



Game changer:

The startling power generative AI is bringing to software development



A newly emerging type of artificial intelligence promises to rewrite the way software is built and maintained

Software development is poised to take a dramatic leap—with big implications for how developers work and how corporations build and maintain the digital code underlying their operations.

The catalyst for this dramatic transformation is generative artificial intelligence (AI), an emerging form of AI that can create original articles, essays, images, music, and yes, code, by building on patterns it finds in existing text, audio files, images, and software.

Generative AI has already gained notoriety for its ability to produce images and text. Late last summer, it generated headlines when an AI-generated image won an art contest. In November 2022, it captured the world’s attention again with the debut of ChatGPT, a wickedly sophisticated, next-generation chatbot created by OpenAI, a research lab in San Francisco. Users order up new content by giving these tools simple prompts like “how did the Space Race change America?” (text) or “a purple rabbit hopping on orange clouds” (image). ChatGPT attracted more than a million users in just five days, making it one of the most rapidly adopted new technology platforms in history. (Exhibit 1)

Exhibit 1: Racing to one million ChatGPT hit one million users in record time



Source: DiploFoundation, Generative AI models - a fun game that can easily get out of hand? December 22, 2022. <https://www.diplomacy.edu/blog/generative-AI-models-a-fun-game-that-can-easily-get-out-of-hand/>

A paradigm shift as significant as the cloud or DevOps

While the public plays with this new toy—and college professors try to figure out how to differentiate ChatGPT essays from student work—similar platforms are being used to generate software, including GitHub Copilot, also based on OpenAI models, and Tabnine, from the Canadian company of the same name. These systems have kinks and won't put developers out of business anytime soon. But they are making it possible for developers willing to work with their limitations to hand off the repetitive parts of their jobs, supercharging their productivity and speeding the time to deployment for new software.

To reiterate, this isn't a few years away. As we write, generative AI models trained against the vast expanse of open-source code available on the internet are already explaining poorly documented code, generating documentation for code, and even writing functions or relatively targeted pieces of code, all with minimal direction from humans. For many developers, generative AI will become the most valuable coding partner they will ever know.

If it lives up to its promise, generative AI will herald a paradigm shift every bit as significant as the cloud or DevOps, the development process that enables faster delivery of more reliable software products and services. In terms of how corporations develop and maintain software, it will prompt changes as big as, and likely even more impactful than, those created by Agile development methods, which enable rapid responses to changing software requirements and customer feedback.

Skeptics are plentiful, to be sure. Some users question generative AI's capabilities. It is a nascent technology whose output can sometimes miss the mark—or simply be incorrect. But in our view, much of today's skepticism is misplaced. The biggest challenge to collaboration between developers and generative AI isn't the technology's capabilities but its

cost for organizations that want to work not just with open-source code but also with the source code and technical documentation the company uses to run its own operations. Training a generative AI model this way, either on their own or with the help of a consultant, could prove an expensive undertaking.

Still, what generative AI is already doing is so compelling that it's hard to believe companies won't keep pushing forward. Take OpenAI's DALL-E, one of several generative AI models. It can produce high-quality images—art—based on natural language descriptions (prompts) entered by users. Like ChatGPT, it's based on a third-generation "Generative Pre-trained Transformer," GPT-3, from OpenAI. GPT-3 features a deep learning neural network with 175 billion machine learning parameters. OpenAI allows limited tinkering with this model for free, but charges for more extensive work.

More directly for this discussion, look what's happening with GitHub Copilot, which uses the OpenAI Codex artificial intelligence model to suggest software code and entire functions based on natural language prompts. It's already affordable in many applications. The cost for individual programmers to use it is \$10 a month. For businesses the cost is \$19 per month per user.¹



¹ Pricing for GitHub Copilot for Business, <https://docs.github.com/en/enterprise-cloud@latest/billing/managing-billing-for-github-copilot/about-billing-for-github-copilot>

Still, for chief information officers and their employers, the most valuable near-term use case for generative AI is likely to be the previously mentioned training of it on the proprietary source code and technical documentation being used by their own organizations. “Fine-tuning” generative AI models this way could lead to huge productivity gains and change forever the economics of creating and maintaining software in a corporate environment.

This would be especially compelling in weak economic environments when companies are faced with added cost pressures, as it would allow them to get more productivity out of their existing software development team. It would be particularly useful to those with the most intransigent technology stacks.

