Developing a Generative
AI Strategy

BY NITIN MITTAL AND TOM DAVENPORT

When new technologies enter the business realm, they catalyze interest in their potential to create efficiencies, drive innovation, and shape the businesses of tomorrow.

With the arrival of generative AI and the immense potential it promises, organizations stand on the brink of dramatic transformation.

Business and public awareness of generative AI has grown dramatically in recent years with the release of large language models (LLMs), which startled many with their capacity to take in natural language prompts and output coherent, conversational replies. Applications that access LLMs (e.g., ChatGPT) captured the imagination, but generative AI is not simply a sophisticated chatbot representing another incremental step in the trajectory of technology. Instead, it is a business-disruptive force, bringing the automation of human cognition into the workforce and throughout the enterprise.

Given the potential, companies across industries are making significant investments in generative AI to deliver business outcomes and competitive differentiation. The race is on, as it seems likely that businesses with a first-mover generative AI advantage will capture a disproportionate share of the market. Still, preparation and caution are virtues, and to orient an enterprise for a future with generative AI, what's needed is clarity on the capabilities and understanding of the business value, as well as a strategy for developing and deploying generative AI use cases.

WHY GENERATIVE AI IS DIFFERENT

Generative AI is a subset of AI characterized by its capacity to take a human input (e.g., a prompt that is a question, gives instructions, or includes code snippets, etc.) and output coherent responses or solutions referencing data at scale across a range of modalities, including text, images, audio, video, code, and 3D models.

It is different from other types of AI in three primary ways. First, it does not require human agency to progress, in that the model can learn by itself over time. Second, it can mimic human cognition in a way no other AI can. And third, the input is not a program language. Rather, the input is natural language, and the output modality mimics human language and the ability to imagine.

Importantly, there is enough variability in the outputs that they can approach a kind of creativity that is unconstrained by the human limitations of speed and scale. For example, a user could ask a generative AI model to create an office layout design that aligns with the business branding and mission statement. The model output could come in the form of an image, video, or even 3D rendering of a space, and the model could also be asked to explain how the branding and mission statement are reflected in the output. This is not just a matter of creative thinking. The model could be prompted to generate tens or even hundreds of designs at a speed and scale exceeding human capacity.

Today's moment with generative AI owes to a convergence of advanced hardware availability, AI capabilities, the integration of these capabilities across an organization's tech stack and, crucially, the ability of non-technologists to directly engage with generative AI tools. It has the potential to disrupt nearly every aspect of the organization, with powerful use cases that can transform operations, influence customer behavior, and create entirely new business models, products, and services. To help inform strategies that guide organizational investment and to understand the business value generative AI can create, organizations should begin with clarity on the archetypes for generative AI application.

SIX ARCHETYPES FOR GENERATIVE AI USE CASES

The transformational impact from generative AI will grow out of these six application archetypes:

Consumer-facing, general purpose. These generative AI models are the popular tools that are built and trained on internet data and made available to the public for general purposes (e.g., LLM-powered chatbots such as OpenAI's ChatGPT or Google's Bard).

Productivity and personal use. These generative AI models augment individual productivity, whether while working or for personal benefit. This may take the shape of use cases such as writing assistants, document query capabilities, automated summaries, and virtual assistants (e.g., Microsoft's 365 copilot).

Business software usage. Generative AI models can be trained for a specific application and integrated with existing enterprise software platforms and solutions. For example, enterprise resource planning (ERP) systems that manage daily business operations can be augmented with generative AI, delivering natural language navigation of complex interfaces, generating financial documents, and autonomously addressing identified issues while escalating critical scenarios to a human in the loop.

Software engineering. Generative AI can be used to autonomously develop software code for business applications. This archetype helps IT teams accelerate their work and address complicated tasks such as configuring and maintaining code across platforms and annotating code for human developers (e.g., AWS CodeWhisperer).

Business-to-consumer applications. Generative AI models can permit new ways of personalizing and targeting customer engagements. Using generative AI to enable hyper-personalized communications, product recommendations, and customer support can have a significant impact on customer satisfaction, retention, and value. An example of this archetype is found in the partnership between NVIDIA and WPP to develop a generative AI-enabled engine for creating commercial content and brand advertising.

