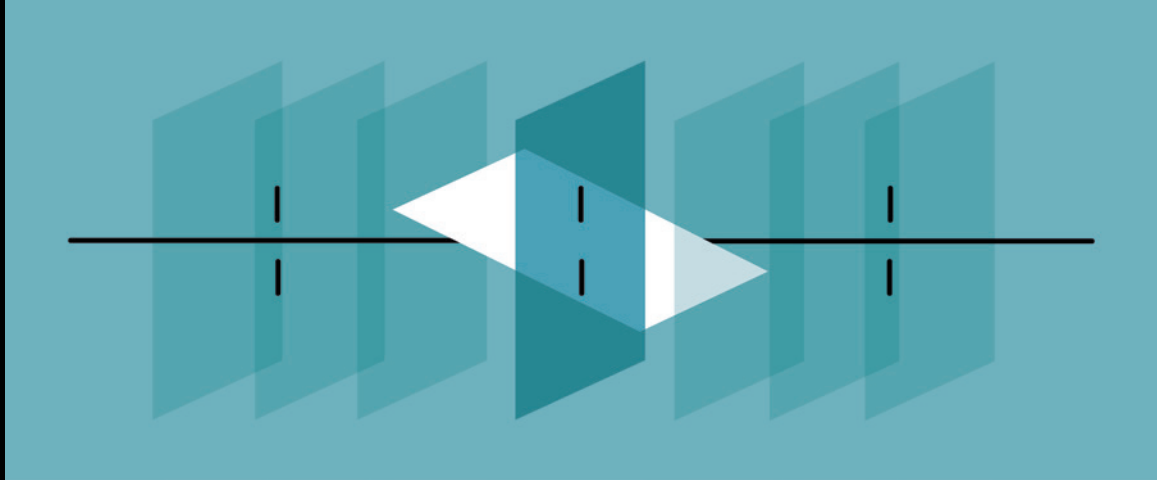
HARVARD BUSINESS REVIEW



Physical AI Gains Ground in Health Care and Life Sciences

HBR Analytic Services **BRIEFING PAPER**

Sponsored by

SoftServe

SPONSOR PERSPECTIVE

The Physical AI Shift in Life Sciences

Every life sciences executive knows what generative AI does. It analyzes, summarizes, and recommends. A person still decides what happens next. But one difficult and consequential question is what happens when AI stops advising and starts acting: guiding a robotic arm through a variable lab process, adjusting a scan in real time, or stepping in only when a clinician needs it. That's physical AI. It changes the calculus for every biopharma and medtech executive deciding where automation ends and clinical judgment begins.

The industries we serve don't get to choose between speed and safety. Every efficiency gained in a lab or on a device must survive regulatory review, clinical scrutiny, and the reality of a patient at the other end of the decision. That is why physical AI's real opportunity in life sciences doesn't replace scientists or clinicians. It returns their time. Fewer hours lost to repetitive, variable tasks. More spent on judgment calls only people can make. Technology succeeds or fails on that distinction.

Getting there takes more than a capable algorithm. It takes validated performance, bounded autonomy, and connected intelligence. Across health care and life sciences, we hear the same question. Decision makers have stopped asking what AI concludes. They want to know whether that conclusion can be used reliably, safely, and transparently inside a regulated environment. We build that same discipline into our own robotics labs in Lviv, Ukraine, and Wroclaw, Poland, where a global team of specialists works hands-on with humanoid robots for automating warehouses, industrial arms to enable dark labs, and digital twins across the processes to prove automation before implementation—building and validating these systems before they reach production. The organizations earning the trust of regulators, clinicians, and patients are the ones that treat physical AI as an ongoing engineering discipline, sustained beyond one milestone.

As a global, established leader in physical AI, SoftServe sponsored this report in association with Harvard Business Review Analytic Services because life sciences executives deserve a clear view of where physical AI has already proved itself, where it still falls short, and what it will take to close the distance between the two.

Those that get this right won't be the ones that move the fastest. They will be the ones that turn earned trust into measurably better outcomes for patients.



Ashish Mishra
Executive Vice President
Segment Head, Health Care and Life Sciences
SoftServe



Vishal Goyal, PhD
Senior Vice President
HCLS Business Consulting Services
SoftServe

Physical AI Gains Ground in Health Care and Life Sciences

The life sciences are heavily regulated industries under relentless pressure to improve clinical outcomes while holding down costs. In recent years, biopharma and medtech companies have pushed AI beyond analytics and digital workflows and into physical processes: robotic systems that run and refine laboratory experiments and devices that assist clinicians during procedures in settings that change minute to minute.

