Global Energy & Materials Practice

Beyond the hype: New opportunities for gen AI in energy and materials

Generative AI can create additional value from other forms of AI and analytics—and the energy and materials sector is uniquely well-positioned to benefit from these advancements.

This article is a collaborative effort by Filipe Barbosa, Soenke Lehmitz, Jukka Maksimainen, Lapo Mori, Bryan Richardson, Erik Roth, Humayun Tai, Sapna Thakker, Ian Wells, and Rodney Zemmel, representing views from McKinsey's Global Energy & Materials Practice.

February 2024

It's nearly impossible to scroll through daily headlines without encountering commentary on generative AI (gen AI)—the latest frontier of artificial intelligence. It seems as if every Silicon Valley personality, venture capitalist, or casual technologist is talking about ChatGPT or Bard, among dozens of other systems, and the potential these tools have to unlock possibilities far beyond imagination.

How closely should leaders pay attention to the hype? This isn't the first time that technology pundits have lined up behind the latest best thing. Should gen AI be dismissed as a fad, or should leaders double down on the latest tools as the panacea for their technical troubles?

The answer is likely neither. Our research shows that organizations that rely on innovation, data analysis, and process automation stand to benefit the most from gen AI. Within the agricultural, chemical, energy, and materials sectors, many companies are now moving beyond straightforward use cases and taking increasingly innovative approaches to adopting gen AI, and estimates show that an additional $390 billion to $550 billion of value can be created in the years to come.

Harnessing the power of gen AI

Gen AI's potential to accelerate growth and reduce costs cannot be ignored (exhibit). This is particularly true for the energy and materials space, which relies heavily on data and analytics for innovation and comprises sectors built upon increasingly nuanced and complicated processes. Simply stated, gen AI adds intelligence to any data, which can then be used to inform decision making—potentially reducing long processes to a single question—and it enables workers to gain previously unknown

Exhibit

Generative AI could create additional value potential above what could be unlocked by other AI and analytics.

AI's potential impact on the global economy, $ trillion

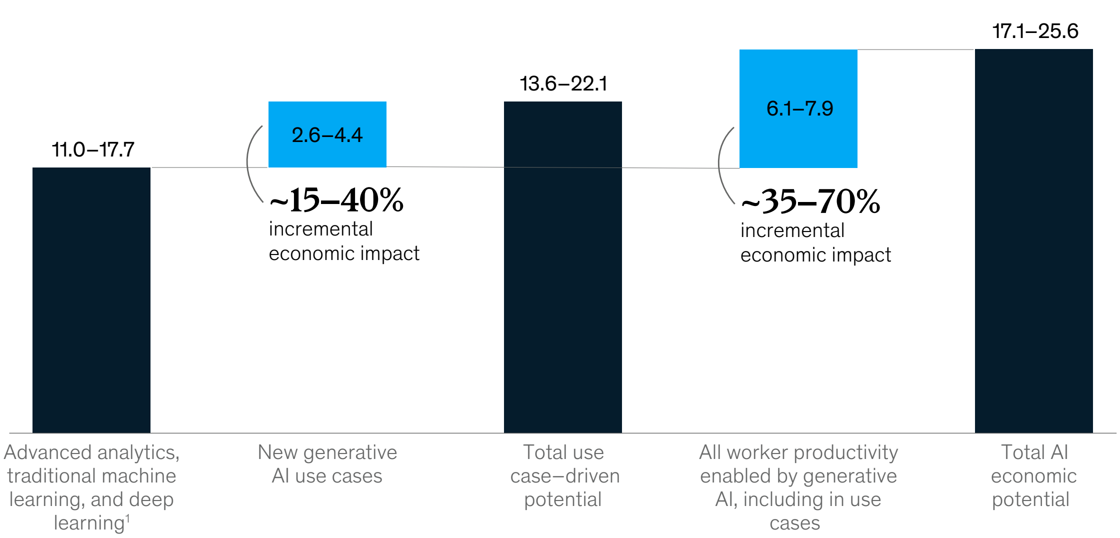

[1]Updated use case estimates from "Notes from the AI frontier: Applications and value of deep learning," McKinsey Global Institute, April 17, 2018.
Source: "The economic potential of generative AI: The next productivity frontier," McKinsey, June 14, 2023

McKinsey & Company

knowledge or capabilities. With this in mind, the growing list of exciting and nontrivial potential use cases in mining, oil and gas, chemicals, agriculture, power, and materials is ample reason for leaders to seriously consider gen AI (see sidebar, “What is generative AI, and why all the excitement?”).

