

AI And Machine Learning

Turn Generative AI from an Existential Threat into a Competitive Advantage

by Scott Cook, Andrei Hagiu, and Julian Wright

From the Magazine (January–February 2024)



Matilda Temperley

Summary. By making it vastly easier and cheaper to improve or create products and services that previously required significant human labor and creativity, generative AI has the potential to disrupt or even commoditize many businesses. Some companies will be able to gain an... [more](#)

By making it vastly easier and cheaper to create or improve products and services, generative AI has the potential to disrupt or even commoditize businesses whose offerings previously

required significant human labor and creativity. To compete, many businesses will have to rethink their strategies and find new ways to add value for customers.

What can firms do to turn generative AI from a threat into an opportunity? How can they use it to build competitive advantage? What types of businesses have the greatest potential? These are the questions we will answer here.

How to Leverage Generative AI

Drawing on our decades of combined experience working with Silicon Valley firms and researching competitive advantage from data and AI through economic modeling, we've found that generative AI can be implemented at three levels that correspond to an increasing degree of competitive advantage.

1. Adopt publicly available tools. Companies can provide employees with off-the-shelf large language models (LLMs) or other types of generative AI, such as Midjourney and Runway, to help them do their jobs more effectively and efficiently. For example, employees can use level 1 tools to improve internal communication or presentation documents, generate ideas for new product features, conduct research on competitors, write customer-outreach messages for salespeople, touch up photos in marketing materials, create social media posts, and so on.

Virtually all firms in all industries—from technology firms to consumer product manufacturers to architects and consultancies—should look at implementing level 1. Just as the internet became essential to surviving and thriving in the late 1990s, so too will publicly available generative AI tools as they become more accurate, applicable, and secure. These include not only general LLMs like OpenAI's ChatGPT and Google's Bard but also industry-specific ones, such as Alexi for law firms and Karbon AI

for accounting firms, and function-specific ones, such as Copy.ai for copyediting and GitHub Copilot for coding. Firms that do not implement such tools will be severely handicapped.

In the short run, some businesses will be able to gain an edge by leveraging these tools better or faster than their competitors. Firms that also drive fundamental business-process change based on using generative AI will achieve an even greater edge; those that are too slow will fall by the wayside. However, relying solely on AI models that everybody (new entrants and existing rivals alike) has access to will bring only temporary advantage before becoming table stakes.

In its basic form, level 1 implementation involves no customization of the AI model. In practice, businesses might achieve some small measure of customization by fine-tuning the models using their own data, a move that advances a company toward level 2.

2. Customize the tools. Businesses can create their own customized generative AI tools that use the data and know-how they've accumulated in the process of serving their customers, either by building them from open-source models or by using models provided by LLM companies such as Anthropic and Cohere. These tools can enhance the customer experience by making products easier to interact with and adding new capabilities, like personalized recommendations.

For instance, providers of software products (especially complex ones) can create product-specific, generative-AI-based chat interfaces. As Des Traynor, cofounder and chief strategy officer of Intercom, a company that provides online tools for customer service, told us, "The beauty of these interfaces is that they free providers from having to trade off ease of use for customizability." Using the software product becomes intuitive for everyone (not just experts), and it can be automatically personalized for a variety of use cases. Instead of having to navigate a maze of drop-

down menus and know exactly what data to enter where, people can speak to the chat interface in natural language, as if they had a personal assistant.



Matilda Temperley photographed big-cat trainer Thomas Chipperfield and his lions. His family has been working with circuses for 300 years, dating back to the Frost Fair, in 1683. His method of training is built on encouragement and compassion.

Consider an enterprise software maker that offers a product for managing travel expenses. It could add a generative AI tool, trained on data from how employees previously submitted expenses, that covers a range of requests by people in various roles. Employees would simply tell the AI that they wish to submit travel expenses for a recent business trip, and it would guide them through the process, indicating which receipts to upload and what other information it needs. It would also allow them to note any special circumstances that should be taken into account—for example, “I took vacation days on Wednesday and Thursday” so that it knows to exclude expenses and per diems on those days. The tool would then process reimbursements in accordance with the company’s policies. If any complications arose, it would ask users for additional information or flag discrepancies.

