The next innovation revolution—powered by AI

AI isn't just for efficiency anymore. It can double the pace of R&D to unlock up to half a trillion dollars in value annually.

This article is a collaborative effort by Alex Singla, Alexander Sukharevsky, Elia Berteletti, Lareina Yee, and Michael Chui, representing views from QuantumBlack, AI by McKinsey, and McKinsey's Operations Practice.

June 2025

The innovation challenge: Good ideas are harder to find

Innovation has been the driver of the extraordinary progress from which humankind has benefited for a couple of centuries, but it faces a largely hidden threat: Innovation is becoming harder and more expensive.

It's instructive here to take the long view. For most of recorded human history, improvements in human welfare from generation to generation have been limited. Take, for example, GDP per capita as a measure of economic prosperity. For most of human history, roughly until the early 1800s, the measure barely moved to $1,200. But since that time, it has grown by more than 14 times (Exhibit 1).[1] Human health has followed a similar trajectory—low for centuries and only significantly improving in recent generations. In 1900, for example, the average life expectancy of a newborn was 32 years. By 2021, this had more than doubled to 71 years.[2]

Exhibit 1

Building on scientific discoveries, the Industrial Revolution sparked great improvements in human welfare.

GDP per capita globally over the past ~2,000 years, $ thousands

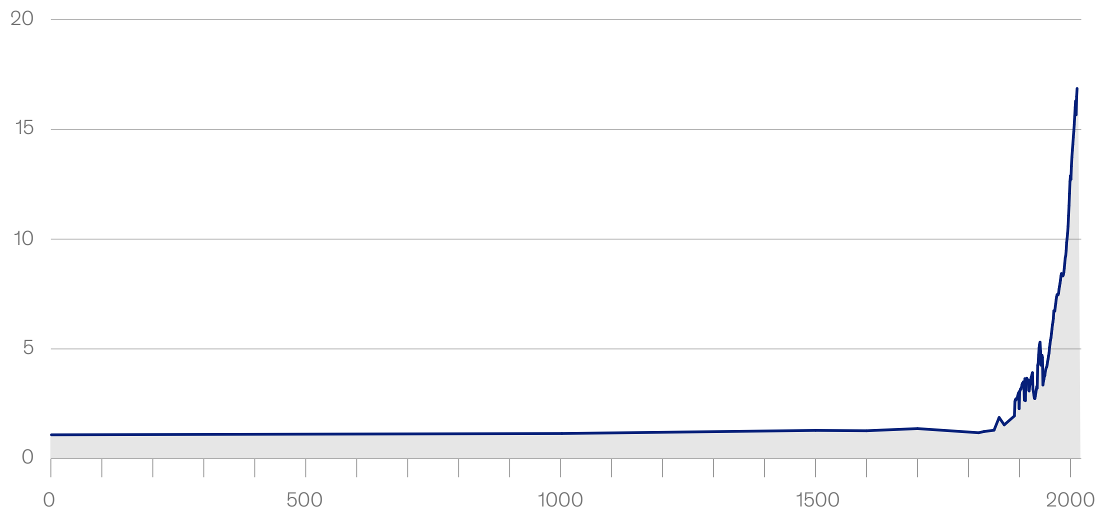

Source: Maddison Project Database version 2023: Jutta Bolt and Jan Luiten van Zanden, "Maddison-style estimates of the evolution of the world economy: A new 2023 update," *Journal of Economic Surveys*, Apr 3, 2024; McKinsey analysis

McKinsey & Company

[1] Jutta Bolt and Jan Luiten van Zanden, Maddison Project Database 2023; Jutta Bolt and Jan Luiten van Zanden, "Maddison-style estimates of the evolution of the world economy: A new 2023 update," *Journal of Economic Surveys*, 2024, Volume 39.
[2] Saloni Dattani et al., "Life expectancy," Our World in Data, 2023.

These and many other improvements in our lives have been driven by a set of scientific discoveries and products engineered based on those breakthroughs. These innovations have enabled economies to grow and people's lives to improve. The steam engine helped power the Industrial Revolution. Vaccines that prevent diseases such as smallpox, measles, and polio continue to save millions of lives each year; infant mortality is estimated to have decreased 40 percent in the past 50 years because of vaccines.[3] The invention of the integrated circuit for computing and lasers for communication through fiber-optic cables helped create the global internet.

But the rate of progress enabled by innovation now faces an under-recognized threat: Innovation is getting more difficult and more expensive.

Even as science advances, R&D productivity is on the wane

By many metrics, and in many fields, each dollar spent on R&D has been buying less innovation over time. In other words, R&D productivity has been declining.

Take the semiconductor industry. With integrated circuits embedded in products that support nearly every part of our lives, this sector has advanced in accordance with "Moore's Law"—the remarkable observation put forward by Intel cofounder Gordon Moore that the number of transistors on an integrated circuit will double about every two years.[4] This is roughly equivalent to an exponential growth rate of 35 percent annually in transistors per dollar.

But this level of performance increase has been bought at the cost of increasing expenditures in R&D. Nicholas Bloom, an economics professor at Stanford University, and his research collaborators published a paper in 2020 that examined the real R&D expenditures of semiconductor companies and equipment manufacturers and estimated that their annual research effort rose by a factor of 18 between 1971 and 2014.[5] In other words, maintaining the performance growth rate in Moore's Law required 18 times more inflation-adjusted R&D spending in 2014 than it did in 1971 (Exhibit 2).

The rate of progress enabled by innovation faces an under-recognized threat: Innovation is getting more difficult and more expensive.

[3] Andrew J. Shattock et al., "Contribution of vaccination to improved survival and health: Modelling 50 years of the Expanded Programme on Immunization," *Lancet*, 2024, Volume 403, Number 10441.
[4] Moore's original 1965 statement suggested doubling every year, but he updated it in 1975.
[5] Nicholas Bloom et al., "Are ideas getting harder to find?," *American Economic Review*, 2020, Volume 110, Number 4.

