McKinsey & Company

Generative AI in the pharmaceutical industry: Moving from hype to reality

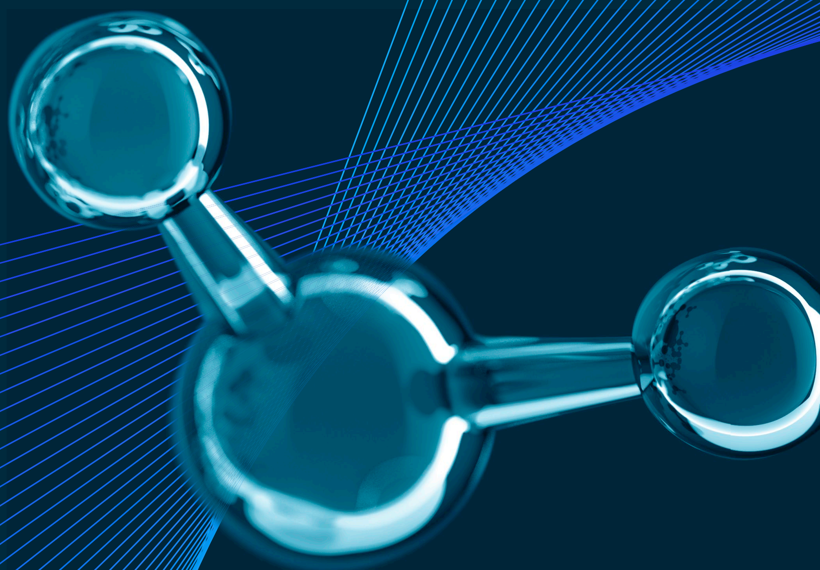

by Chaitanya Adabala Viswa, Joachim Bleys, Eoin Leydon, Bhavik Shah, and Delphine Zurkiya

January 2024

Accelerated drug discovery. More efficient clinical trials. Quicker regulatory approvals. Ultratargeted marketing materials generated in-house. Generative AI is transforming nearly all aspects of the pharmaceutical industry, revamping the way companies operate and potentially unlocking billions of dollars in value. The McKinsey Global Institute (MGI) has estimated that the technology could generate $60 billion to $110 billion a year in economic value for the pharma and medical-product industries (Exhibit 1), largely because it can boost productivity by accelerating the process of identifying compounds for possible new drugs, speeding their development and approval, and improving the way they are marketed.

Pharmaceutical companies, of course, have long been in the vanguard of artificial intelligence. Even before last year's explosion of interest, researchers were applying complex AI models to unlock the mechanisms of disease. AlphaFold2, ESMFold, and MoLeR, for example, all use deep learning to help predict the structures of nearly all known proteins, transforming our understanding of their underlying diseases.

With compelling new gen AI use cases emerging, we were eager to understand, in a more granular way, the most promising areas of potential value (see sidebar "A guide to our methodology"). So we dug deeper into MGI's data and modeling of 63 generative AI use cases in the life sciences and calculated the potential economic impact for five industry domains: research and early discovery, clinical development, operations, commercial, and medical affairs. Those numbers in hand, we tapped the experience and knowledge of McKinsey's leaders in each domain. These leaders used their extensive work helping clients deploy gen AI over the past year to identify the use cases most likely to spark meaningful near-term productivity gains and economic value.

Yet numbers tell only part of the story. The impending gen AI–driven life-science revolution promises unquantifiable effects on human health and well-being. An accelerated drug discovery process, for example, will help cure more diseases more quickly, opening additional resources that could then be applied to currently underserved areas. The ability to generate insights and patterns from vast quantities of patient data will spark more personalized treatments—and improved patient outcomes. Gen AI tools could also make patient care more consistent by reducing deviations in the manufacture and delivery of therapeutics. Finally, by automating tedious and time-consuming tasks like document creation and record keeping, gen AI stands to boost the productivity of researchers and medical liaisons, so they can better serve clinicians and patients.

Exhibit 1

Generative AI is expected to produce $60 billion to $110 billion in annual value across the pharmaceutical industry value chain.

Expected value annually, $ billion

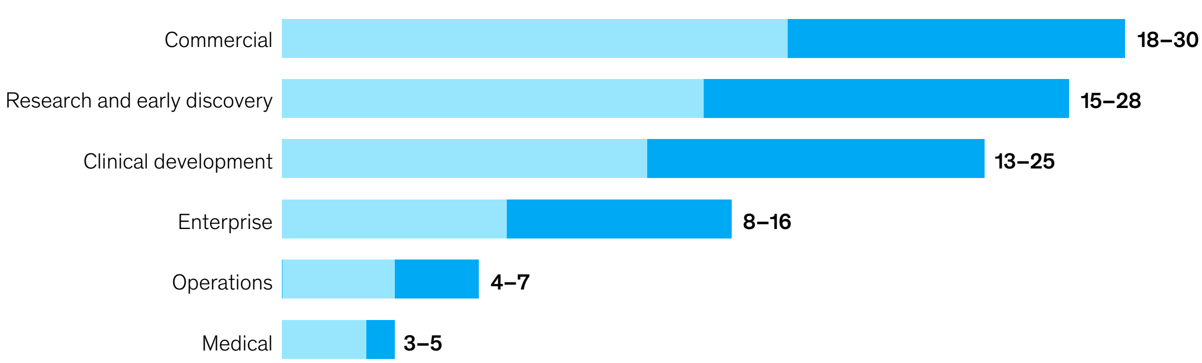

Source: McKinsey analysis

McKinsey & Company

Harnessing this powerful new technology, of course, will not be easy. Executives will have to grapple with tricky strategic decisions and operational challenges in an uncharted landscape marked by fast-changing technology and emerging risks. Here we offer guidance on getting started and analyze the most promising use cases and the elements needed for gen AI to transform them.

The way the industry works to treat disease is changing rapidly. Below, we offer a glimpse of the shape that new world will probably to take.

Moving from hype to reality

Before pharma companies can seize the opportunities generative AI presents, they must step back and understand exactly what it can and cannot do—in other words, differentiate the reality of gen AI from the hype that has come to surround it. Below, we debunk four of the most powerful misconceptions business leaders have about the technology.

Hype. Generative AI, on its own, will deliver the bulk of the value to be created.

Reality. This is a disruptive moment for the entire field of artificial intelligence, not for gen AI alone. Traditional analytical AI models, such as those currently used to promote stakeholder engagement and help diagnose diseases, will continue to capture value. The difference is that new gen AI applications will significantly enhance their capabilities..

Hype. Generative AI can easily be plugged into existing data sets to unlock key insights.

Reality. Gen AI cannot deliver results unless a proper data architecture is in place. Companies will need to build an intelligence layer that can understand issues such as molecular structures, clinical operations, and patient data. A multipronged approach will be necessary to create a data infrastructure that can run internal and external data sets. This is more than a purely technical matter: data scientists will need to collaborate closely with leaders in business strategy, medical affairs, and legal and risk to set priorities and execute strategies.

Hype. Selecting the right large language model (LLM) will be a key strategic differentiator.