Or consider banks and other financial service companies, which could implement level 2 by creating generative AI private bankers. These would be trained on the recommendations made by human private bankers to existing clients before being rolled out to clients looking for investment advice customized for their circumstances and needs.

At its simplest form, level 2 implementation consists of combining a generative AI tool with internal data to provide customers with a more intuitive user interface and more-personalized services. Many businesses will be able to keep improving their customized model by incorporating user feedback. This starts moving them in the direction of level 3.

3. Create automatic and continuous data feedback loops. When fully implemented, generative AI tools produce reliable signals in the natural process of customers' using the product or service, which are automatically fed back into the model to improve its capabilities with minimal human intervention. Companies that allow the AI tool to infer how helpful it has been to customers in achieving their goals will garner the most reliable signals.

Creating a feedback loop that is unique to each firm's product or service is the holy grail. The more customers use the offering, the more feedback signals they generate, which allows the generative AI tool to further improve itself, leading to more users, more usage, more feedback, and so on. The result is a powerful form of compounding competitive advantage.

Consider how Chegg, an online education company, is using generative AI to create a personalized learning assistant for students. This custom-built online tool is trained on Chegg's proprietary library of educational content and data from usage of its products. Students interact with a conversational user interface, which can tailor guidance to their individual needs. The tool continuously improves its study content, such as flash cards and practice tests, by figuring out the main sticking points

common to users. This creates a data feedback loop across Chegg users. A feedback loop specific to each user is also created, enabling the tool to adjust the pace of instruction and the content and difficulty of practice tests in accordance with the individual's performance. In doing so, the tool helps students learn better while generating reliable feedback signals, which allows it to keep improving.

Purchasing generative AI capability from third-party providers will become similar to buying cloud computing services, with prices constantly falling and security increasing.

Companies in all kinds of industries can benefit from level 3 implementation. Publishers of online video games, for instance, can use generative AI to create and constantly enhance customized nonplayer characters (or NPCs, those not controlled by human players), using information on how users react to and engage with different NPCs during gameplay. Companies can also use generative AI to experiment with different screenplays and environments and learn, almost in real time, which ones increase user engagement.

Companies that own both streaming services and content production (such as Disney, HBO, and Netflix) are also good candidates for level 3 solutions. For instance, a company could develop a proprietary AI solution to customize its shows by adjusting certain scenes for different audiences. It could use a tool like Runway to produce content that better targets viewers' tastes, taking into consideration, for instance, tolerance for adult content or language, sensitivities of viewers in different countries,

or attributes of other shows that users watch. The model would update itself based on consumer behavior, such as whether customers continue watching after the adjusted scenes, as well as any direct feedback they might give, and the model would incorporate those insights into future projects. The feedback loop wouldn't be as strong and continuous as it would with video games, however, where learnings can be incorporated in real time.

Professional services firms, such as research and advisory consultancies or marketing agencies, should also consider going to level 3. A consulting firm could create a generative AI model that produces first drafts of research reports, for example. Employees would then make revisions to the early drafts, and that feedback would be continuously digested by the AI, which would use it to produce better initial drafts. (A caveat: Firms must ensure that the AI tool does not use one client's confidential data when engaging with another client, especially one in the same industry.) Such a tool would essentially turbocharge the learning process for firms and their professionals.

The main challenge in reaching level 3 lies in figuring out how to obtain feedback without disrupting the customer experience. Ideally, the signal should be generated during the natural course of using the product. Consider again Chegg's AI tool. It observes how well a student does on every practice test on a question-by-question basis. It gets immediate insights on where the student has problems—allowing it to change the way concepts are explained or adjust the difficulty and nature of questions on future tests—without negatively affecting the student's experience in taking the test.