Exhibit 2

Increased real R&D expenditures are required to maintain Moore's Law on growth in the density of semiconductor chips.

Inflation-adjusted semiconductor R&D expenditures, 1971–2014,[1] multiples (factor increases)

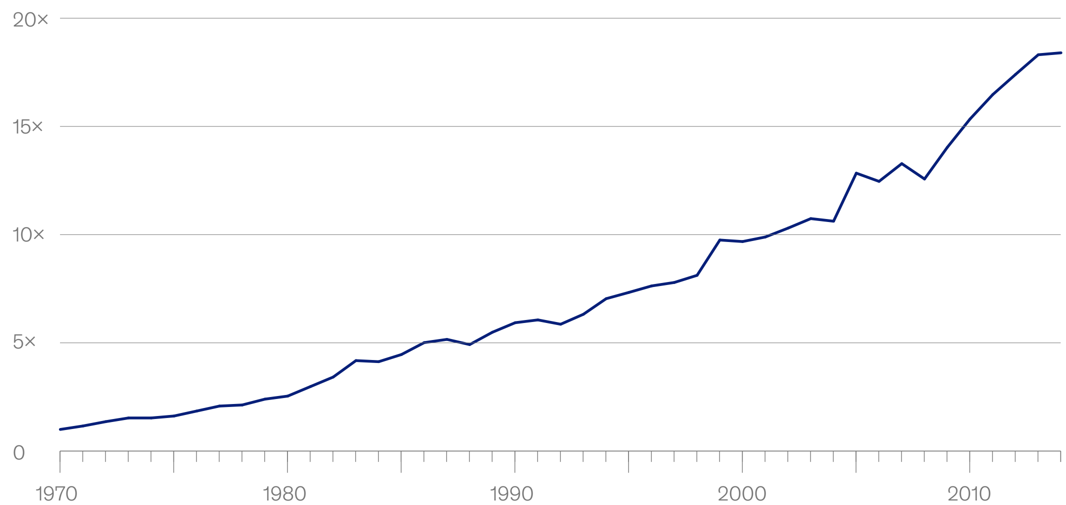

[1]Nominal semiconductor R&D expenditures deflated by nominal wage of high-skilled workers.
Source: Nicholas Bloom et al., "Are ideas getting harder to find?," *American Economic Review*, Apr 2020

McKinsey & Company

It's not just semiconductors. The biopharmaceutical industry has produced innovative products used to prevent and treat many diseases, enabling millions of people to live longer and healthier lives. But the challenge of declining R&D productivity in that industry led Jack Scannell, a multidisciplinary life sciences analyst, researcher, and entrepreneur, to coin the term "Eroom's Law" (that is, the reverse of Moore's Law) to describe the fact that drug discovery has become slower and more expensive over time.[6] He and his research collaborators found that the number of new drugs approved per billion US dollars spent on R&D halved roughly every nine years between 1950 and 2011, falling around 80-fold in inflation-adjusted terms (although the decline appears to have stabilized somewhat in the past decade) (Exhibit 3).

Declining R&D productivity has been reported in other fields, such as agriculture, where higher yields for multiple crop types require increasing levels of R&D spend. Using company-level data across all sectors in the United States, Bloom and his team found that R&D productivity declined in general, with output measures including revenue, market capitalization, employment, and revenue per employee. (This was not the case for all companies: While most experienced a decrease in R&D productivity, 22 percent of organizations increased research productivity.)[7]

6 Jack W. Scannell et al., "Diagnosing the decline in pharmaceutical R&D efficiency," 2012, *Nature Reviews Drug Discovery*, Volume 11, Number 3.

7 From a macroeconomic standpoint, Bloom and coauthors also compared the growth in total factor productivity (TFP) in the US economy with the level of research effort (gross domestic investment in intellectual property products in the national income and product accounts) and found that increasing levels of research investments were necessary to sustain TFP growth from 1930 to the end of the century. Nicholas Bloom et al., "Are ideas getting harder to find?," *American Economic Review*, 2020, Volume 110, Number 4.

Exhibit 3

‘Eroom’s Law’: Drug discovery is becoming more expensive over time.

Innovative drugs per $ billion in inflation-adjusted R&D spend, 1950–2023, number (logarithmic scale[1])

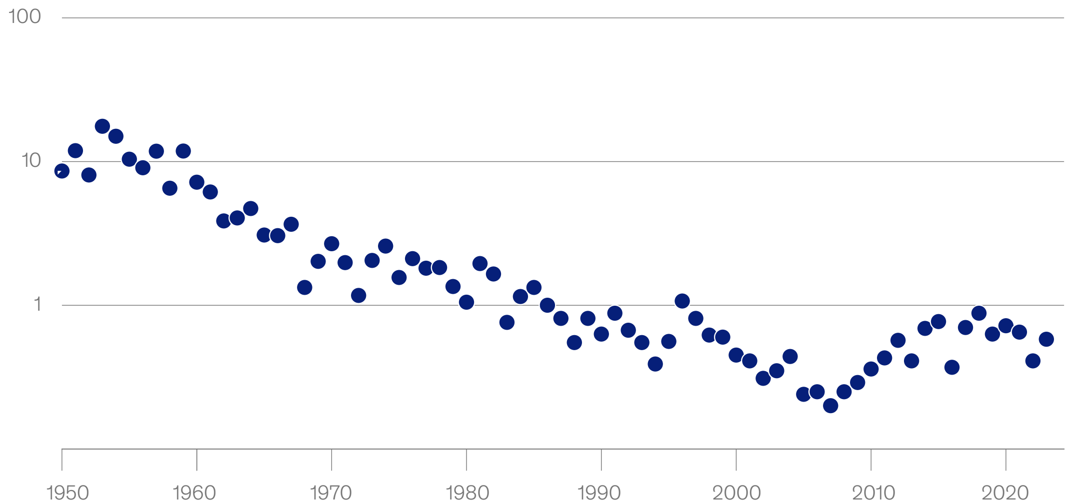

[1]A log scale is being used to show the exponential decay.
Source: *Nature* reviews 2012 and 2015; PhRMA member companies R&D expenditures; McKinsey analysis

McKinsey & Company

AI has the potential to bend the curves of R&D productivity, not only unlocking more economic growth but also boosting the chances of solving some of the most important human challenges, from preventing and curing diseases to reducing the level of carbon emissions.