Reality. Gen AI models account for only about 15 percent of a typical project effort, according to McKinsey research. Most of the work involves adapting models to a company's internal knowledge base and use cases. That is particularly true in the pharmaceutical industry, given the complexity of its data and the uniqueness of its regulations and technology. To succeed with gen AI, companies must integrate it across complex workflows to promote adoption and impact—a reality that highlights the need for effective change management. Indeed, McKinsey has found that 70 percent of digital transformations fail because of technical issues but because leaders ignored the importance of managing change.

Hype. Generative AI will instantly affect every part of the organization.

Reality. As with any digital transformation, leaders must apply an end-to-end lens and prioritize only the use cases and applications that make sense for overall business goals. Those leaders must create a strategic road map to ensure that they are optimizing the overall impact, the time to impact, and other important considerations. A "2x2 approach" is an effective strategy for companies getting started: begin with two use cases that require minimal disruption to the business, can build excitement across the organization, and have an impact most rapidly, as well as two other use cases that are potentially more transformational as longer-term goals.

Emerging use cases across life-science domains

Gen AI can create value across the entire business value chain through its ability to synthesize myriad sources of data, both structured and unstructured, and generate bespoke visual, textual, and even molecular content. Although the technology will affect all industries, it will have a particularly strong impact in pharmaceuticals. The reason is gen AI's truly multimodal nature: foundational models are built not just on language but also on images, omics, patient information, and other types of data—and these are all required to explain and solve the processes of diseases and how best to treat compression—the decreasing amount of time they have to capture a new drug's value. Over the past two decades, that time frame has fallen by almost 18 months: to 9.8 years, from 11.7, McKinsey research has found.

Most specific use cases will fall in one of four main categories: knowledge extraction, content and compound generation (for instance, generative chemistry), customer engagement (such as services for healthcare providers and patients), and coding and software generation. Organizations must move quickly to seize the competitive edge. Yet attempting too much too soon—for example, by launching multiple pilots without first drafting a clear road map to scale—can be problematic. In addition, leaders must understand that gen AI cannot be a peripheral initiative; it should be one of the top priorities. Given the conservative nature of many large organizations, C-suite executives must make a genuine commitment.

Below, we dive into 21 individual use cases that McKinsey domain experts regard as having the greatest potential for a near-term impact across five life-science domains (Exhibit 2). Many of these use cases cannot be realized unless some degree of digitalization is already in place, and not all of them will necessarily apply to all companies. While we recognize that gen AI remains an emerging technology not yet fully deployed at scale in most instances, we have also tried to estimate the potential impact for each use case. These estimates reflect what we have learned from our work over the past year, developments taking place throughout the industry, and interviews with gen AI experts and leaders.

Exhibit 2

Generative AI could propel holistic results in the life sciences sector in a number of ways.

Expected value annually (not exhaustive), $ billion

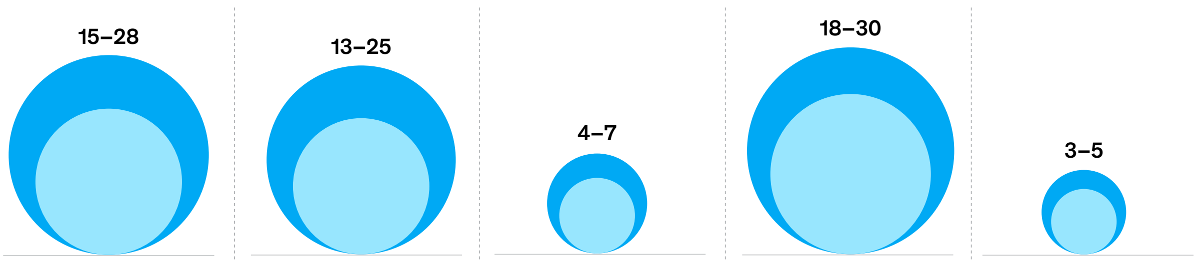

Research and early discovery
- Scientific knowledge extraction
- In silico compound screening
- Large molecule optimization and drug-vector design
- Indication selection for asset strategy
- Trial and portfolio optimization

Clinical development
- Trial performance copilot
- Smart data management
- Regulatory intelligence engine
- Major submission content writer

Operations
- Augmented sourcing
- Generative AI virtual assistant for manufacturing
- Reimagined investigations in quality
- No-touch planning and real-time inventory optimization

Commercial
- Personalized content creation
- Medical and legal review assistance and automation
- Customer-enablement copilot
- Strategic insights generation
- Patient experience optimization

Medical affairs[1]
- Customer-insight generation
- Sharper, more efficient medical writing
- Rapid summaries of scientific and medical literature