When a natural and high-quality feedback loop is not attainable, as will be the case with many offerings, companies will need to resort to directly asking for customer feedback. This should be done in a minimally intrusive way while making it clear to

customers how they benefit from giving it (for example, “Please tell us on a scale of 1 to 10 how helpful our AI chatbot was so that it can learn to better serve you in the future”). And if asking for direct customer feedback is too disruptive or it provides an unreliable signal, companies could instead involve humans in the loop, as was described above in creating draft reports.

Implementation Considerations

Taking full advantage of generative AI requires ever-higher expenditures and technical expertise as companies progress through the levels. It also involves more-extensive adjustments to online products, services, and internal processes.

At level 1, the key issue companies must address is how much internal data employees should be allowed to share with publicly available generative AI tools. Whenever these tools use proprietary data, it raises data-security and competitive concerns; however, we are already seeing generative AI providers take steps to solve these problems—for example by ring-fencing each business customers’ AI and data.

At level 2, companies will need to focus on training and fine-tuning customized AI tools to ensure that their outputs are as accurate as possible (minimizing “hallucinations”) and are highly relevant to customers.



Matilda Temperley

As companies implement level 3, they must focus on redesigning their online products and services, as well as some internal processes, to seamlessly integrate generative AI throughout the entire customer experience. In doing so, they will maximize the breadth, depth, and quality of the feedback signals they extract.

At first glance, it might seem that only large firms will have the resources to achieve levels 2 and 3, but in reality, customization using proprietary company data is becoming increasingly easy and affordable. Companies can now obtain generative AI capability from providers such as Open AI (ChatGPT Enterprise), Microsoft (Bing Chat Enterprise), and Anthropic (Claude for businesses) and fine-tune it based on their own data. They can also be reasonably confident that their data will be kept secure and private and that learnings from it won't spill over to anyone else (to the generative AI provider or to competitors). Purchasing generative AI capability from third-party providers will become similar to buying cloud computing services, with prices constantly falling and security increasing. Alternatively, companies can build their own generative AI capabilities by adopting open-source models, such as Meta's Llama, and training

them on their own data. That will keep everything under the company's control, as is the case when doing cloud computing on one's own servers.

What Businesses Will Benefit Most?

As we've noted, virtually all businesses should adopt generative AI at level 1. Most should at least consider moving to level 2 and possibly to level 3. To determine whether doing so would confer an advantage, leaders should ask two questions:

How much of our current offering can be replaced by generative AI? This question determines the scope for potential disruption by generative AI, which, of course, indicates how urgent it is for a company to start moving up the implementation levels.

The impact of generative AI on companies that make simple consumer products such as cleaning supplies and kitchenware will be minimal. But even those companies can benefit from level 1 adoption by using the tools to increase internal productivity and creativity—for instance, to come up with new product ideas, designs, or formulations. (See “How Generative AI Can Augment Human Creativity,” HBR, July–August 2023.)

Software companies will find that the extent to which their products can be replaced by generative AI will vary a lot. All else being equal, complex software products are more likely to be affected, because a natural-language chat interface could make them a lot easier to use for a wider customer base. On the other hand, software that offers value through integration with other products or services is more immune to disruption. For example, the payment platform Stripe is built on complex and highly secure connections to financial institutions that are not easily replicable by generative AI.



The tools you need to craft strategic plans — and how to make them happen.

Sign Up

For online services currently provided by humans, the threat of disruption is high and immediate, because generative AI will be able to handle the bulk of the work. Examples include logo- and website-design services (such as Webflow, Fiverr, and Upwork) and online education (Chegg).

For firms where after-sales service and support are important, generative AI can play a key role in building competitive advantage. Specifically, it can be used to supercharge customer service, which is a central part of the value proposition for users looking to get the most out of the products or who face technical challenges in using them.

If the threat of disruption is high, businesses obviously should move as quickly as possible to levels 2 and 3. If it is not, they can take a slower approach to implementing generative AI.