How AI can reignite innovation productivity

Over the past decade, we have seen how AI, when coupled with complementary management practices to rewire the way organizations operate, can generate real business value. Even prior to the advent of gen AI, analytical AI was being used by roughly half of the enterprises represented in McKinsey’s Global Survey on AI.[8] Those organizations have been deploying the technology across a variety of business functions—from increasing revenue through more targeted marketing to reducing costs in supply chain operations. Since ChatGPT became available in late 2023, the percentage of organizations reporting that they use AI has spiked upward by 20 percentage points, with companies implementing gen AI in use cases from customer service to software engineering.

Most of these applications of AI have been aimed at improving the efficiency of existing tasks and workflows. But boosting efficiency and productivity is just one way that AI promises to unlock a new era of growth and opportunity. Our research shows that AI also can be deployed to accelerate innovation to create entirely new products and services. To put it another way: AI can be used to bend the curves of the declining R&D productivity we documented in the previous section.

[8] “The state of AI: How organizations ar rewiring to capture value,” McKinsey, March 12, 2025.

We have identified three primary channels through which AI technologies can accelerate innovation, each with a corresponding type of model: increasing the velocity, volume, and variety of design candidate generation; accelerating the evaluation of candidates through AI surrogate models; and accelerating research operations.

Increasing the velocity, volume, and variety of design candidate generation

A simplified model of the R&D process consists of identifying a set of customer needs, generating candidate designs, and then evaluating those ideas to identify the most promising ones that will best meet the needs of the customer or user. One of the highest potential opportunities for AI to enhance innovation is to more quickly generate a greater volume and variety of design candidates.

Gen AI technology is based on foundation models—very large simulated neural networks that are trained on vast collections of data to take unstructured data (that is, data that isn't best stored in rows and columns like a spreadsheet, such as human language) as inputs and then generate unstructured data as output. Large language models (LLMs) are the best-known types of foundation models, underpinning the chatbots that have made gen AI such a compelling technology.

However, foundation models can be trained to produce outputs other than human language. They can be trained to generate chemical compounds, drug candidates, computer code, electrical designs, physical designs, and other types of potential solutions. With sufficient computing power, these models can generate design candidates far more quickly than researchers, designers, or engineers can on their own—increasing the number of "shots on goal" that could potentially produce a successful design.

For example, a retailer used gen AI tools to create dozens of alternative 3D store configurations, rendered with photorealistic fidelity. Using traditional computer-aided design and rendering tools, a designer might have only created a handful of sketches, and at a much lower level of fidelity. Without the ability to quickly generate a variety of alternative designs, many of these options would likely not have been considered. An unexpected side benefit of the AI-generated 3D renderings was the discovery of certain aesthetic decor features inserted by the foundation model to fill out the rendering—features that appealed to consumers but were not in the initial design parameters.

Thus, not only can AI quickly generate a greater volume of candidates, but AI systems also can generate a greater *variety* of candidates—in particular, designs that a human researcher or engineer would be less likely to produce, given the biases that stem from their training and on-the-job experiences. Provocatively framed, AI can be more creative than humans.

An early example of AI's ability to generate ideas that a human would not have considered occurred in March 2016. DeepMind had trained an AI-powered system called AlphaGo that squared off against the world's top Go player, Lee Sedol, in Seoul. Go is considered one of the most complex and strategic board games in the world. Perhaps more remarkable than the fact that the AI triumphed in the best-of-five match was the now-famous "Move 37" in game two: AlphaGo made such an unexpected move that several commentators believed it was a mistake. It was a completely outside-the-box move, defying centuries of conventional Go strategic principles. It was, as one commentator noted, "a move no human would ever make." It was fresh, it was novel—and it was foundational to AlphaGo's victory.

Some R&D organizations have described AI as generating similarly innovative ideas in the lab. For example, David Baker, a researcher at the University of Washington, has led a team that uses deep learning models to design novel proteins that bind and catalyze other reactions. More specifically, Baker and his team are creating entirely new proteins, complex functional molecules that interact with other molecules at a subatomic level and don't already exist in the world—a goal that has been beyond the capability of scientists to accomplish without AI tools. Among the applications of these custom-designed proteins: new vaccines and medicines, biosensors for hazardous materials, and agents that can capture or break down environmental pollutants. For leading this pioneering work, Baker shared the 2024 Nobel Prize in Chemistry.

The ability to use AI to creatively generate a greater variety of candidates hasn't only been applied at the molecular level; it is also being applied in physical engineering disciplines. Generative models, for example, are currently being used to design rocket engines with novel geometries, particularly their cooling channels, which are becoming manufacturable with 3D printing.

Accelerating the evaluation of candidates through AI surrogate models

A complementary activity in the product development life cycle is the evaluation of candidate designs. For physical products, this has historically meant manufacturing prototypes and then subjecting them to a regimen of physical tests—for example, the crash tests that automobile manufacturers perform to test the safety of their vehicles. But these tests tend to be both costly and time consuming.