However, this promise can only be realized if there is a clear vision of how to harness the power of gen AI—and understanding how to cut through the noise can be difficult. Industry players will need to take a hard look at how gen AI fits within their current digital strategies. This includes whether the organization has the digital capabilities to enable these technologies, whether to roll out commoditized solutions as they become available, or whether to design something totally new and exceedingly ambitious. Leaders will also need to understand the risks that gen AI entails and how to manage them in a way that ensures the organization is protected.

Establishing these capabilities early and securing a foothold in gen AI now could give companies an advantageous ability to adopt more-advanced models rapidly in the future; large language models (LLMs) are expected to exponentially increase in size and power over the next two to three years. In fact, cutting-edge models are already showing substantial improvement compared with models from the first few months of 2023, expanding the range of technically feasible use cases. As with traditional analytics—and digital technologies before those—organizations will likely choose to view gen AI not as the destination but rather as a powerful new tool to enable the organization’s full potential.[1]

What is generative AI, and why all the excitement?

Like many perceived overnight successes, generative AI (gen AI) has actually been around for years. Although OpenAI’s ChatGPT, Google’s Bard, and other large language model (LLM)–based tools burst onto the scene in late 2022 and early 2023, they all have common origins in advancements in deep learning, which have been familiar to research scientists for the past few years.

At the risk of oversimplification, gen AI refers to algorithms that can be used to create new content or synthesize existing content from large quantities of audio, code, images, text, or data sequences.[1] LLMs, such as the one used to power ChatGPT, can learn and reproduce relationship patterns in data that are more complicated than previous state-of-the-art models, showcasing abilities to generate compelling content seemingly from whole cloth.

[1] For more on gen AI and the differences between machine learning and AI, see our previous explainer, “What is generative AI?,” McKinsey, January 19, 2023.

The energy and materials sector is uniquely positioned to take advantage of gen AI

Sophisticated heavy industry has come to rely on data and analytics to push into the next frontier of efficiency. Oil and gas, agriculture, electric power, chemicals, and the materials and mining sectors are uniquely positioned to harness the power of gen AI to transform parts of the business, in terms of both their back offices and cross-functional cases and their core business and operations.

This reliance on analytics comes with huge amounts of data. In fact, nearly every modern plant, mine, or farm has years of data in sensor historians, as well as databases for failure modes and effects analysis, engineering reports, work orders, and maintenance logs detailing daily operations. Resource exploration and extraction comes with terabytes of electromagnetic and seismic measurements. And OEM manuals and troubleshooting guides fill dusty shelves in storage rooms.

This preponderance of structured and unstructured data is ripe for exploration and analysis via gen AI and is proprietary enough to offer a specific advantage to companies that move to exploit it.

[1] For more, see “The economic potential of generative AI: The next productivity frontier,” McKinsey, June 14, 2023.

In many of these industries, the asset intensity alone matters, as well as use cases that optimize asset utilization (assuming the data are available), compress processes, and predict outcomes over time. With this in mind, there are two categories of use cases that will likely apply to most subsectors within industry: those that are more straightforward and those that are considered "moonshots."

The straightforward use cases generally do not require a lot of technical expertise or specialization to deploy and will likely be rapidly commoditized. Examples include standard back-office functions, such as virtual assistants that automate administrative functions or customer-facing chatbots, or even "copilots" for software developers and data scientists, which can provide a substantial boost to productivity but which have only recently been available off the shelf.

By contrast, the moonshot use cases are more innovative and consequently require more customization—sometimes even a trained-from-scratch LLM. And although moonshot use cases have the potential to deliver significantly more value, they also require substantial upfront investment in capabilities and infrastructure. As a result, the application of gen AI for these use cases can vary based on the particular nuances of each subsector as well as on each part of the value chain.

Examples of moonshot use cases in industry include the following:

— *Utilities.* Organizations with thousands of miles of transmission lines, pipelines, and other remote and sometimes inaccessible infrastructure often spend millions of dollars on asset integrity. Corrosion and predictive-maintenance models can be retrained with previously unusable, unstructured inspection records, improving performance. This includes the integration of many sources of data, including traditional records, such as past damage, visual inspections, and data from sensors on the asset itself. Here, gen AI can significantly improve the effectiveness of a core business function necessary for both continuity of operations and public safety. Other sources of data, such as drone, aerial, and satellite-based images can be substantially improved by gen AI–powered computer vision.