[1]Via efficacy gains on expenditures.
Source: McKinsey analysis

McKinsey & Company

Research and early discovery: A new generation of treatments

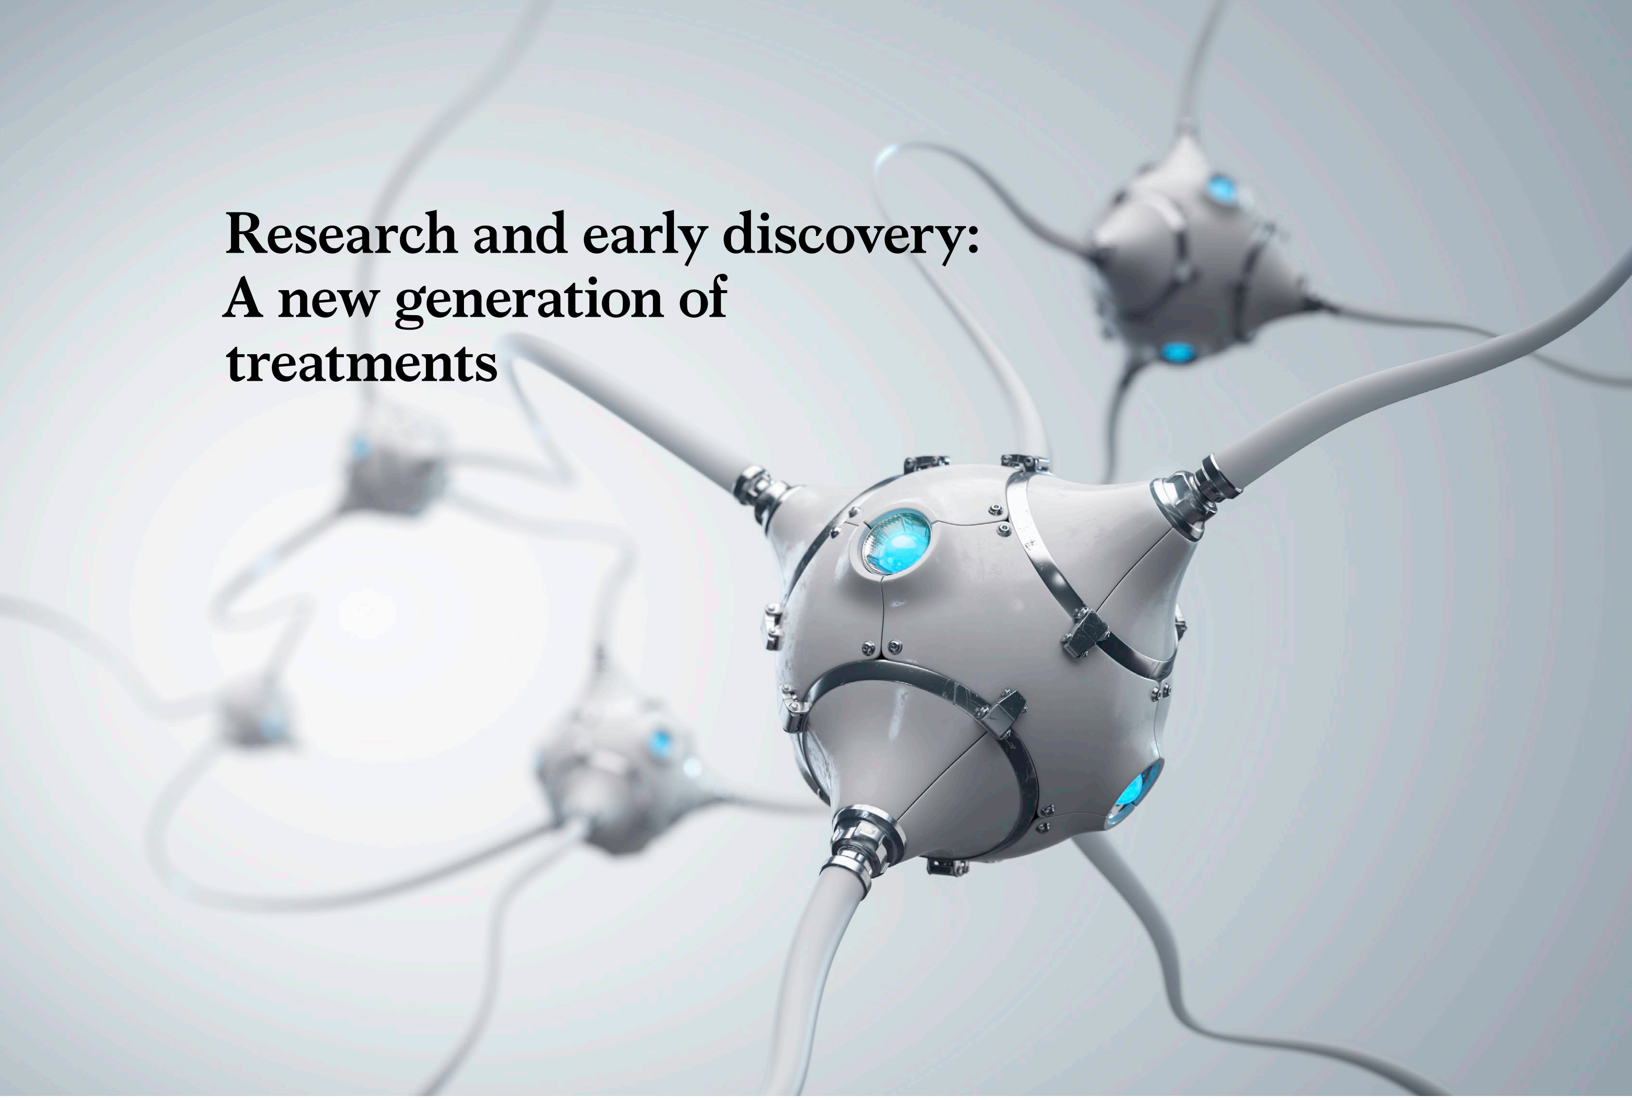

Potential opportunity: $15 billion to $28 billion

Our understanding of disease may be advancing rapidly. Yet the process of discovering and developing new drugs has been moving in the opposite direction: the average cost of bringing one to market is rising, not falling. Gen AI could help accelerate the process of identifying targets, developing validation assays to test compounds, singling out the most promising leads, and assisting in preclinical testing to determine their effectiveness.

Pharmaceutical manufacturers are already using foundational models for these purposes. In addition to natural-language models such as BioGPT and Med-PaLM, the models that researchers use include image models to analyze microscopy and pathology data, chemistry models to improve predictions for functional readouts of small-molecule data, large-molecule models for protein folding and predictions, patient journey models to focus development efforts on promising indications, and multimodal models to combine these modalities and thus enable in silico experiments. Augmenting such foundational capabilities with broader, more established forms of AI—such as computer vision, virtual screening, and knowledge graphs—can accelerate gen AI's impact across pharmaceutical research, perhaps cutting drug discovery time lines in half. We have identified five key use cases with a strong potential for a significant near-term impact.

Use case one: Extracting scientific knowledge. To better understand disease and drug targets, scientists spend much time extracting and summarizing information in documents such as patents, scientific publications, and trial data. That is not only arduous but also often provides incomplete or inaccurate information, given the sheer volume of data that must be processed. GPT-powered knowledge extraction—which uses AI algorithms to analyze unstructured data, including text, images, and other forms of information—

can alleviate this burden. Unlike earlier solutions based on natural-language processing (NLP), new gen AI tools offer a much deeper and broader understanding of both the medical context and intent. Researchers can therefore pose open-ended Q&As, easily shift between different tasks, and frictionlessly integrate additional evidence through prompt engineering. Little to no additional training is required to tailor information to specific use cases.

Potential impact: A more than 30 percent increase in initial manual assessments of drug targets

Use case two: In silico compound screening. Drug development can be hindered by the difficulty of identifying and prioritizing the chemical compounds that are most likely to successfully treat a particular disease and are thus most worthy of testing in laboratories. Gen AI accelerates the screening process with state-of-the-art foundational chemistry models that can map millions of known chemical compounds by their structure and function and overlay this information with known results for tested molecules (Exhibit 3). Like GPT-4, which is trained to predict the likely next word in a sentence, these models predict the next part (for instance, an atom) in the structure of a small molecule or a large molecule (such as an amino acid). Through many iterations, the model learns fundamental principles of large- and small-molecule chemistry. This knowledge can then be used to train bespoke machine-learning models that offer still more precise predictions—even in largely unexplored areas of chemistry—that companies can prioritize for subsequent screening.

Potential impact: An up to 2.5 times increase in the performance of chemical compound activity models; a more than fourfold boost in speed (from months to weeks) in the time needed to identify new leads

Use case three: Optimizing large molecules and drug-vector design. The tools described in the preceding use case help researchers to predict the disease-treating potential of small chemical compounds. An even thornier problem is doing the same thing for complex chains of molecules: for instance, proteins such as antibodies and mRNA. Although these molecules hold substantial promise for more targeted therapies and better vaccines, their complexity substantially increases development costs and discovery time lines. Again, next-generation LLMs can help by learning to predict the next substructure of large molecules (for instance, nucleic or amino acids) and generating insights about large-molecule chemistry. These insights can be used for the in silico design of new drug vectors and for predicting their efficacy in various drug discovery assays.

Potential impact: An acceleration of more than three times in large-molecule design

Exhibit 3

Generative AI may allow researchers to generate molecules with enhanced quality, safety, and novelty.

Chemist capabilities, present and future

	Present	Future
Drive innovation	Use traditional methods to identify new assets (eg, trial and error, personal knowledge, literature)	Use in silico analysis to proactively identify new assets and optimize their development
Synthesize compounds	Conduct chemical analysis and quality control check on products, relying on prior knowledge and education	Review reports created by generative AI to identify chemical trends and adjust analysis and synthesis methods accordingly

Source: McKinsey analysis

McKinsey & Company

Use case four: Indication selection for asset strategy. Gen AI's knowledge extraction capabilities (detailed in use case one, above) can also help researchers determine which conditions, or indications, to target with a specific molecule—one of the most important decisions facing biopharma companies. To make these calls, researchers must draw information from multiple sources, such as opinion leaders, literature reviews, omics analyses, trials data, and the activities of competitors. Yet given the vastness of this information, indication selections often cover only part of the available evidence base, so conclusions may not be optimal. Gen AI can help to address this issue by analyzing a wide range of structured and unstructured data sets.