GENERATIVE AI CAN mine trial data or respond to inquiries in a chat window, providing recommendations that a human can choose to accept, modify, or disregard. In contrast, physical AI uses sensory inputs and algorithms to determine what machines and medical devices actually do. It provides the logic for a robot navigating a laboratory to perform DNA sequencing, and it can also support decisions within MRI systems about which data to acquire, help CT scanners optimize radiation dose, and guide ultrasound operators toward the correct view. But it's not just one technology. These complex applications may apply AI, computer vision, deep-learning reconstruction, and signal-processing algorithms. The system constitutes physical AI when it applies various inputs to directly control or alter a physical process.

"Health care has been building physical AI for years, even if we don't always use the term," says Dr. Taha Kass-Hout, global chief science and technology officer at GE HealthCare, a Chicago-based health care solutions provider. "Think about all the AI that already runs on medical devices—whether it is improving image quality, guiding image acquisition, supporting clinicians, or making sure equipment is healthy and functioning. Now sensors, edge compute, cloud, and connectivity are coming together. That is the next phase."

The life sciences sector has automated in waves, and most of those waves stopped at discrete, highly structured tasks. The systems worked

HIGHLIGHTS

Physical AI is turning medical devices, wearables, and robots into **adaptive systems that sense patients and their surroundings, read conditions in real time,** and respond with precise physical action.

Physical AI creates strategic value by addressing meaningful problems, **redesigning workflows, and distributing work between humans and machines.**

Despite technical, regulatory, and cultural hurdles, physical AI is already showing up in **patient-facing devices like smart prosthetics and in biotech applications** such as automated laboratory equipment.

> **“Health care has been building physical AI for years, even if we don’t always use the term.”**
>
> Dr. Taha Kass-Hout, global chief science and technology officer, GE HealthCare

when processes were standardized. They struggled with variable samples, changing experiments, and the exceptions that sent a human back to the bench. The result was islands of automation, not connected workflows.

Now AI arrives alongside machine vision, better robotics, sensors, simulation, high-performance computing, and lab infrastructure that actually shares data. Systems built on that stack can perceive what’s in front of them, adapt when conditions shift, and act across a workflow instead of running a script. As a result, physical AI can automate a 50% broader scope of work than traditional robotics, according to a May 2026 article by Boston Consulting Group, and the firm estimates that it has cut 70% of the engineering time needed to train robots. The group also estimates a three-times faster payback period for cutting-edge robotics investments, which it says has gone from approximately five to seven years to one to three years.[1]

Much is at stake. A McKinsey & Company report, “Technology Trends Outlook 2026,” released in September, projects a boom in economic benefits for “the broader physical AI ecosystem,” as it “unlocks $1 trillion or more in economic value by 2035, with general-purpose robotics scaling to $370 billion by 2040.” McKinsey adds that “the software and model layers of physical AI alone could represent $330 billion by 2036.”[2]

Yet the technology must prove itself to regulators even as it gains traction in the life sciences and other sectors. Governments must work out how to encourage investment in medical devices, robotics, and AI while also regulating the technology. The U.S. Food & Drug Administration (FDA), which is charged with ensuring AI-enabled medtech devices operate safely and effectively, granted clearance in 2025 to 295 of these devices.[3] Yet the agency says its evaluation methods are “not designed to predict behavior in dynamic, real-world environments.”[4]

That is no minor footnote. It helps explain why life sciences’ adoption of physical AI applications is concentrated in highly controlled settings such as the lab bench, the manufacturing line, the imaging lab, and the logistics floor. The closer robotic and diagnostic technologies get to patient intervention, the slower the adoption. As physical AI enters unpredictable environments, success will depend on earning and keeping the trust of clinicians, regulators, and hospital buyers. Each may demand quantitative evidence of its value and safety.

“Trust begins with evidence,” says Jay Lee, professor and director of the Industrial AI Center in the Department of Mechanical Engineering at the University of Maryland. “No AI system or human can be expected to be correct under every condition. The key is to validate performance against prior evidence, define where the system can be trusted, and continuously improve it through real-world feedback. Reliability should not be a one-time certification; it should be a continuous learning and validation process.”