Unsurprisingly, over many years, scientists and engineers have developed mathematical and computational models to simulate the performance of physical systems to perform in silico testing. So, rather than putting an airplane design into a wind tunnel or a racing yacht design into water, designers use computational fluid dynamics (CFD) to evaluate the performance of a particular configuration. Instead of building a prototype structure to determine how the design candidate might perform, engineers can use finite element analysis (FEA) to predict how forces will affect a structure. Rather than setting up physical experiments, for example, radio engineers can use computational electromagnetic (CEM) modeling to predict how an antenna design might perform.

While these acronym-laden, physics-based mathematical models are often less expensive than physical experiments, these simulations are often extremely computationally intensive and can take many hours, or even days, to run. But a recent discovery found that it is possible to repurpose the neural network technology developed for AI systems to train models that can act as proxies for more computationally intensive physics-based models. These AI-style surrogate models do not imitate the "thinking" that people do; instead, they predict the outcomes of physical phenomena in the world. When used to predict the behavior of a complete system, these models are akin to a "digital twin."

Take weather forecasting. Over the years, scientists have developed complex and detailed models of the Earth's weather that have enabled increasingly accurate forecasts. However, because of their computational intensity, these physics-based simulations must be run on powerful supercomputing clusters. DeepMind, for example, trained a neural-network-based machine learning model that predicts the weather faster (eight minutes versus hours) and more accurately on a single AI-optimized processor than a top operational physics-based weather-forecasting system running on a supercomputer with tens of thousands of processors.[9]

The same kind of techniques are being used to evaluate product designs. As previously noted, for example, computational fluid dynamics is used to simulate the aerodynamic performance of aircraft (and automobiles, including for racing). Designers are now using neural network models trained on wind tunnel and CFD data to predict hundreds of results in a few seconds for a range of flow velocities and angles that were not included in the wind tunnel testing or CFD simulations that otherwise would have taken hours or days to produce. The benefit here is not simply increasing the speed of a single simulation run per se, but the ability to test a panoply of possibilities. In a CFD case, engineers can test many alternatives to optimize the design of a turbine compressor. They can then use other automated systems to check for manufacturability, reliability, product cost, et cetera, and run through iterations that would otherwise not have been possible in a reasonable time.

In the life sciences, researchers are using similar techniques to study the proteins that exist in the world. Predicting the 3D-folding structure of proteins from their known sequence of amino acids has historically been incredibly challenging, involving myriad quantum-level interactions at the subatomic level. British AI researchers Demis Hassabis and John Jumper won the other half of the 2024 Nobel Prize in Chemistry for training a model that can predict the 3D structure of proteins, which has now been used to predict the structure of around 200 million proteins, covering almost every known protein.[10] The ability to predict molecular structures and their interactions can enable the testing and evaluation of various biological products, from therapies to treat disease to biological production of materials.

Some design challenges require evaluating and optimizing designs across multiple physical phenomena that interact with one another, so-called multiphysics problems. Requirements to analyze multiple modalities multiply the complexity of modeling them together. For instance, designing an aircraft antenna could require an understanding not only of the design's radio frequency characteristics but also its aerodynamic and thermal properties, all of which can interact with one another. Integrated neural-network-based models, given sufficient training data, can integrate a variety of modalities, multiplying their potential to accelerate design candidate evaluations.

Accelerating research operations

In addition to generating and evaluating design candidates, there are several additional ways that LLMs, sometimes coupled with other AI technologies, are being used to accelerate various activities in the product development process:

Identifying and analyzing customer/user needs, products, and features. LLM-powered software solutions are being used, particularly by consumer companies, to synthesize a vast array of product reviews, social media posts, customer service transcripts, and other sources of customer data to identify addressable market segments and the product categories and features/functions that would best address the as-yet unmet needs of customers.

---

9 Ilan Price et al., "Probabilistic weather forecasting with machine learning," *Nature*, 2025, Volume 637.

10 Mihaly Varadi et al., "AlphaFold Protein Structure Database in 2024: Providing structure coverage for over 214 million protein sequences," *Nucleic Acids Research*, January 25, 2024, Volume 52, Number D1.

Exploring and synthesizing existing research and data. In industries such as life sciences, chemicals, and materials, there is a vast and rapidly growing body of published research and databases. It can be challenging for scientists to keep up with the literature in their own subdiscipline, not to mention the adjacent or even distant areas of other research, which could bring insights for breakthroughs in their field. Oftentimes, the volume of machine-readable data being made available is growing even more rapidly than published papers.

Tools enabled by LLMs and analytical AI can synthesize insights from published literature and databases, both to inform innovation practitioners and to suggest potential avenues for creating solutions. Google, OpenAI, Perplexity, and Anthropic, for example, have all introduced knowledge agent products that perform multistep research tasks that one might otherwise assign to a research assistant: creating a work plan, searching a set of sources on the web, producing a well-structured research report.

Streamlining internal knowledge management. Not only is there a burgeoning volume of publications and data available publicly, but large corporations hold a huge amount of both codified knowledge in various databases and tacit knowledge in the minds of employees. LLM-powered tools can help to codify tacit knowledge—say, transcribing and capturing recorded meetings and other communications (with the permission of the participants, of course). Tools similar to the publicly available research products previously mentioned can then help product development practitioners find relevant corporate knowledge, which can be combined with externally sourced data to generate syntheses and insights.

Automating documentation tasks. In many product development processes, particularly in highly regulated industries such as pharmaceuticals and aircraft manufacturing, there are significant documentation requirements—for example, for regulatory filings, engineering change orders, and other required documentation. LLMs can accelerate the process of both generating and synthesizing these documents. (Of course, systems must be put in place, including human review, to ensure that these documents meet requirements for accuracy and fidelity.)

Collaborating with humans for ideation and concept development. Product managers, scientists, engineers, designers, and other participants in the product development process can "converse" with LLMs to stimulate ideas, get "opinions," and have their ideas challenged, much as they would with a colleague. These experiences illustrate that it is possible for humans and AI to collaborate, but the human skill in using AI tools can significantly influence the effectiveness of these collaborations (see sidebar, "Agents in R&D").

