or chatbots. They potentially offer greater levels of understanding of conversation and context awareness than current conversational technologies. Facebook’s BlenderBot, for example, which was designed for dialogue, can carry on long conversations with humans while maintaining context. Google’s BERT is used to understand search queries, and is also a component of the company’s DialogFlow chatbot engine. Google’s LaMBA, another LLM, was also designed for dialog, and conversations with it

convinced one of the company's engineers that it was a sentient being— an impressive feat, given that it's simply predicting words used in conversation based on past conversations.

None of these LLMs is a perfect conversationalist. They are trained on past human content and have a tendency to replicate any racist, sexist, or biased language to which they were exposed in training. Although the companies that created these systems are working on filtering out hate speech, they have not yet been fully successful.

## **Knowledge Management Applications**

One emerging application of LLMs is to employ them as a means of managing text-based (or potentially image or video-based) knowledge within an organization. The labor intensiveness involved in creating structured knowledge bases has made large-scale knowledge management difficult for many large companies. However, some research has suggested that LLMs can be effective at managing an organization's knowledge when model training is fine-tuned on a specific body of text-based knowledge within the organization. The knowledge within an LLM could be accessed by questions issued as prompts.

Some companies are exploring the idea of LLM-based knowledge management in conjunction with the leading providers of commercial LLMs. Morgan Stanley, for example, is working with OpenAI's GPT-3 to fine-tune training on wealth management content, so that financial advisors can both search for existing knowledge within the firm and create tailored content for clients easily. It seems likely that users of such systems will need training or assistance in creating effective prompts, and that the knowledge outputs of the LLMs might still need editing or review before being applied. Assuming that such issues are addressed, however, LLMs could rekindle the field of knowledge management and allow it to scale much more effectively.



## Deepfakes and Other Legal/Ethical Concerns

We have already seen that these generative AI systems lead rapidly to a number of legal and ethical issues. “Deepfakes,” or images and videos that are created by AI and purport to be realistic but are not, have already arisen in media, entertainment, and politics. Heretofore, however, the creation of deepfakes required a considerable amount of computing skill. Now, however, almost anyone will be able to create them. OpenAI has attempted to control fake images by “watermarking” each DALL-E 2 image with a distinctive symbol. More controls are likely to be required in the future, however — particularly as generative video creation becomes mainstream.

Generative AI also raises numerous questions about what constitutes original and proprietary content. Since the created text and images are not exactly like any previous content, the providers of these systems argue that they belong to their prompt creators. But they are clearly derivative of the previous text and images used to train the models. Needless to say, these technologies will provide substantial work for intellectual property attorneys in the coming years.

From these few examples of business applications, it should be clear that we are now only scratching the surface of what generative AI can do for organizations and the people within them. It may soon be standard practice, for example, for such systems to craft most or all of our written or image-based content — to provide first drafts of emails, letters, articles, computer programs, reports, blog posts, presentations, videos, and so forth. No doubt that the development of such capabilities would have dramatic and unforeseen implications for content ownership and intellectual property protection, but they are also likely to revolutionize knowledge and creative work. Assuming that these AI models continue to progress as they have in the short time they have existed, we can hardly imagine all of the opportunities and implications that they may engender.

**Thomas H. Davenport** is the President's Distinguished Professor of Information Technology and Management at Babson College, a visiting scholar at the MIT Initiative on the Digital Economy, and a senior adviser to Deloitte's AI practice. He is a coauthor of *All-in on AI: How Smart Companies Win Big with Artificial Intelligence* (Harvard Business Review Press, 2023).

NM

**Nitin Mittal** is a principal at Deloitte Consulting, the leader of its analytics and cognitive offering, and a coleader of Deloitte's AI strategic growth offering. He is a coauthor of *All-in on AI: How Smart Companies Win Big with Artificial Intelligence* (Harvard Business Review Press, 2023).

## Recommended For You

---


**How Generative AI Could Disrupt Creative Work**



**Generative AI Has an Intellectual Property Problem**



**A Framework for Picking the Right Generative AI Project**



PODCAST

**How Generative AI Changes Creativity**