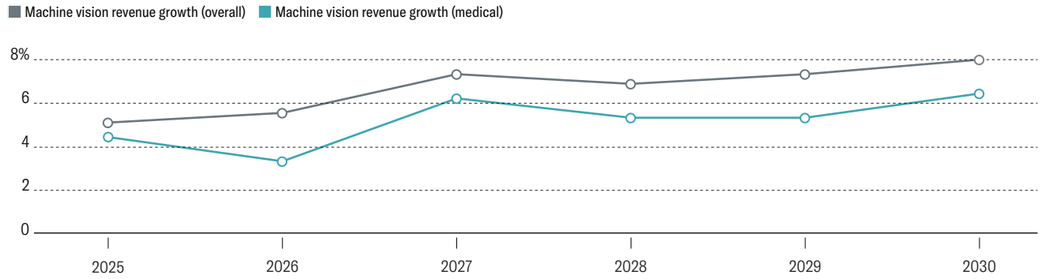
Despite technical, regulatory, and cultural hurdles, physical AI is already showing up in patient-facing devices like smart prosthetics and in biotech applications such as automated laboratory equipment. In the first half of 2026, robot orders in the life sciences, pharmaceutical, and biomedical sectors rose 32% over the same period the previous year, according to the Association for Advancing Automation (A3), a nonprofit trade group for the automation sector based in Ann Arbor, Mich. The global robotic surgery market is projected to climb at a 10% compound annual growth rate (CAGR) between 2025 and 2032, according to MMR Statistics, a market research firm based in Pune, India.[5] Demand is growing, MMR reports, because robotic platforms can reduce hospital stays by one or two days and cut postoperative complications, “outcomes that directly affect hospital cost structures.”

Adopters are applying the technology to repetitive data-rich tasks with measurable cost-saving objectives. Biotech companies and device makers are buying robotics, digital twins, machine vision, and AI-enabled automation to raise productivity and quality and to make their operations harder to disrupt. These efforts may include automated visual inspection on fill-finish lines, robotic liquid handling in labs, autonomous intralogistics in warehouses, and

FIGURE 1

Medical Machine Vision Growth Trails the Overall Market

The overall market for machine vision technology is projected to rise at a 6.91% compound annual growth rate vs. a 5.4% CAGR for the medical field, where adoption may lag for regulatory reasons.



Source: Interact Analysis, Machine Vision 2026 study, May 2026

computer vision–based compliance monitoring. The medical machine vision market will grow at a 5.4% CAGR between 2025 and 2030, according to Interact Analysis, a market research firm based in Wellingborough, U.K. Machine vision is moving beyond traditional inspection into advanced biomedical applications such as AI diagnostics, pathology, and automation.[6] FIGURE 1

“Automation hardware has been deployed in life sciences for a long time,” says Alexander Shikany, executive vice president of A3. “What’s changing with physical AI is the intelligence behind that hardware. These systems can now perceive and adapt to greater variability, work more flexibly across smaller batches and changing processes, and take on tasks that traditional automation struggled to handle. That allows scientists to spend less time on repetitive processes and more time on the higher-value work that drives research and innovation.”

This report examines what physical AI deployments deliver in biotech and health care and what it takes to succeed. It also addresses the “sim-to-real” opportunity, ensuring models trained in digital twin simulations behave safely and consistently after they leave the simulation. Validation is especially critical in health care: Manufacturers need to ensure accountability, provide evidence, and manage risk. As physical AI enters commercial deployment in the medical device and biotech industries, intelligence isn’t the main obstacle. The hard part is achieving human-level reliability, safety, and cost discipline that hold up in unpredictable environments.

“Physical AI changes both the product and its life cycle,” says Kass-Hout of GE HealthCare. “The dual advantage comes from bringing together the hardware, software, AI, cybersecurity, clinical evidence, and service around the workflow from the beginning.”

Valuing Automation

Physical AI is rarely justified by task count alone. Its value lies in what those tasks make possible: faster research, more clinical capacity, less variability, and experts’ time freed from routine operational work. Job replacement is not the primary business case, either. As organizations evaluate physical AI’s capabilities against new or existing business needs, they must model how a deployment could generate measurable clinical, operational, or economic value.