— *Oil and gas companies.* Specialized models, extended from image processing applications, can process, interpolate, and interpret expensive seismic data to identify key attributes (such as horizon tracing, fault location, or direct hydrocarbon classification). As a result, the amount of data needed for high-resolution exploration can be decreased while the quality of the results can be increased.

— *Mining companies.* Mines with fleets of complex and widely distributed machines in the field can power models with libraries of maintenance manuals, historical work orders, procedures, tooling inventories, and parts databases. Doing so can enable a powerful AI assistant for maintenance technicians, helping to streamline work and increase reliability. Although it may seem like a straightforward application of off-the-shelf models, special care must be taken to ensure the advice presented is correct and useful to skilled technicians, and integration with existing systems is necessary to achieve full value.

— *Chemical companies.* Vast chemical databases can be leveraged to create models that can predict properties of new chemicals, substantially decreasing the search space for a physical laboratory and quickening the pace of molecule discovery. Similarly, new synthesis pathways can be digitally prototyped, helping to solve for low-cost, low-energy, or low-carbon emissions.

— *Agricultural companies.* By mining data on weather, soil conditions, pest pressure, and more, agronomy companies can build virtual advisers powered by gen AI. These advisers can identify personalized risks and opportunities for growers, farm managers, and farm operators, which they can access 24/7 at a low cost. Gen AI can also synthesize wide arrays of data points to generate testing scenarios for analytics programs, which can allow agronomy companies to simulate different events and give more accurate outputs.

Such moonshot applications are designed to be aspirational. However, forward-thinking players have already begun developing some of these use cases, among other examples, recognizing the substantial potential for value in their core business activities.

How industry leaders should think about prioritizing and implementing gen AI

The challenge in the years to come will be implementing use cases that drive meaningful value for the organization. On this point, a common pitfall for companies engaging with new tech is to kick off multiple pilots in several areas of the organization without first designing a comprehensive digital strategy. Leaders should instead be laser focused on spending the necessary time and resources on use cases that are both high impact and feasible. These bets are likelier to produce real results, feeding high adoption rates and increasing stakeholder support.

Companies should take care to consider if gen AI is the correct choice to help solve a given problem. Although models informed by gen AI have provided novel benefits in many functions and a clear boost to the status quo in others, the problems it can help solve weren't always previously intractable, and "traditional AI" solutions (such as simpler forecasting, system modeling, or optimization applications) were often more than suitable—and may even be more appropriate going forward. There is still impact on the table for many industrials that can be captured more simply in, for example, common process-optimization use cases before gen AI is an appropriate choice.

Although off-the-shelf models are easy to develop and deploy, they can limit the organization's ability to differentiate itself in the market. As a general rule, core operational use cases will likely need major customization in industrial sectors. This is due in part to the highly complex and technical nature of industrial processes. For example, recognizing and developing a chemical formula requires a tailored model infrastructure and specific database of often proprietary data to ensure the model provides accurate results. Similarly, almost all applications will require additional deep operational and vertical industrial knowledge to be made accurate enough to be useful, as well as involvement from key end users to ensure use cases are capable of addressing real pain points in existing processes.

That said, industry leaders thinking about prioritizing and implementing gen AI can define the core elements to make the digital strategy successful and understand the potential risks.

Core elements to make the digital strategy successful

Gen AI is only one facet of an organization's broader digital strategy, just as LLMs are only one of many models that companies can use to pursue new opportunities. To fully extract gen AI's potential value to an organization, leaders can consider the following points:

— *Business-led strategic road map.* The organization can build the strategic road map by understanding where the value is, what is achievable, and what can be logically sequenced to achieve real value quickly and incrementally. The senior team should be aligned to ensure that the road map and the resourcing matches the aspiration. And the path to create a tangible competitive advantage should be clear.

— *Talent.* A road map for in-house skills and capabilities will need to work in sync with the strategic road map. Gen AI is an extension of traditional analytics, and existing talent can be shifted or upskilled to implement a number of use cases. That said, particularly sophisticated or complex applications, such as those requiring a bespoke LLM, may require hiring for specific new skills. Data scientists, machine-learning engineers, and cloud architects will continue to be in high demand as new specializations emerge.

— *Agile delivery.* Gen AI use cases should be treated like any other digital project, meaning companies should build incrementally and quickly create a minimum viable product, learn lessons early, and adapt as things change. The pace of release for products or services as well as risk appetite for the organization is typically set, and the necessary "control functions," such as legal, finance, or risk, can also work in an agile way.