Among other things, we see developers at forward-thinking organizations turning to generative AI models to:

- 1 Explain the functionality of code** they did not write but are now tasked with maintaining or upgrading.
- 2 Provide guidance** on which classes or libraries of code have the necessary functionality to implement specific requirements.
- 3 Improve the consistency of code** through automated refactoring, style changes, and automatically generating explanatory comments.
- 4 Simplify code** by suggesting less complex alternatives that can ease future maintenance.
- 5 Explore alternative solutions** to coding challenges, widening the developer’s perspective on possible options.

- 6 Generate unit tests** that can help reduce human error in coding and resolve corner-case errors that occur outside normal operating parameters.
- 7 Translate code from one language to another**, helping with legacy migration and development of use cases.
- 8 Write documentation** such as user stories and success criteria.
- 9 Accelerate debugging** with generative AI’s ability to quickly identify coding mistakes.
- 10 Write code that goes beyond templates or stub code** to jump-start new tasks.

Do things that were previously impossible

By lowering barriers to entry for new developers on complex codebases, generative AI also will allow companies to do things that were previously impossible, such as:

- **Rapidly onboard large groups of developers** to accelerate new features or major changes to software. These developers would be productive quickly and require less guidance from existing developers.
- **Make developers more fungible** across frameworks, platforms, products, and systems of record. Generative AI will provide the scaffolding and guidance they need to work on a wider range of projects than they would normally be able to handle.
- **Make corporations less beholden to developers of their legacy codebase** as more of that legacy knowledge is built into a fine-tuned generative model.
- **Make the purchase or sale of software-related intellectual property less dependent on the retention of developers.** Generative AI will be able to codify much of their legacy knowledge. An existing fine-tuned generative AI model also would improve the quality of technical due diligence during the acquisition.
- **Use offshore resources more effectively.** A contextually informed generative AI assistant could provide guidance to an offshore coding team while senior onshore engineers and product owners get some sleep. This also may enable changes to the resource mix to take greater advantage of offshore resources.

The level of effort required to fine-tune a commercial generative AI model on a corporation's proprietary codebase, and the efficacy of that fine-tuning on specific commercially available generative models, remains unknown at the moment. It is possible that the effort, and associated cost, will prove beyond the means of some organizations, at least in the near future. But the potential value is so great that we believe any large organization with a lot of developers that isn't urgently exploring the opportunity is doing its stakeholders a disservice. A proof of concept at any large corporation with a significant legacy codebase would go a long way toward demonstrating the technology's potential—and likely give that organization a leg up on its competitors.

Remaining challenges and risks

As with the adoption of any new technology, a broad move to use generative AI in the corporate environment will surely hit some bumps along the way. For all of their sophistication, current iterations of generative AI can still generate flawed results. Even where results are rock solid, building in traceability and transparency may be important because many who come to rely on the technology will not fully understand how its outputs are achieved—and, as such, not trust it. It will be important for organizations to identify risks early and understand how those risks may impact building trust in the use of generative AI.

Overcoming management skepticism could be a particular obstacle. It may be hard to get senior executives to believe developers will not blindly follow their AI models and leave their organizations open to potential missteps triggered by what the executives may see as black-box technology platforms. Developers and their bosses will need to sell senior leadership not only on the technology but also on the safeguards to be built around it.

All that said, generative AI will inevitably become more accurate and reliable over time. What's more, the entire software development lifecycle process in a modern enterprise is already designed to find and fix bugs caused by mistakes developers have made. Errors attributable to generative AI can be mitigated by the same processes. With additional training of the technology, along with monitoring to ensure that outcomes align with expectations, we believe businesses will build trust in its capabilities.

In the meantime, intellectual property issues around generative AI remain unresolved, too. These models are trained on open-source code, with many different types of licenses, and it remains to be seen what will happen if software they generate is deemed too similar to open-source code.

Residence considerations for intellectual property are a potential sticking point, too. A large bank concerned about trade secrets, for example, may be reluctant to share its data with a start-up generative IT vendor operating in a multitenant cloud environment. This could create a market for a larger vendor—perhaps Microsoft, Oracle, or IBM—to offer a single-tenant product more sensitive to those concerns. Microsoft, which invested in OpenAI in 2019 and again in 2021, has since announced an additional multiyear, multibillion-dollar investment into the organization.