When a biopharma company deploys physical AI to automate its own laboratory or production operation, it may evaluate its impact on throughput, yield, quality, and labor. Consider visual inspection on a fill-finish line, where every vial must be checked for particulates, cracks, and fill-level defects. An AI-powered vision system can inspect vials at line speed and apply consistent criteria without the fatigue and variability of manual inspection. The value is measurable in fewer false rejects, better detection of critical defects, faster lines, and shifting labor from routine inspection to exception handling and system oversight. The costs are measurable, too, including model validation, change

> **“For major changes in clinical or biotechnology workflows, we should look for step-change improvements—potentially two or three times better in the right applications.”**
>
> Jay Lee, professor and director, Industrial AI Center in the Department of Mechanical Engineering, University of Maryland

control, monitoring for performance drift, and workforce retraining. Although the technology is new, plant managers can evaluate the potential risks and desired returns using familiar metrics.

From a hospital’s perspective, buying a physical AI–enabled medical device is a more complicated investment than buying automation for its own operations. The hospital must assess not only whether the technology is safe and clinically effective but also whether it will generate enough value to justify its cost. That calculation can be difficult because the hospital often does not control the payment associated with its use. Insurers generally reimburse the procedure rather than the specific technology used to perform it, leaving the hospital to capture the benefits indirectly through better outcomes, fewer complications, greater efficiency, or increased procedural capacity.

Physical AI creates strategic value by addressing meaningful problems, redesigning workflows, and distributing work between humans and machines. Organizations using the technology give machines an appropriate level of “bounded” autonomy, allowing them to act within predefined operational limits. GE HealthCare’s AI-powered Caption Guidance system for cardiac ultrasound illustrates the principle of bounded autonomy. The AI in an ultrasound device provides real-time guidance to clinicians on probe positioning and automatically captures an image after it reaches a diagnostic-quality threshold. The AI optimizes image acquisition while the technician runs the procedure and the cardiologist interprets the results.

By shifting a technically demanding part of the workflow to the AI-powered device without transferring clinical decision making, the system improves image quality and consistency. Well-designed deployments should therefore operate with bounded autonomy—giving machines greater autonomy where they demonstrably outperform humans while preserving human authority where judgment, accountability, and clinical context matter. Incremental gains may not fully justify the effort of physical AI deployments.

“The evidence has to be compelling, not simply a one or two percent improvement,” says Lee of the University of Maryland. “For major changes in clinical or biotechnology workflows, we should look for step-change improvements—potentially two or three times better in the right applications. The value must be measurable: lower cost, lower risk, shorter time, reduced workload, or less uncertainty in critical decisions.”

Shikany of A3 points to recurring sources of value such as speed, consistency, flexibility, safer work, traceability, and reuse. Physical AI matters most where variability makes conventional automation difficult, he says. “It’s particularly useful when the organization can reuse the resulting institutional capabilities across multiple products or processes.”

A well-structured proof of concept tests not only whether the system works but also whether the first deployment leaves behind something reusable such as data infrastructure, integration patterns, validated software components, or the operational expertise to run it repeatedly.

“The device is always where you start when it comes to diagnosis and monitoring, but the way we look at it is the entire workflow—all the steps in the chain of events that need to happen around the patient,” Kass-Hout says. “A radiologist or specialist can assess and interpret an image with the aid of AI for checking structures, detecting lesions, measuring them, and looking over time to see whether a lesion has grown or shrunk. AI can then help capture and structure that information and draft a first report. The broader value is more consistent exams, less repeat work, faster interpretation, better equipment availability, and broader access.”

Workflow optimization also extends to life sciences R&D and manufacturing, Shikany says. His list includes sample handling, instrument loading, assay preparation, precision assembly, inspection, and material movement in clean rooms. These applications matter, he says, because they combine repetitive work with enough product or process variation that conventional fixed automation may be hard to justify—or may not reach sufficient volume to scale in R&D.

Adopters must weigh multiple objectives. “The goal is not to claim that we can remove doctors,” Lee says. “We should say that we have a system that can continuously improve, reduce waste and uncertainty, and create a better decision window. If doctors do not have to spend two hours every day on repetitive problems, they can spend those two hours with patients.”

The Road to Autonomy

Robotics is perhaps the most discussed aspect of physical AI, particularly in health care, where talent shortages affect critical patient care and diagnostic roles. Unfortunately, robots aren’t coming anytime soon to plug those gaps, and physical AI and robotics aren’t synonymous. To plug those gaps in clinical labor, physical AI combines AI-driven decision making with machines and devices that sense and act in the real world, though they typically aren’t designed to walk out the door of their own volition. Robotics, by contrast, is the broader field of designing and operating machines that perform physical tasks. Many robots rely on fixed programming, predefined rules, or human control rather than AI. A robotic arm that repeatedly moves a vial from one position to another can be highly automated without being AI-powered.