Estimating the economic potential of using AI to accelerate R&D

Our research finds that AI could substantially accelerate R&D processes across a set of industries that make up 80 percent of large corporate R&D expenditures. For industries whose products consist of intellectual property (IP) or whose R&D processes are closest to scientific discovery, the rate of innovation could potentially be doubled. For industries that produce complex manufactured products, R&D processes could be accelerated by 20 to 80 percent, depending on the industry (Exhibit 4). Overall, our analysis estimates that the potential annual economic value that could be unlocked by using AI to accelerate R&D innovation is about $360 billion to $560 billion. Next, we examine how this value capture could potentially play out across a range of different industry sectors.

Agents in R&D

AI developers have begun to create gen AI–enabled agents that use foundation models to execute complex, multistep workflows. The application of AI agents in R&D workflows has captured a lot of interest. Imagine agents able to plan and manage an entire testing and validation process: identifying the candidate variants to be tested and the parameters under which the tests should occur; executing the tests (potentially aided by physical robotics if using physical prototypes); and then optimally iteratively adjusting the next set of tests in a closed-loop, active learning cycle. Such agents could accelerate innovation even further beyond the opportunities we have documented to make each of these individual steps faster and more effective.

Exhibit 4

AI can boost R&D throughput by accelerating design generation, research operations, and design evaluation.

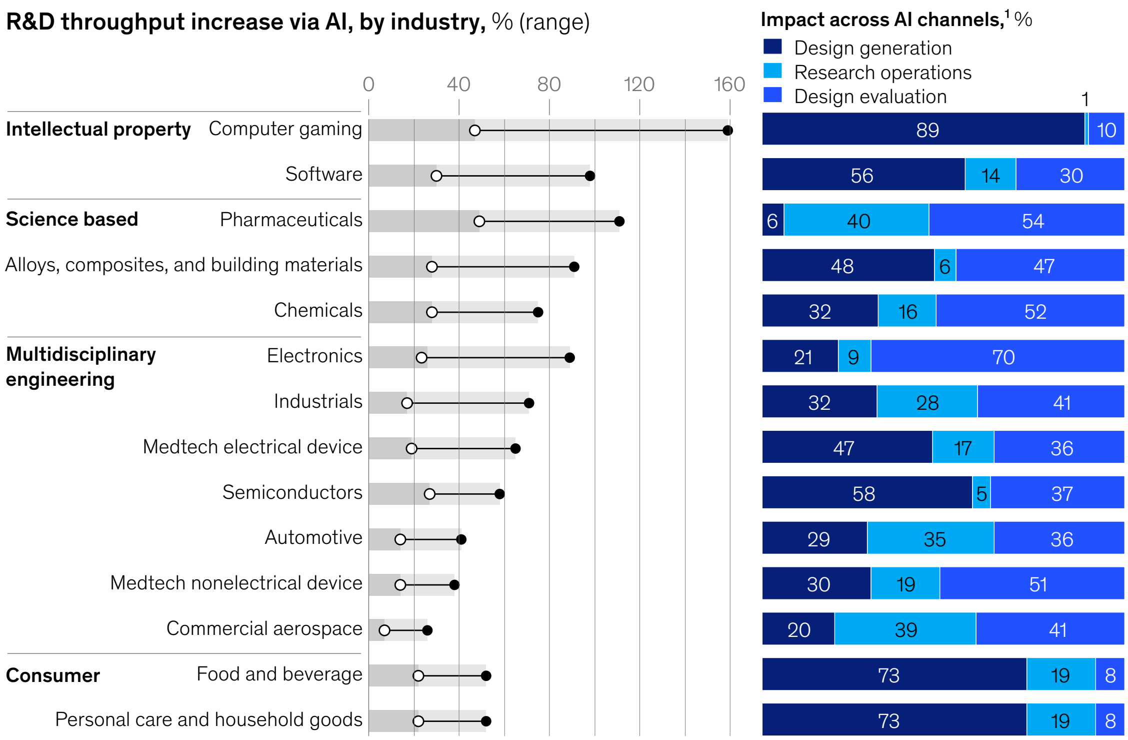

[1]Figures may not sum to 100%, because of rounding.

McKinsey & Company

IP product industries

Computer gaming and software are industries in which products fundamentally consist of IP. With no need for physical prototyping or manufacturing, all the AI-driven R&D acceleration levers can be directly applied. Furthermore, the application of gen AI for developing software and creating visual content are among the most mature use cases for that technology, which could lead to a potential doubling of R&D throughput, or even more. The bulk of the impact of AI in these industries stems from accelerating the process of generating designs—in this case, computer code and game visualizations.

The generation of design candidates in these industries is particularly well matched to the most advanced capabilities of gen AI models. Frontier gen AI foundation models have been advancing rapidly in their ability to generate computer code. At the time of writing, the CEOs of Google and Microsoft have both estimated that 30 percent of the new code produced at their companies was written by AI.

That said, while these generative-software-development tools are now widely available, the actual acceleration we have observed in many companies has been modest so far (although some start-ups have reported truly remarkable levels of software development productivity). This suggests that unlocking the full potential of AI will require a set of complementary organizational changes, above and beyond the technology itself—a theme to which we will return. But if these tools and the complementary organizational changes are put in place, we estimate that the software industry could double the rate at which it produces new products.

Another significant portion of the work in developing computer games, particularly those that are immersive, is the design and rendering of the virtual worlds in which game play takes place. Some of the earliest applications of gen AI, even before LLMs became widely available, consisted of being able to generate images. The rapidly evolving capability to generate visual content can directly accelerate the content creation process for computer games. Although the gaming industry is small, this illustrates that combining the ability to generate software and create content could increase its output by 150 percent.