What is the scope for improving our competitive position using generative AI? Businesses will be able to use generative AI to enhance the value they provide customers only if there is sufficient potential for them to leverage their own data to build a model that surpasses what can be achieved using publicly available AI tools like ChatGPT or Midjourney. To gauge that potential, companies should ask three additional questions:

How idiosyncratic is our firm's data? To provide an advantage, a business's internal data must not be easily replaceable by alternative, publicly available data that can be used to train similar AI models. For example, website- and logo-design companies are unlikely to amass specialized data that would confer an advantage greater than what competitors could achieve

using data derived from public websites. By contrast, the value Chegg offers its students can't be replicated by rivals using public LLMs. That's because its data comes from the unique actions of its students as they take quizzes, interact during instruction sessions, ask questions when they are confused, and so on.

Internal data is usually more valuable when it is idiosyncratic to the relevant product or industry vertical and to the specific customer need. Such data produces recommendations, content, and other output that cannot be replicated by generalist AI tools trained on publicly available and historical data. In such cases, we agree with Dan Rosensweig, president and CEO of Chegg, who told us, "True verticalization should win." In other words, companies that gain specialized insights from highly segmented customer groups will have a distinct advantage.

How reliable is the feedback we get from customers? Businesses that can closely observe the extent to which their products help customers reach their ultimate goals are in a strong position to leverage generative AI at the highest level. Video game makers receive very reliable signals by observing people's gameplay and can use that information to determine the extent to which AI-powered elements heighten engagement. In other cases, companies must make do with less-reliable signals, such as click-through rates. Consider Adobe's Firefly, a generative AI tool integrated into Adobe Photoshop and Illustrator for creating and manipulating images using text prompts. The only signals Adobe Firefly has for gauging whether the tool is serving customers well are how frequently each customer uses it and how many times each person tries to generate different images, which are noisy measures of satisfaction.

Note that the same type of feedback may be more reliable for some products and less for others. For instance, thumbs-up/thumbs-down is a much more useful indicator of whether

customers liked a movie recommendation than of how useful they find a particular feature of an enterprise software product.

How costly is it to get reliable customer feedback? Some companies will be able to obtain a very reliable signal by simply observing their customers as they use their products. Other businesses will need to engineer customer feedback loops, which can be costly for the company or its customers. (See “To Get Better Customer Data, Build Feedback Loops into Your Products,” HBR.org, July 11, 2023.) A company may have few opportunities to observe useful customer feedback, or privacy rules may prohibit it from making use of customer activity to train its model. Getting around such issues by asking customers to give feedback during usage or having employees act as “humans in the loop” in order to amass training data for the generative AI tool can be tricky. The first approach can annoy customers, and the second can drain company resources. When obtaining reliable feedback is onerous or costly, businesses should not attempt to achieve level 3.

...

Given the expected widespread adoption of generative AI technologies, the firms best placed to achieve durable competitive advantage are the ones with access to unique customer data that can be continuously replenished via self-reinforcing feedback loops. This has important implications not just for managers but also for investors, who should focus on the companies that have the potential to reach level 3 implementation.

To realize the opportunities that generative AI is creating, however, requires proper execution, which is where some incumbents may fall short. CEOs and senior leaders must make sure that this technology is treated as a fundamental part of company strategy, not just a technological issue to be delegated to IT. In that sense, it’s different from cloud computing, where the decision of how much to rely on external cloud providers versus build in-house hosting capability is not quite strategic. Generative AI, by contrast, can and should affect the customer value

proposition directly. Even if generative AI is not an existential threat, leaders should push their organizations to add capabilities preemptively to make sure they don't fall behind.

A version of this article appeared in the January–February 2024 issue of *Harvard Business Review*.

SC

Scott Cook is a cofounder and the chairman of the executive committee of Intuit.

Andrei Hagi is an associate professor of information systems at Boston University's Questrom School of Business.

 @theplatformguy

Julian Wright is the Lim Chong Yah Professor of Economics at the National University of Singapore.