— *Technology and tooling.* Gen AI may require new assets that are not already set up, as well as access to new tools. Cloud capabilities will be necessary. The organization should have an MLOps (machine-learning operations) approach that allows the scaling up of AI in a safe and stable way.

— *Data management.* Gen AI is built on data first, and in this way, data are the real source of competitive uniqueness for industrial companies. Gen AI has the power to make data vastly more useful, but those data must first be available and reliable. Many energy and materials companies have already begun centralizing and de-siloing their data to support traditional analytics applications, and these efforts should be bolstered to enable gen AI use cases, as well. The organization therefore must balance managing data centrally versus data from the outside world. Doing so is critical because trust for a company's digital products is earned from its customers.

— *Adoption and operating-model change.* End users are engaged to help develop products that will accelerate, augment, or automate parts of their work. Business leaders should be held accountable for the adoption of technology.

Organizations should start answering these questions as soon as possible, but before rushing to encourage experimentation to test and learn gen AI, a better approach is to consider the above and stay laser focused on a few use cases that will deliver real value.

Potential risks of gen AI in industry

The potential risks for gen AI are significant and wide-ranging. The following risks are industry agnostic and should be carefully thought through before engaging with gen AI:

— *Accuracy.* Gen AI models can "make up" answers that seem credible (sometimes referred to as "hallucinations") or may have poor abstract reasoning. Responses that might seem high quality to executives may in fact be of little value to end users. For instance, consider a maintenance technician's need for highly specific, highly accurate advice related to replacing an engine component. The threshold of accuracy for usefulness is much higher when providing specific work instructions versus, say, offering troubleshooting advice.

— *Security.* Gen AI is susceptible to backdoor attacks, which can sometimes go unnoticed. There have been instances of hackers hijacking models, creating misinformation, stealing data, or committing fraud.

— *Privacy.* Competitively sensitive or confidential data can be leaked via public LLM APIs. Data available to LLMs may need to be compartmentalized to reflect access controls that already exist, particularly with confidential or classified data.

— *Fairness.* As with traditional AI, gen AI can sometimes produce biased outputs. It can also be misused to bypass intentionally added safeguards.

— *Legal.* There is potential risk of intellectual property infringement, copyright violations, and liability from misuse. How laws apply to outputs created by gen AI is still ambiguous. Many jurisdictions are still debating how to effectively regulate this technology.

Although these risks exist in all industries, the energy and materials sector should be particularly conscious of the potential risk of foundational models' initial lack of accuracy, especially given the potential implications of an imprecise output. Inaccurate model responses that jeopardize safety cannot be accepted, and careful safeguards for gen AI tools workers are using must be established. Model accuracy must be as high as possible while limiting the potential impact of an inaccurate output. Human operators should be kept in the loop and should be ultimately responsible for making operational decisions informed by gen AI. Hallucinations should also be managed. Potential mitigation could involve tuning performance and objectively measuring the accuracy of outputs in combination with other tools to prevent or flag hallucinatory responses. And teams developing tools on top of LLMs can enrich or engineer prompts using advanced techniques to improve accuracy of responses.

In addition, the sector should take note of the risks to privacy and take steps to protect proprietary, confidential, and sensitive data when using foundational models. Potential mitigation should include validating data and prompts to safeguard from malicious inputs, performing adversarial testing to filter responses, and deploying guardrails to protect and enhance data security features.

Many parts of the energy and materials sector have shied away from overhyped technological breakthroughs. But the speed of advances made in gen AI shouldn't be waved off as merely the "flavor of the month." In fact, executives in the sector should be first in line to explore how their organizations can benefit from gen AI. LLMs have the potential to deliver high value in industries that are technical and rely on complex processes. Moving forward, the question leaders should ask themselves is not "Where can I apply generative AI?" but rather "How can I apply AI to deliver transformational value?"

Filipe Barbosa is a senior partner in McKinsey's Houston office, where **Ian Wells** is an associate partner; **Soenke Lehmitz** and **Erik Roth** are senior partners in the Stamford office; **Jukka Maksimainen** is a senior partner in the Helsinki office; **Lapo Mori** is a partner in the Denver office; **Bryan Richardson** is a partner in the Vancouver office; **Humayun Tai** and **Rodney Zemmel** are senior partners in the New York office; and **Sapna Thakker** is consultant in the Toronto office.

The authors wish to thank Richard Bates, Rebecca Carothers, and Lydia Terisno for their contributions to this article.

Designed by Global Editorial Services