Science-based product industries

Another set of industries where AI can accelerate innovation are those where the product development process is very close to scientific discovery, including industries like pharmaceuticals, chemicals, and alloys, composites, and building materials.[11]

Leading companies in the pharmaceutical industry have already been deploying AI in their R&D processes. These companies are training, adapting, and customizing foundation models for omics-based target identification (determining what molecular processes in a disease could be modulated to mitigate its effects) and in silico molecule design of drug candidates. AI surrogate models are also being used for in silico screening, molecular optimization based on structure and property predictions, and potentially preclinical analyses of pharmacokinetics (what a body does to a drug) and pharmacodynamics (what a drug does to a body). Other applications of AI in the drug discovery process include the ability to mine the extensive databases and literature in the field and to apply techniques such as computer vision to enable high-throughput experimental screening.

---

[11] "Scientific AI: Unlocking the next frontier of R&D productivity," McKinsey, May 6, 2025.

AI also has the potential to reduce the total cost of drug discovery and increase the probability of success of candidates that reach the clinical trial stage—that is, the likelihood that a candidate will be approved through the clinical trial process. (While we estimate the relative impact of generating new candidates on the overall timeline of drug discovery to be relatively small—about 5 percent—generating higher-quality candidates could improve the probability of success.)

It is important to note that the actual potential for increasing the number of drug candidates that become approved therapies is constrained by the clinical trial process, which has its own challenges (including cost, patient recruitment, and clinical and regulatory capacity). While AI could potentially accelerate this process, we only examined the drug discovery process within the scope of this study.

The potential to accelerate the R&D process is also high in industries that produce materials used as inputs into other industries, such as chemicals and alloys, composites, and building materials. AI surrogate simulations—for example, for physicochemical modeling—can be used for property prediction and analysis in that it helps to predict properties such as structure, strength, toughness, ductility, permeability, conductivity and resistance, and corrosion, depending on the type of material and its intended application. These techniques can also be used to optimize the processes for synthesizing/manufacturing these materials. As in pharmaceuticals, LLMs can be used for market analysis and to mine scientific literature and databases during the initial conceptualization and specification phase. And in all these industries in which some degree of physical experimentation is still required (in addition to in silico simulations), there is a potential for agentic AI to automate the process of managing experiments, though this capability remains nascent.

Overall, the top end of the ranges of throughput acceleration ranges from 75 percent for chemicals R&D to more than 100 percent for pharmaceutical discovery.

Complex manufactured product industries requiring multidisciplinary engineering
There is a wide swath of industries in which the product design process requires a variety of engineering disciplines (for example, electronics, industrials, medical technology, semiconductors, automotive, and commercial aerospace). Designing a commercial aircraft or an automobile, for example, requires engineers who specialize in aerodynamics, as well as in areas such as structural dynamics, propulsion, and electrical systems, among other disciplines. In electronics, product development requires not only an understanding of electrical and electronic circuits (often including the intended and unintended effects of electromagnetic radiation) but also the ability to predict and manage the thermal properties of a product. And as software is increasingly embedded in and becoming a larger part of the value delivered by physical products, software engineering is becoming a critical capability.

Across these disciplines, AI-powered generative-design systems can create a set of design candidates, often more quickly and from a wider search space than would be considered without these techniques (akin to AlphaGo's Move 37). Multiphysics AI-style deep learning surrogate models (those that incorporate multiple modalities of analysis, such as structural, fluid dynamics, thermal, and electromagnetic) can be used to predict the performance characteristics of design candidates more quickly than other numerical simulation methods (finite element analysis, computational fluid dynamics, and electromagnetic modeling). As engineering organizations develop an understanding of the relative strengths and weaknesses for each of these methods, they can better allocate their simulation and testing efforts across traditional numerical simulations, deep-learning surrogates, and physical prototypes.

In some of these industries, meeting the documentation and reporting requirements is critical, especially those in which safety and regulatory considerations predominate, such as aerospace, auto, and medical technology. LLMs can assist in meeting those requirements in a timely and efficient way, provided their outputs can be appropriately validated. As in other industries, LLMs and predictive machine learning can also assist in the initial concept development phases with market research and crafting product specifications.

Overall, while these industries share many characteristics (such as complex supply chains and integrated materials and electrical and software designs to create finished products), they also represent a wide variety of different products and markets. This is reflected in our estimates of potential impacts of AI accelerating R&D. In electronics, for example, the pace could nearly double, while in commercial aerospace the potential impact is 25 percent. One shared area of potential impact across all industries requiring multidisciplinary engineering is in the process known as verification (did I build the system right?) and validation (did I build the right system?). Some form of verification and validation occurs in all these industries, accounting for as much as half of the R&D timeline. The potential to transition from physical prototyping and testing to in silico testing in verification and validation could be one of the largest potential levers for accelerating the entire innovation process (though this is sometimes gated by regulatory requirements).

Consumer goods

The applications of AI for accelerating R&D in consumer goods (such as food and beverages and personal-care and household goods) parallel those for analyzing market trends and generating design candidates. LLMs and analytical AI can be used to generate and synthesize data-driven insights to provide direction for new-product development. These levers are becoming increasingly valuable as the quantity and variety of digital data about consumers continues to grow. Specialized foundation models, for example, can generate candidate recipes for food and beverage, candidate formulations for cosmetics, and candidate designs for other product categories such as apparel and household goods.

The potential for using AI surrogates for modeling consumer preferences (versus actual consumer testing) lies largely in the ability to create “digital twins” of consumers. Our estimates for the current potential of using AI surrogates (in place of consumer testing) to accelerate R&D in consumer product industries are conservative, though the figures likely will increase. About three-quarters of the AI impact we estimated for these industries comes from the generation of new-product candidates.