For example, real-world data (RWD)—drawn from visits to doctors, insurance claims, electronic medical records, hospital data, and other sources—is often underused to select indications. With gen AI, foundation models that treat medical events as words and patient medical histories as documents allow researchers to uncover the semantic similarity of different events, and that makes it possible to estimate the biological proximity of one indication to another from a patient and clinical perspective.

What's more, information from molecular knowledge graphs can be tapped to reveal new connections (say, between entities such as proteins or human biological pathways) already identified in the literature or public data. These approaches can help uncover novel indications that can be rapidly validated through in vitro or animal models, increasing the likelihood of finding indications with a high probability of success and reducing the number of blind alleys (and their opportunity cost).

Potential impact: Successfully prioritizing indications that have gone on to be approved in clinics; eliminating those that would probably have failed; identifying entirely novel indications later validated in animal models

Use case five: Optimizing trials and portfolios. Once an asset has been matched with an indication, testing in a clinical setting begins. But identifying the appropriate patients to study is not easy, so clinical trials often include participants who may not respond to the treatment, and that can slow down its development.

One of the few fields to address this problem is precision oncology, where researchers now use biomarkers to stratify patients according to their probability of progression or to predict their response to different treatments. Gen AI will help other fields do the same thing, which will help companies assemble more diverse and representative clinical-trial populations. What's more, the use of models that examine genetic and phenotypic data alongside real-world data (drawn from sources such as patients' medical records) will allow researchers to better understand why different subgroups of patients respond differently to the same treatments. Finally, by using gen AI models built on data from medical-imaging techniques such as X-rays, CAT scans, and MRIs, scientists could even identify entirely new biomarkers: deeply hidden visual signatures of disease activity and severity that lead to unforeseen new treatments. The cumulative result: shorter, more efficient trials with a greater likelihood of success.

Potential impact: About a 10 percent increase in the possibility of success (PoS) for trials; about a 20 percent reduction in their cost and duration; time to approval accelerated by one to two years—all leading to a potential double-digit impact on the net present value (NPV) of assets or portfolios

Clinical development: Smarter trials, better data, quicker results

It takes, on average, ten years and $1.4 billion in out-of-pocket costs to bring a single drug to market, and about 80 percent of those costs are associated with clinical development, according to researchers at the Tufts Center for the Study of Drug Development. Clinical development is the process of bringing therapies from lab to patient by rigorously testing a potential medication’s safety and efficacy in human subjects, a process characterized by lengthy clinical-trial time lines and rigorous regulatory requirements. Gen AI addresses these pain points by increasing efficiency across the entire clinical-development process, unlocking economic value across three dimensions: up to 50 percent cost reductions enabled by the streamlining of clinical-trial processes and auto-drafting trial documents; a 12-plus month acceleration in the time it takes to conduct a trial; and at least a 20 percent increase in NPV, thanks to enhanced health authority interactions, quality control, and improved signal management. Across these three dimensions, we have identified four use cases with a strong potential for near-term impact.

Potential opportunity: $13 billion to $25 billion

Use case one: Trial performance co-pilot. Gen AI can rapidly analyze vast quantities of structured and unstructured data. It is therefore a powerful study-team companion, which can share insights and suggest effective interventions to improve the outcomes of clinical trials. Several best-in-class pharmacos have already created “clinical control towers”—advanced analytics platforms that support operational decision-making during clinical development by providing a single source of insights to speed up clinical trials. These AI co-pilots accelerate the trials by empowering study teams in at least three important ways: conversational AI capabilities provide tailored, actionable insights in an engaging format; smart alerts promote proactive, early interventions; and the automatic drafting of communications makes coordination with cross-functional team members more effective. These tools also promise to speed enrollment by automating analyses, proactively addressing enrollment challenges, and boosting collaboration.

Exhibit 4

Clinical operations managers could deliver trials faster and more efficiently with actionable insights and personalized engagement through generative AI.

Clinical operation manager capabilities, present and future

	Present	Future
Analyze trial performance insights	Manually collect and analyze data across multiple sources to understand trial performance	Leverage automated analyses and visualizations, tailored to trial characteristics including therapeutic area, phase, and next milestone
Act early and effectively	Intervene in a reactive manner based on subjective judgment and incomplete information	Receive proactive alerts from copilot, with data-driven flagging of risks and prescriptive recommendations of early interventions that can support on-time trial delivery
Personalize site engagement	Engage principal investigators (PIs) across sites with the same message and content	Automatically draft personalized messaging to engage PIs and site coordinators, with high engagement based on tailored content and timing
Prepare trial documents and reports	Manually prepare thousands of pages of documents	Refine auto-drafted documents with documents tailored for each trial, localized for each site, and translated for each country

Source: McKinsey analysis

McKinsey & Company

Potential impact: Approximately 20 percent cost efficiency improvement; 10 to 20 percent faster enrollment

Use case two: Smart data management. Data management today is a highly labor-intensive process, requiring manual trial-by-trial configuration of electronic data-capture systems, as well as detailed review and reconciliation of incoming patient data. By combining traditional and generative AI capabilities, data management can be automated across multiple steps. Databases can be created with one click based on the protocol, and case report forms can be auto-generated based on protocol, patient profile, and visit type. Data can then be reviewed and automatically cleaned in real time and queries intelligently and efficiently auto-generated based on trial context, patient status, and site actions—focusing attention on the most critical data gaps to be addressed before database lock.

Potential impact: 30 percent–plus cost savings for data management; a 50 percent–plus reduction in time to database lock; 70 percent fewer manual queries

Use case three: Regulatory intelligence engine. During the clinical-development process, pharma companies must answer questions and requests from regulatory agencies. These are known in the industry as Health Authority Queries (HAQs), and they often create bottlenecks that can delay the approval and market entry of new therapies.

Gen AI–enabled intelligence engines can help across three fronts: predicting potential HAQ patterns for a given submission; rapidly crafting appropriate sponsor responses; and providing deeper intelligence to submission strategies. Gen AI's predictive analytics, for example, can help teams proactively anticipate HAQs, thus reducing their number, both initial and follow-up. The insights generated will go on to inform regulatory strategy, risk management, and the broader R&D strategy.

Potential impact: Approximately 30 percent faster responses; 50 percent fewer HAQ follow-ups

Use case four: Major submission content writer. As the final step before regulatory review, submission writing must be done as quickly and accurately as possible to achieve or accelerate launch time lines. Yet drafting the clinical-study reports typically requires eight or more weeks to complete. Gen AI–based tools can cut this time almost in half by generating an "80 percent right" first draft from the underlying protocol, statistical-analysis plan, and tables, listings, and figures—within minutes instead of days—enabling rapid, collaborative iteration by cross-functional stakeholders through conversational prompting. Medical writers are thus freed up to focus on sections of the dossier that require a more complex clinical interpretation.