Purpose-specific use cases. Some generative AI-enhanced models are accessed or built by organizations for domain- or sector-specific use cases. Trained on enterprise data for discrete applications, these models leverage the capabilities of generative AI to address select challenges or opportunities unique to a specific enterprise. As an example, a large retailer may train and deploy a model to autonomously negotiate vendor contracts under a certain size, freeing up human capital to focus on larger, higher-impact vendor agreements.

With this understanding of generative AI archetypes, the enterprise is equipped to take the next step in its strategy development and inspect the value that could grow out of each archetype, modality, and use case. There is great potential for first-order business outcomes, including cost reduction, process efficiency, revenue growth, accelerated innovation, and the discovery of new ideas and insights. The benefits may be cross-cutting or specific to a business unit. Sales teams may see opportunity to unlock customer engagement in business-to-consumer applications; human resources professionals may identify use cases that enhance productivity; and IT teams may desire a tool to accelerate software engineering. Business units will voice varying needs and innovative ideas, and the applications are limited only by the imagination.

Conceiving the archetype and value in this way, enterprises are prepared to make informed decisions around the approach to enabling these models (e.g., build in-house versus acquire from a third-party and customize in a secure environment). To be sure, like all AI, generative AI raises questions around risk, governance, and compliance. As such, an informed generative AI strategy necessarily will account for the trustworthiness of the tools, methods, models, and frameworks that are deployed.

CONSIDERATIONS FOR RISK AND TRUST

While enterprises today are still evaluating and experimenting with generative AI, first movers are heading toward production and deployment at scale. This subtype of AI is still relatively new, and along with exploring the capabilities, organizations are learning to navigate the risks around things such as

> "While generative AI outputs are susceptible to inaccuracies (aka "hallucinations"), companies are discovering ways to mitigate these limitations by leveraging supplemental technologies and process guardrails."

security, privacy, and reliability, as well as the implications for AI governance and regulatory compliance.

Some of the concern around generative AI may be rooted in risks associated with publicly available LLMs and perhaps less so with models trained on enterprise data in secure environments. While generative AI outputs are susceptible to inaccuracies (aka "hallucinations"), companies are discovering ways to mitigate these limitations by leveraging supplemental technologies and process guardrails.

Ultimately, trustworthiness is an essential component of a value-driving generative AI strategy, and as such, use cases need to be considered across the domains of impartiality, responsibility, accountability, privacy, transparency, explainability, and safety and security. A framework can help in this regard. Deloitte leverages its Technology Trust Ethics framework internally and with clients to evaluate the risks associated with emerging technologies and make decisions on whether and how to implement tools that unleash capacity, efficiency, and growth. Assessing risk informs necessary changes to processes, human capital training and hiring, and technology evaluation and acquisition. These important changes enable AI governance, which in turn supports compliance and trustworthy technology. Every generative AI use case is different, and each should be assessed and reassessed with a mind toward risk mitigation.

MOVING FORWARD WITH A GENERATIVE AI STRATEGY

With emerging technologies, the road to adoption can be unclear at first. However, by viewing generative AI through the lens of archetypes for application, businesses can begin to craft a strategy for building, adopting, and deploying this transformational technology.

These are still the early days of generative AI, but the technology is maturing quickly. The enterprises that can seize the first-mover advantage with generative AI can find powerful differentiators and outsized market capture. A strategy that accounts for the archetypes, outcomes, and risks can set the organization on the path to realizing the greatest potential from this business disruptive force. With this, the core question shifts from "What can this model do for my business?" to "What can I do with this model by training it on enterprise data in a secure environment with trustworthy guardrails in place to generate something that is meaningful in my business context?" AQ

Nitin Mittal is a principal with Deloitte Consulting LLP and the U.S. AI Strategic Growth Offering Consulting leader.

Tom Davenport is the bestselling author of Competing on Analytics *(Harvard Business Review Press, 2017) and* Big Data at Work *(Harvard Business Review Press, 2014).*

About Deloitte

MCE / KEY FACTS

MCE is your reliable partner for continuous success with agile people development solutions.

10,000,000

participants on AMA & MCE programmes in the last 10 years

92%

of Fortune 1,000 companies are our business partners

96%

of participants report they are using what they have learnt at AMA & MCE

1,000+

expert facilitators globally

100+

Open Training Programmes running throughout EMEA

98

year's experience working with our clients around the globe