How accelerating R&D with AI could boost earnings

Those AI-driven productivity gains could unlock real economic impact, both for the companies that produce new products and their customers. Overall, we estimate that $360 billion to $560 billion of potential annual economic potential could be unlocked from using AI to accelerate R&D (Exhibit 5). As a proxy for those economic benefits, we estimated the potential EBIT impact of incremental new products being developed under the acceleration scenarios modeled—assuming unconstrained demand, no additional bottlenecks after product

Exhibit 5

$360 billion to $560 billion of annual economic potential value could be unlocked by AI accelerating R&D in large product companies.

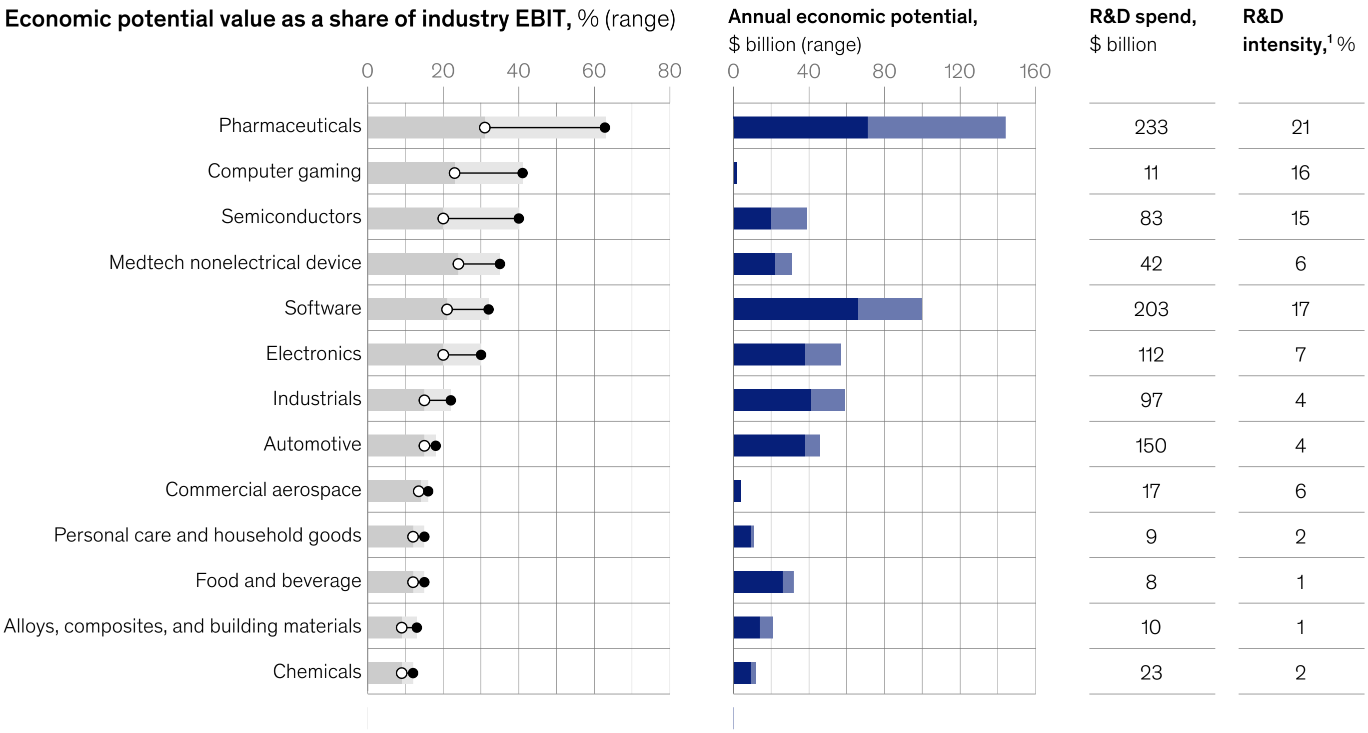

[1]R&D spend as a percentage of revenue.

McKinsey & Company

development (say, regulatory constraints such as clinical trials and supply chain or production capacity), and that new products will maintain current EBIT margins within each industry. We scaled this potential by the percentage of revenue accounted for by new products over the past five years. For the industries we studied, this would suggest a value at stake equivalent to a double-digit percentage increase in EBIT.

The industries likely to experience the greatest incremental economic potential (such as pharmaceuticals, semiconductors, and software) have high potential for accelerating their R&D processes by using AI. In other science-based materials industries (such as chemicals and alloys, composites, and building materials), the economic potential of using AI in R&D, expressed as a percentage of their current EBIT, is relatively lower, though still substantial. This reflects the fact that in these industries, a substantial share of revenue comes from the sale of existing commodity products (the chemicals industry, for example, will continue to produce ethylene, ammonia, and the like). Absent transformational shifts in those industries (which sometimes do happen), the impact of accelerating R&D will likely be muted. Instead, such industries will likely benefit most from using AI to improve their production processes (which wasn't the focus of our study).

At the individual company level, these overall industry trends might not hold. Specialty chemicals companies, for example, tend to derive higher percentages of their revenue from new products than commodity producers do. Composites are a much newer class of materials than chemicals and, as a result, present considerable opportunities for innovation. Another observation is that industries with the greatest economic potential from the use of AI to accelerate innovation are those with the highest R&D intensity (that is, R&D expenditures as a percentage of revenue). These are also industries with the strongest reputations for innovation. While perhaps unsurprising in hindsight, the high R&D intensity reflects investments in areas that hold the highest potential for improving their companies' bottom lines.

All that said, it is unlikely that all the economic value we have sized will be captured as corporate profit, even in the industries that introduce more novel products. Not only is demand constrained, but in general, many of the benefits that companies deliver is actually captured by customers in the form of higher-value products. While this constitutes genuine economic value, only a fraction of it is captured as profit by the producers of those products.