As noted, costs will be a stumbling block for some organizations over the near term, especially those seeking access to state-of-the-art generative IT models like OpenAI's GPT-3.

Ultimately, we see generative AI becoming an invaluable addition to the advanced analytics and AI toolbox. However, organizations will need to develop governance protocols around the responsible use of AI. The aim will be to ensure the technology is used in line with their standards, with consideration of the impact on customers and society, and to build trust in AI-generated outcomes.

The importance of getting started now

Although already impressive, generative AI technology is advancing rapidly. Another large leap in capabilities is expected within this year. (Exhibit 2)

Even if organizations deem the technology insufficiently mature in its current state, doing proofs of concept now to become familiar with the technology is prudent and will help position them to respond quickly when more advanced models are released. With all of the challenges and risks posed by generative AI, the long term upside still remains extremely positive.

Exhibit 2: A fast ascent The power of generative AI models is increasing at a rapid rate. Graphic shows model size in billions of parameters for key generative AI models introduced since 2018.

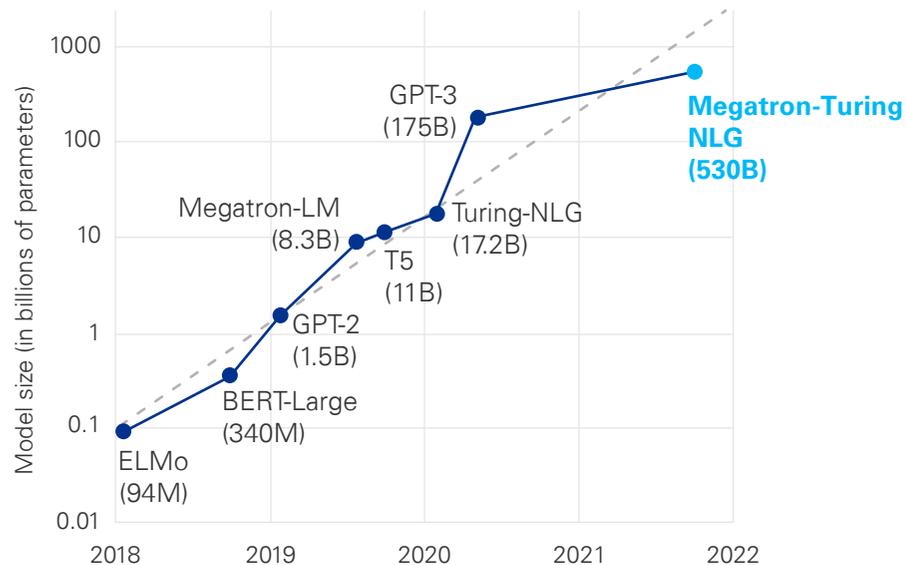


Chart courtesy of Nvidia. Source: Nvidia Developer Technical Blog, Using DeepSpeed and Megatron to Train Megatron-Turing NLG 530B, the World's Largest and Most Powerful Generative Language Model, Paresh Kharya and Ali Alvi, October 11, 2022, <https://developer.nvidia.com/blog/using-deepspeed-and-megatron-to-train-megatron-turing-nlg-530b-the-worlds-largest-and-most-powerful-generative-language-model/>

How KPMG can help

An early and enthusiastic advocate for the power of artificial intelligence, KPMG is positioned to help your organization leverage generative AI to transform software development and maintenance. Drawing on our deep expertise in machine learning and natural language processing, we can guide your organization through strategy development, platform selection, and implementation—and then provide ongoing support to help you optimize your investment in this transformative technology. We understand both the promise of generative AI and the process and cultural changes—including the embrace of responsible AI practices—that will be required to realize its full potential.

Contact us



Mark Shank

Principal, Advisory
Lighthouse
KPMG in the U.S.
mshank@kpmg.com



Kelly Combs

Director, Leader of Responsible AI
Lighthouse
KPMG in the U.S.
kcombs@kpmg.com

Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates or related entities.

The information contained herein is of a general nature and is not intended to address the circumstances of any particular individual or entity. Although we endeavor to provide accurate and timely information, there can be no guarantee that such information is accurate as of the date it is received or that it will continue to be accurate in the future. No one should act upon such information without appropriate professional advice after a thorough examination of the particular situation.

© 2023 KPMG LLP, a Delaware limited liability partnership and a member firm of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

DASD-2023-11580

kpmg.com/socialmedia