Potential impact: 40 percent faster regulatory submissions; a 50 percent improvement in cost efficiency across regulatory organizations; a two times reduction in quality issues

Operations: Reimagining the next horizon of operational performance

Potential opportunity: $4 billion to $7 billion

The pharmaceutical-operations value chain encompasses sourcing, manufacturing, quality, and the supply chain—and gen AI is expected to improve them all. First, the technology's ability to search and analyze large bodies of text, visuals, and other data sources will generate a wealth of new insights. Its content-generation capabilities will then enable teams to develop complex data representations—in text, visual, audio, and other formats—tailored to specific contexts. Finally, gen AI's conversational abilities, again in multiple formats, will make interactions with users more efficient and satisfying. We see four priority use cases across the operations value chain.

Use case one: Augmented sourcing. The procurement process tends to be inefficient and time-consuming: category managers must research and analyze market conditions, supplier information, and pricing data before evaluating RFPs or other submissions. A gen AI negotiation advisory bot can improve the process by generating first drafts of RFPs, purchase orders, invoices, and review responses. The technology can also aid negotiators by finding and analyzing relevant patterns in past negotiations and outcomes. Once suppliers have been brought on board, gen AI enables smarter contract management and the proactive monitoring of category and supplier performance.

Potential impact: A 5 to 10 percent reduction in procurement management costs; productivity gains (depending on roles and categories) of 50 to 80 percent

Use case two: A gen AI virtual assistant for manufacturing. Gen AI–based virtual assistants will help optimize drug manufacturing by quickly locating relevant standard operating procedures, automatically generating checklists and guides for repeatable right-first-time operations, and helping supervisors to monitor and manage line performance in real time. Virtual assistants also enable predictive maintenance—

averting shutdowns by flagging potential line failures, automatically generating intervention and troubleshooting plans and maintenance tickets, and optimizing repair and replacement schedules through planned maintenance.

Potential impact: A 10 to 15 percent improvement in overall equipment effectiveness (OEE) by reducing key losses; a 30 percent–plus increase in productivity for line leaders; a 15 to 35 percent workload reduction for maintenance technicians; a 5 percent reduction in quality costs through the detection of anomalies

Use case three: Reimagined investigations in quality. Deviation management is critical for all pharmacos, since they must adhere to good manufacturing practices (GMP) and stringent regulatory requirements. Identifying and capturing deviations is now an arduous manual process. Investigating them, for example, is a challenge given the limited availability of integrated data and cross-functional resources, so it is difficult to take effectual corrective and preventive action and thus to mitigate risk. Gen AI helps pharma companies reimagine their end-to-end deviation-investigation and -management process by providing tools to help clarify trends, the severity classifications of deviations, potential root causes, and corresponding corrective actions (Exhibit 5). All necessary reports can then be automatically generated and reviewed in compliance with corporate quality policies, thus making investigators more effective and productive.

Potential impact: A 35 percent–plus increase in productivity; a 30 to 40 percent improvement in the effectiveness of investigations

Use case four: No-touch planning and real-time inventory optimization. Particularly when raw materials are in short supply, inventory issues can have a substantial impact on production time lines. But adjusting supply and production plans in real time is tricky at best, requiring expert, on-the-spot judgment

Exhibit 5

Quality assurance specialists can use generative AI to produce faster corrective actions and reduce deviations with potential remediation measures.

Quality assurance specialist capabilities, present and future

	Present	Future
Identification of investigation requirement	• Manually search for deviations by sifting through thousands of deviation reports	• Conduct quick and effective searches of relevant deviations and create summaries using generative AI (gen AI)
Effective and efficient investigation	• Conduct investigations from scratch, with long lags between the event and the investigation's findings	• Expedite the investigation process • Improve quality of investigation output by leveraging gen-AI-driven summary of potential root cause and corrective actions
Report generation	• Manually write reports by gathering details about deviations, root causes, and corrective actions	• Leverage gen-AI-developed drafts of reports that can be reviewed and finalized
Standard operating (SOP) procedure development	• Manually write SOPs identified in investigations	• Use gen-AI-developed first drafts of SOPs to expedite the development process and improve quality

Source: McKinsey analysis

McKinsey & Company

to track orders, identify bottlenecks, and optimize networks. Gen AI–based planning tools address such issues by analyzing historical and market trends to anticipate spikes in demand, to predict supply chain bottlenecks and disruptions, to generate plans for proactive intervention, and to help draft production plans in real time by taking stock of available materials, current customer demand, and operational constraints. The tools also automatically monitor supplies to achieve optimum levels of inventory.

Potential impact: A 2 to 3 percent decline in supply chain costs; a 15 percent increase in the accuracy of inventory planning and forecasts; a 20 to 30 percent workload reduction for demand planners

Commercial: Better relationships with partners, providers, and patients

Potential opportunity: $18 billion to $30 billion

Any treatment, of course, works only if it is prescribed and taken correctly, which is why pharmacos spend considerable energy cultivating relationships and building trust with care providers, pharmacists, insurers, and patients. Gen AI tools can deepen all of these relationships. Internal brand leaders will be able to build on their proximity to healthcare providers, patients, and other partners to fine-tune campaign strategies in real time. Pharma professionals and patients will have a wealth of data at their fingertips to make faster, smarter decisions. Finally, the technology's content-generating capabilities will allow many aspects of the marketing process to move in-house. Outside advertising agencies will continue to play an important role but probably one focused on tasks at which they excel, such as innovation, creativity, media planning, and negotiations. Below, we present five use cases with a high potential for a strong near-term impact.

Use case one: Personalized content creation. Today, the creative and production process is almost completely outsourced to designers in agencies, where creators research and draft marketing materials based on content briefs provided by clients. Multiple iterations reflect feedback from marketers and from medical and legal reviewers (MLR). Current large language and multimodal models can streamline this tedious, resource-intensive process. Generative AI can, for example, standardize and accelerate the upfront creative design process while leaving room for innovation (Exhibit 6). In as few as five days, marketers can create first drafts and creative concepts that are ready to share with MLR reviewers before being passed on to agency partners for creative elevation and production. Gen AI can also help provide upfront campaign ideas for marketers to refine based on quantitative feedback from traditional AI and data analytics models.

Exhibit 6

Sales representatives who use generative AI can more quickly create personalized engagement plans.

Sales rep capabilities, present and future

	Present	Future
Collect and analyze data/ insights	Manually collect, analyze, and synthesize data and insights to prioritize/select accounts and healthcare providers (HCPs)	Use on-demand data and insights based on vetted, high-quality data sources and cross-function market access information to prioritize/select accounts and HCPs
Tailor an personalize messaging	Use advanced analytics and basic demographic and behavioral data to customize HCP messaging	Use generative AI to synthesize and curate brand data and HCP-specific data
Coordinate omnichannel	Use limited data to coordinate HCP outreach across channels	Utilize real-time insights around HCP interactions to ensure seamless coordination and optimization of sales efforts
Attend training and upskilling sessions	Rely on quarterly trainings with headquarters and meetings with regional and district managers	Augment established training cadence with real-time on-demand coaching, and just-in-time insights during HCP office/account visits

Source: McKinsey analysis

McKinsey & Company

Potential impact: A multifold increase in marketing efficiency; as much as a 30 to 50 percent reduction in content creation costs; an accelerated pipeline with a more than 20 percent increase in velocity

Use case two: Medical and legal review assistance and automation. The medical, legal, and regulatory process, undertaken to ensure that all advertising and promotional materials are accurate and comply with regulatory requirements, is a key component of life-science product marketing. At present, MLR reviewers manually review submitted content, highlighting new claims, identifying precedents, and tracking similarities across different types of data sources. Gen AI tools streamline the process in a number of ways: tracking the reuse of previously approved materials and references, automatically reviewing materials for potentially problematic language, drawing on previously approved language to instantly offer compliant rephrasing options, and accelerating compliance checks and speeding up the approval process by automatically ensuring that reviewers stay in the loop.