Overall, viewed in terms of economic value, our estimates could be considered conservative, as our analysis is focused on the value of generating more products through accelerated R&D while assuming the same margins as previous products in those industries. We have not attempted to quantify the value of AI enabling the design of higher-value products that could be developed, nor the potential lower R&D costs. Our estimates are similarly conservative in terms of the degree to which physical testing can be replaced with AI surrogate simulations. Nor have we attempted to estimate the value of truly breakthrough innovations that transform markets (if, for example, nuclear fusion was to enable limitless, clean electricity production).

Most important, these figures do not capture the broader benefits that innovation can have in society. Oftentimes, the most transformative impact of innovations are the downstream effects they have in other areas: The classic example is all the new industries that were spawned by inventions created for space exploration. The value of some innovations isn't even best expressed through economics: The value of saving lives through healthcare innovations could be considered incalculable. And returning to the theme of human welfare, if AI can help to bend the curves of innovation productivity, that would improve the quality and duration of life of future generations.

What business leaders can do to harness the power of AI in R&D

Advances in science and technology, however impressive, won't move the needle alone. Realizing the potential of using AI to accelerate innovation will also require organizational changes. This section identifies four key levers that leaders must consider.

Move quickly and scale rapidly

While AI is rapidly advancing in its capabilities, all the levers we describe here are not only available today but are already being deployed in corporations. However, these technologies require time and focused effort to use effectively. Climbing this learning curve sooner—and faster—can help you to gain a competitive edge over others. In this way, speed is a strategy. But successful pilots are no guarantee of success. Too many companies end up in pilot purgatory. There are a set of practices that lead to capturing value at scale, and you should build those organizational muscles.

Rewire your organization beyond just tech

The organizational practices that underpin value capture at scale require AI and its associated tech stack, but they also encompass a much broader rewiring of your organization.[12] They include aligning strategy; building the right talent and organizational models; agile delivery, adoption, and scaling; and change management and governance. Even without AI, there can be significant differences in the R&D throughput between global competitors in a single industry. Rewiring your R&D organization for AI can be an opportunity to achieve a step change improvement in performance.

One example of a specific organizational shift that can help to unlock AI acceleration in R&D is putting the groups that do prototyping/testing and those responsible for simulations, which often are separated in today's organizations, into one unified organization. As technology advances, deciding when physical testing is required and what can be done in silico (whether with numerical methods or AI surrogates) and in sequence should be a set of holistic decisions under one management.

In addition to transforming and streamlining processes within R&D, it's also important to note that the path to creating impact could be constrained by factors outside the R&D organization—or even outside the company. For instance, we discussed the constraints that the clinical trial process has on bringing new healthcare therapies to the clinic; sometimes work in other areas (including the use of AI) is required to more fully capture the value of accelerating R&D.

Build a core competency around models

The AI models that are used to create and evaluate design candidates are critical to using AI to accelerate the R&D process. Thus, a new critical core competency will be evaluating, integrating, training/adapting, and making build-versus-buy decisions about models, including open-source models, procured models, and even internally trained models, as part of the R&D process. This will also require deepening data sourcing and engineering capabilities.

Be thoughtful about incorporating humans in the loop

While there are intriguing scenarios in which R&D is fully automated, for the foreseeable future in most of the industries we have analyzed, people will still have roles in the R&D process. But those roles are likely to shift considerably in an AI-enabled future, requiring reskilling. Organizations will have to identify when it is critical to have a human in the loop, for example, to ensure safety or to sign off on various decisions where having an accountable individual is critical. And understanding how the deployment of these technologies affects the employee value proposition and employee working experience—for example, do they feel it "gives them superpowers" or like they "serve the machine?"—is important to recruiting and retaining the best talent.

---

[12] Eric Lamarre, Kate Smaje, and Rodney Zemmel, *Rewired: The McKinsey Guide to Outcompeting in the Age of Digital and AI*, New York: Wiley, 2023.

In an economy driven by innovation, there may be no more potent currency than fresh ideas to explore. Deployed strategically and accompanied the key organizational shifts previously outlined, today's emerging AI capabilities promise to unlock new pathways toward the growth, progress, and prosperity that characterized the previous century—but only if leaders embrace this new era of imagination and act now.

Alex Singla is a senior partner in McKinsey's Chicago office, **Alexander Sukharevsky** is a senior partner in the London office, **Elia Berteletti** is a partner in the Seattle office, and **Lareina Yee** is a senior partner in the Bay Area office, where **Michael Chui** is a senior fellow.

The authors wish to thank Arun Mittal, Arvind Venkataraman, Dean Deng, Holli Dobay, Keertan Kini, Neenu Babu, Nine Chantawasinkul, Rob Konkel, and Tiffany Walther for serving on the project team.

They also wish to thank Alessandro Mattozzi, Alex Peluffo, Aleyse McNealy, Anna Herlt, Anshul Sinha, Ben Meigs, Bill Wiseman, Brendan Gaffey, Brittany Presten, Chris Anagnostopoulos, Clemens Cepnik, David Champagne, David Naney, Delphine Nain Zurkiya, Diana Tang, Elea Medina, Germán Marín Estañ, Humayun Tai, Jacomo Corbo, Jan Paul Stein, Jeffrey Algazy, Jennifer Hou, Jeremy Wallach, Jordan Bank, Katarzyna Smietana, Kimberly Borden, Levent Tuter, Lieven Van der Veken, Martin Harrysson, Nathan Flesher, Nimit Patel, Nina Hirsens, Obi Ezekoye, Patrick Summers, Paul Chi, Pengfei Zhan, Peter Cholewinski, Philipp Kampshoff, Prakhar Dixit, Raj Rajendran, Rob Loughlin, Robert Ferris, Robert Linden, Sajal Kohli, Sara Cinnamon, Sebastian Göke, Stephan Lidel, Steven Prast, Tara Balakrishnan, Thao Dürschlag, Thomas Kilroy, Tore Johnston, Varun Marya, and Yann Phoon for their contributions to this article.

This article was edited by Larry Kanter, a senior editor in the New York office.

Designed by McKinsey Global Publishing