Potential impact: A two to three times acceleration of the content approval process; an increase in the speed of review

Use case three: Customer-enablement co-pilot. Marketers, field reps, and other customer-facing team members have an unprecedented amount of information at their disposal. What they often lack is an effective way of synthesizing the right insights at the right time. To increase the productivity of market analysts, generative AI tools are being developed to provide on-demand retrieval, summarization, and synthesis of both unstructured data (such as text and images) and structured information (for instance, tables and databases). That makes campaigns more effective and enables richer, more targeted conversations between field staff and care providers.

Potential impact: A 10 to 15 percent improvement in the productivity and effectiveness of field teams, resulting in 1 to 2 percent topline growth

Use case four: Generating strategic insights. Today, brand leads and marketers spend significant time and resources synthesizing business and market insights. They aim to link brand questions to hypotheses, to identify approaches for analysis, to uncover competitive intelligence about rivals, and to create compelling documents with cohesive brand narratives. One recurring problem is that marketers spend too much time synthesizing diverse sources of information and not enough interpreting data to make key decisions about a brand's direction. Gen AI's interactive search capabilities can help marketers draw deeper insights from sources such as customer research and data sets, data on physicians and patients, policy changes, legal developments, and formulary implications. These insights will help identify and sharpen understanding of patterns in customer journeys while informing and enriching customer segmentation.

Potential impact: A 10 to 30 percent improvement in the understanding of key insights and trends; embedding those insights across customer segments

Use case five: Optimizing the patient experience. Many patients stop taking their prescribed medications—or never fill the prescriptions in the first place. (In immunology, for example, about half of all patients stop taking new medications during the first year.) This happens increasingly often—even among patients who benefit from therapies—for a number of reasons, including the complexity of getting reimbursed for prescriptions and the proliferation of new medicines available. Gen AI can help address this significant problem by providing patients and physicians' offices with on-demand insights about reimbursement and proper care options. The technology can also help escalate critical issues to experts while empowering patients and physicians' offices with a range of self-service tools. All of this can lead to increased patient adherence and improved outcomes, in part because the technology can address unmet needs by upskilling patient service teams.

Potential impact: A 5 to 10 percent decrease in the number of patient drop-offs through better assistance for patients

Medical affairs: Smarter providers, healthier patients

Potential opportunity: $3 billion to $5 billion through efficiency gains

Unlike other the other life-science domains, medical affairs neither generates revenue nor focuses on enhancing commercial performance, sales, or profitability. Instead, medical-affairs teams play a key role supporting patient care by enabling trustworthy scientific exchanges between biopharmaceutical companies and healthcare providers (see sidebar "A note on managing gen AI risks in medical affairs"). The heart of this highly regulated effort is a series of processes for reviewing, synthesizing, and communicating scientific knowledge—all tasks amenable to improvement through gen AI.

Here we present three use cases now generating considerable excitement among medical-affairs leaders. These applications will support ongoing trends toward digitally enabled ways of working and greater personalization of customer engagement models. They could also free up time from mundane content generation tasks and allow teams to focus on creating deeper, high-value personal interactions with internal and external stakeholders.

Use case one: Generating customer insights. Medical-science liaison (MSL) professionals spend hours interacting with customers to create a richer, more holistic understanding of external scientific and clinical perspectives on disease areas and patient care. These interactions are crucial learning opportunities for the pharma industry, and the ability to systematically capture insights from them has long been a goal for medical-affairs teams.

With appropriate consent and privacy guardrails, large language models can help make the most of the rich data created during these interactions: they can synthesize and analyze insights and improve the performance of existing natural-language processing methods and help summarize and clarify thousands of discussions each year—say, by identifying common themes raised by clinicians across tens of thousands of conversations. These insights can then better inform development areas, ensuring that therapies meet the needs of patients and providers more effectively.

Potential impact: A multifold increase in the capture of insights, because typical medical-affairs organizations systematically analyze less than 1 percent of field interaction notes

Use case two: Sharper, more efficient medical communication. Gen AI models can dramatically reduce the time needed to craft technical and scientific documents, by accelerating the composition of medical-information request responses, scientific reports, lay summaries, and day-to-day communications with internal and external stakeholders (Exhibit 7). In addition to generating content, gen AI tools could help review and moderate it, thus boosting productivity and accelerating workflows. For example, medical-information professionals must often respond to emails from clinicians asking for information on clinical trials or evidence to improve the treatment of particular subpopulations of patients. At present, this time-intensive task often requires multiple people to understand the questions, identify the appropriate literature, and synthesize responses. Gen AI can help screen incoming requests, digest relevant content, and even produce the first draft of an answer, so medical teams can respond more rapidly, using a potentially broader base of scientific evidence.

Potential impact: 20 to 30 percent savings in medical-writing costs (and potentially 50 to 70 percent once solutions and systems mature); a 50 to 70 percent reduction in time to deliver medical–legal reviews

Use case three: Rapid summaries of the scientific and medical literature. Another challenge for medical-affairs teams is engaging external stakeholders with scientific content tailored specifically to their own requirements. This content's technical nature and the need to comply with regulatory standards mean that MSLs must often provide stakeholders with lengthy scientific documents (such as published research papers or long presentation decks), though they would actually prefer more focused responses—say, the answer to a specific query. Using gen AI tools trained on approved content, medical-affairs teams could rapidly pull together tailored materials, including text, data tables, figures, infographics, videos, and audio.

Potential impact: A two to three times boost in engagement, associated with personalized content in other industries (and among early adopters in medical affairs)

Exhibit 7

Generative AI may provide higher-quality interactions between medical science liaisons and healthcare providers.

Medical science liaison capabilities, present and future

	Present	Future
Shape strategy for healthcare provider (HCP) engagement	• Collect available information about HCPs to develop engagement strategies • Prepare materials for HCP meetings using available data and experience	• Use insights created by generative AI about HCPs to better shape engagement strategy • Review generative-AI-created materials to provide customized in-depth information to HCPs

Source: McKinsey analysis

McKinsey & Company

From use case to scale: Rewiring the organization for gen AI

The use cases we have described are compelling pilots for life-science companies taking their first steps in gen AI. Collectively, they could create a paradigm shift in the discovery, development, and delivery of therapies to patients. But such a shift will not happen by accident. To truly capture gen AI's value, it is critical to design systems for eventual scaling rather than implementing the technology as a series of isolated solutions.

What will be needed is a comprehensive strategy clearly linked to scientific and business value, supported by a robust technology infrastructure that includes an end-to-end tech stack and an awareness of the evolving AI landscape. In many cases, the most successful model is the product–platform approach: the IT enterprise group leads the development of the appropriate platform infrastructure—with critical enterprise components, such as security, data governance, and cost tracking—and provides an API for creating use cases (or products). Companies develop a minimum viable platform to enable the initial use cases and improve it over time so that it can support increasingly complex ones.

To make all this work, pharma companies will need to reimagine their operating models. One common mistake leaders make is to embrace either of two extremes for managing digital transformations. One is a highly decentralized approach, in which the organization simultaneously launches multiple use case pilots simultaneously. While this strategy—think of it as allowing 1,000 flowers to bloom—lets companies move quickly, it often leads to quality, cost, and sustainability problems, as well as operational silos that inhibit the sharing of knowledge and the development of synergies. The opposite approach—a top-down platform-based model with centralized decision-making and a phased rollout of use cases—is also problematic. Although it is cost efficient and allows leaders to build for scale from the outset, it is also slow and often frustrating.

Leaders adopting the product–platform approach must therefore move between the two models, striving

A guide to our methodology

We based the potential economic impact of gen AI in different domains of the life sciences on the McKinsey Global Institute's analysis of its impact in 63 individual use cases. Each of them was then mapped to a specific life-science domain—an exercise based on the typical activities performed within it. Generative AI co-pilots for sales representatives, for example, were mapped to the commercial domain.

For generative AI use cases (such as synthesizing scientific literature) that affect multiple life-science domains, we based our allocation of the economic impact on the relative size of the relevant domains, as well as expert opinion. Our quantification of the economic impact reflected a bottom-up approach: expert opinion, the results of McKinsey's internal experiments, and published research helped us estimate the potential quantitative range of impact, in both cost savings and revenue uplift, for specific use cases (for instance, the percentage of time saved when generative AI tools review the clinical-research literature).

Finally, we scaled the impact to the life-science industry, basing our estimates on the percentage of domain operating costs spent on tasks related to a generative AI use case (for example, the percentage of life-science R&D costs allocated to reviewing clinical research) and the percentage of industry revenues spent on a domain (such as the percentage of life-science revenue allocated to R&D). Readers will find full details in the appendix of "The economic potential of generative AI: The next productivity frontier."

constantly to balance speed and innovation with the need for quality control, financial discipline, and organizational symmetries. Wherever leaders fall on this spectrum, they will also need to ensure that they have the proper mix of talent on hand, as well as the necessary guidelines and controls to mitigate risk and ensure trust and buy-in from stakeholders.

To make this shift from implementing use cases to generating value at scale, pharmaco leaders will have to reimagine each step of the value chain (Exhibit 8). That will require them to ask important questions about structures, processes, technologies, data, people, and change management.

Structures

When an organization has established a high-level gen AI strategy and a blueprint for implementation, it will need to address whether it has the right structures in place. Which organizational structure will best support the strategic requirements—for instance, a product-centric approach or the centralized center of excellence (CoE) model? How will high-level decisions about critical strategic issues (such as roles, governance, and the size of the workforce) be finalized? How will risk and compliance considerations be woven across the organization?

Of course, once the necessary structures are in place, the organization must be open to change and allow them to evolve over time.

Processes

New gen AI tools will involve both incremental changes and fundamental shifts in the way pharma professionals work. Strategists and decision makers will need to exploit gen AI's ability to rapidly analyze and synthesize information, yet remain mindful of the technology's limitations. So will researchers and creatives when they brainstorm. As communications grow increasingly automated, teams may need to rethink the ways they collaborate. Implemented correctly, these adaptations will focus employees on high-value activities and fundamentally reshape roles throughout the organization.

Technologies

Many organizations have been experimenting with foundational models such as ChatGPT, but it's important to recall that the LLM itself accounts for

Exhibit 8

Life sciences workers of tomorrow will need to be creative thinkers, knowledgeable in the latest generative AI developments.

Life sciences employee capabilities, present and future

	Present	Future
Must know:	• Have knowledge of pharmaceutical/medical products and industry landscape • Proficient in traditional data analysis techniques (eg, using Microsoft Excel)	• Employ generative AI (gen AI) to produce insights from various sources of information • Proficient in gen AI tools (eg, natural-language processing, prompt engineering)
Must be:	• Compliant by following regulatory requirements and guidelines • Collaborative with cross-functional teams	• Data-driven: make data-driven decisions that enhance their effectiveness • Effective communicator: convey complex gen AI insights to nontechnical stakeholders in a clear and concise manner
Must do:	• Employ iterative methods, experiment to refine approaches, and invest time in creating documents	• Use gen AI for automated report and documentation generation, focus on strategic tasks and goals, and dedicate significant time to mentoring junior colleagues

Source: McKinsey analysis

McKinsey & Company

as little as 15 percent of an overall gen AI solution, McKinsey research has found. To truly scale this technology, organizations must design and adapt a comprehensive, end-to-end tech stack, prioritizing the selection of models and considering specific needs for information security, task-oriented performance, and latency. With many simpler applications trending toward commoditization, organizations must also judiciously balance buying solutions from outside vendors with building them in-house. Finally, organizations scaling gen AI initiatives frequently face unforeseen costs. That highlights the need for robust financial governance and a financial-operations (FinOps) framework for meticulous budgeting, vigilant monitoring, and efficient management of the resources for implementing gen AI.

Data

The effectiveness of gen AI depends on the quality of an organization's data, which must be continually enriched to order for sharing across internal functions. For strategic differentiation, organizations must invest in proprietary data. In both cases, they will have to create and manage labeled data sets to quantify, measure, and track the performance of gen AI applications. As they do so, they will also need to deal with existing data quality and curation challenges (such as redundant, outdated, and conflicting information) while ensuring that AI models have the right context for responding accurately to prompts.

People

The accelerated pace of change that gen AI promotes will require organizations to think differently about which skills they need (Exhibit 9) and where to find them—recruiting new people, upskilling existing employees, or dynamically allocating the right talent to the appropriate priorities. External recruiting will not be easy. Since ChatGPT's release, in late 2022, the number of AI-related job listings has quadrupled. In biopharma alone, the number of AI-related job postings has grown by 43 percent annually across the top ten pharma companies since 2018, according to McKinsey research.

Attracting and retaining AI talent will be possible only if organizations maintain the right kind of work environment: dynamic (and potentially disruptive) gen AI use cases, a modern tech stack, and an agile working model that promotes experimentation and the development of skills. Hiring is only part of the equation. Many current employees will need to be upskilled to meet the evolving job requirements and new demands of their current roles; others must develop new skills to fill new roles. Companies will have to develop at-scale

Exhibit 9

Scaling generative AI requires a holistic strategy across the organization.

Sampling of key strategic questions that must be answered

Strategy

- **Road map**
  - How do we align our generative AI (gen AI) strategy with our technology aspirations?
  - How should we approach the transformation in a way that ensures value capture and unlocks competitive advantage?

Capabilities

- **Data**: How do we set up a robust data foundation to scale gen AI across the organization?
- **Operating model**: How do we organize ourselves and teams to deliver on our gen AI strategy?
- **Talent**: How do we manage talent to stay ahead of the gen AI skills gaps?
- **Technology**: How do we set up a scalable tech stack and infrastructure to support multiple gen AI use cases and solutions?

Change management

- **Adoption and scaling**
  - How do we design our scaling plan to ensure easy reusability and scalability across the organization?
  - How do we deliver effective training to support skill building and manage culture change at scale?
  - How do we think about risk and responsible use of gen AI across the organization?

Source: McKinsey analysis

McKinsey & Company

technical-upskilling programs, led by technically informed leaders, to identify the roles that can be upskilled, to monitor the performance of recently trained employees, to assess the program's impact on business objectives, and to continually source new candidates as needs evolve.

Change management

Even incremental change can be destabilizing, which is why change management is essential for any organizational transformation. To ensure the sustained adoption of gen AI at scale, organizations should embrace an influence model that promotes shifts in both mindsets and behavior. Leaders throughout the organization, including the C-suite, must provide guidance, resources, and motivation for stakeholders affected by gen AI. These leaders must not only communicate a transformation story to build awareness and excitement, foster trust, and address misconceptions or concerns but also participate in continuous exchanges to help ensure the sustained adoption of gen AI at scale.

In addition, organizations must adapt their workflows and roles to incorporate gen AI into their daily activities. Before launching it to users, they should train employees in its limitations, strengths, and risks, as well as the structuring of prompts. Finally, to build momentum for change, organizations should create teams of early-adopter champions to shape the deployment of gen AI use cases and prove their value.

Responsible AI and mitigating risk

It's impossible to discuss the promise of generative AI without also considering its inherent risks. While those risks are considerable for all industries, they are particularly elevated for pharmacos, which face not only a unique and complex regulatory environment but also heightened concerns about intellectual-property infringement and data privacy. Factor in the high-stakes nature of treatments for disease and it becomes even more essential that life-science firms carefully assess the risks gen AI poses and construct policies and guardrails to mitigate them. Off-the-shelf gen AI models

A note on managing gen AI risks in medical affairs

The unique context of medical affairs creates specific risks requiring careful management. Crucially, medical-affairs content is highly technical, and widely available foundation models may not perform well enough initially on tasks involving the medical literature, scientific concepts, and regulatory-standard documentation. Task-specific models using general-purpose or specialized foundation models (such as BioBERT) may, with fine-tuning, come to play a large role in medical use cases; in this event, companies will need to partner with other organizations or build their own way toward a level of functionality that creates meaningful value.

A high bar will be needed to generate content for scientific and clinical exchanges as well. Clinicians, regulators, and patients all expect to receive scientifically accurate content free from hallucinations and mischaracterizations. Since medical-affairs teams also handle commercially sensitive data, such as unpublished research findings, gen AI solutions must meet a high bar for data security and information governance. That will probably require the development of specialized risk modules for medical gen AI tools. In all likelihood, such tasks will require a human-in-the-loop approach. Gains from automation will therefore be limited unless risk management processes evolve a good deal. What's more, there will probably be no meaningful productivity gains without a thoughtful redesign of the operating model of medical-affairs organizations and investments in creating a culture of agility for technology-enabled transformations.

will be difficult to deploy. Instead, these companies should seek to mitigate risk through bespoke technology solutions with proper controls and robust human oversight.

Pharma companies must also understand that different use cases, and different domains, face different kinds of risks. Medical affairs, for example, is a challenging environment because recommendations can directly affect the lives of patients. In the research domain, by contrast, the risk of a failed experiment resulting from a gen AI hallucination (or inaccuracy) has less serious consequences. In some cases, hallucinations might actually be helpful—for example, if a model suggests a chemical that has never before been considered as a potential treatment.

These are some of the risks that executives need to consider.

Inaccurate models. Pharmaceutical companies are exposed to the risk of gen AI hallucinations resulting from poor or incomplete data. They can mitigate that kind of risk by placing guardrails around gen AI content—for example, ensuring that humans review it before it is distributed to providers or patients. The bottom line: gen AI should never be the final decision-maker; it should instead accelerate decision-making by human workers.

IP infringement and data privacy. Foundational models typically include large volumes of internet-based data, and that has led to alleged copyright violations, plagiarism, and other forms of IP infringement. This risk is particularly high among life-science companies because of the exceptionally stringent data privacy regulations surrounding patients' medical data; many countries, for example, require this information to remain on domestic servers. To avoid infringing IP, businesses using foundational models need proper guardrails, such as training models on their own intellectual property and writing IP protections into contracts with external vendors.

Regulatory compliance. AI regulations with specific ramifications for generative AI are being planned for near-term enactment, although they vary by country—for instance, the European Union's proposed AI Act, US federal and state law, and China's draft administrative measures for generative AI services. Commercial pharma faces additional risks as a result of increased regulation by the US Food and Drug Administration and similar agencies on the contents of advertising and promotional materials. Pharma companies can mitigate these risks by embedding guardrails directly in content-generating models and ensuring that humans always make the final decisions.

Conclusion: Lessons from electricity

As the use cases discussed in this article show, pharmaceutical companies have already embarked on their generative AI journeys. But the technology alone does not guarantee success. Companies won't unlock the full potential of gen AI until they understand how to use it correctly. That will require both time and patience.

It's helpful to look at gen AI through the lens of another disruptive technology, which swept through society at the turn of the last century: electricity. The first light bulbs were invented in the 1870s, and several years later the first electric motors started to power manufacturing machinery. Yet as business historians and innovation experts like to point out, by 1900 electricity powered less than 5 percent of manufacturing, and even as late as the 1910s most facilities continued to rely on steam. Why weren't business leaders more eager to use the new technology? After all, electric motors were far cleaner, safer, and more efficient than steam engines.

"Why didn't electricity immediately change manufacturing?," an article on the site of the British Broadcasting Corporation, explains that "to take advantage of electricity, factory owners had to think in a very different way." They couldn't get the benefits of the new technology "simply by ripping out the steam engine and replacing it with an electric motor. You needed to change everything: the architecture and the production process. And because workers had more autonomy and flexibility, you even had to change the way they were recruited, trained and paid." Of course, that's precisely what happened eventually, and by the 1920s productivity was soaring.

What does this history mean for the pharmaceutical industry and generative AI? For starters, we may not see massive industry-wide productivity gains immediately. Nor is it possible to know whether gen AI will unleash an electricity-level—or, for that matter, an Internet-level—degree of transformation across the business world. Pharma also faces its own unique challenges, including the complex nature of biological systems, the difficulty of treating disease, and a stringent regulatory environment.

Nonetheless, gen AI could give the pharma industry a once-in-a-century chance to address those long-standing obstacles and create new breakthroughs in science and patient care. Much as it would have been foolish for factory owners to stick with steam in the 1910s, pharma companies would be unwise not to recognize the transformative potential of AI. They should begin working now to understand, implement, and scale it.

Chaitanya Adabala Viswa is a partner in McKinsey's Boston office, where **Delphine Zurkiya** is a senior partner; **Joachim Bleys** is a senior partner in the Charlotte office, where **Bhavik Shah** is a partner; and **Eoin Leydon** is a partner in the Dublin office.

The authors would like to thank Mahmoud AbuEid, Asli Aksu, Chris Anagnastoupoulos, Salma Ayoub, Eric Bruckner, Lucia Darino, Supreet Deshpande, Alex Devereson, Aliza Dzik, Anas El Turabi, Lionel Jin, Matej Macak, Abhi Mukherjee, Raj Rajendran, Boyd Spencer, and Hann Yew for their contributions to this report.