“When the average person thinks of physical AI, they probably picture a humanoid robot,” Shikany says. “That’s not exclusively what physical AI represents. It can also mean applying AI in an automation system that includes vision systems, robots, grippers, conveyors, and other technologies working together.”

In regulatory and industry circles, AI safety, including cybersecurity and privacy, is among the chief policy issues. Machine safety is another key consideration, particularly how organizations determine and limit an AI-powered device’s autonomy when human workers or medical patients are present. Unrestricted autonomy is unlikely on capability grounds alone, and even if not, insurers, industry standards, and state or federal rules may forbid it outright. Anyone building or deploying an AI-powered device must consider appropriate levels of autonomy, the organization’s risk tolerance, and how to retain human oversight.

“We think about this as bounded AI,” Kass-Hout says. “Systems increasingly need to perceive the environment they are in, guide positioning, automate routine steps, check quality, and adapt within defined clinical and safety limits. There are checks around what needs to happen, when the system should escalate to a human, and which steps it can repeat. These algorithms do not go into the wild and do things on their own without a control plane around them.”

PHYSICAL AI TRANSFORMS ROBOTICS

50%

Physical AI can automate a 50% broader scope of work than traditional robotics, according to Boston Consulting Group.

70%

Boston Consulting Group (BCG) estimates that it has cut 70% of the engineering time needed to train robots.

3x

BCG estimates a three-times faster payback period for cutting-edge robotics investments, from approximately five to seven years down to one to three years.

> “Responsible AI must be able to say how certain it is and, when necessary, stop and say, ‘I need a human to intervene.’”
>
> Kass-Hout of GE HealthCare

Reliability is a fundamental industry concern and, in some cases, remains a work in progress. Shikany says some of the most mature applications today are those where the task, environment, product, and exception conditions are relatively well defined. He adds that even if conditions vary, applications such as vision-guided part handling, medical device inspection, and autonomous material movement can meet demanding requirements when properly engineered and validated. “As the time horizon expands and the tasks turn into workflows or unstructured environments, the necessary [operational] technology remains emerging,” Shikany says.

Near-flawless uptime is the minimum requirement for AI-powered devices in regulated health care environments and biotech labs, and it is not the only prerequisite. The systems must be explainable (able to support their actions transparently), monitored over time, secure, traceable, and, when faced with uncertainty, able to escalate decisions to human experts.

“This is not something you deploy once and forget about,” says Kass-Hout. “You have to continuously monitor these algorithms after deployment. We need evidence that performance holds across many populations and sites, predictable change control as models evolve, and cybersecurity and privacy protection throughout the algorithm’s life cycle. Responsible AI must be able to say how certain it is and, when necessary, stop and say, ‘I need a human to intervene.’”

Trust is not something a system earns by passing a few tests. Developers must define the system’s intended use and operating boundaries, test the system not only in normal conditions but also in edge cases and foreseeable failures, and monitor its performance across the system’s life cycle. In physical AI, trust ultimately comes not from assuming the system will always be right but from designing it to behave predictably when it is uncertain or wrong. Shikany says engendering trust also requires smart cybersecurity, change control, performance monitoring, safe fallbacks, and defined points of human review. He adds, “You have to look at the complete system and make sure it has been appropriately tested and validated under operating conditions where it’s actually going to operate.”

As physical AI deployments earn their place in biotechnology and clinical environments, Lee believes their greatest long-term impact may come from creating closed-loop learning in which real-world experience continuously improves future decisions. Health care already generates large amounts of physiological data. Connecting this data with treatments, patient characteristics, and clinical outcomes could improve drug development, treatment selection, and personalized care.

“Consider an ICU patient connected to multiple sensors,” Lee says. “When a drug is administered, we may observe changes in blood pressure, temperature, respiratory function, and many other physiological signals. If AI can connect the treatment, the patient’s characteristics, and the physiological response with the eventual clinical outcome, we can begin learning which treatments work best, for which patients, and under which conditions.”

The Key Role of Digital Twins

A digital twin is a computational model that represents a physical entity or system, such as a patient’s anatomy, a scanner, a prosthetic device, or a care pathway. This model is calibrated using data and is continuously updated as that data changes. Organizations can use digital twins to model their operational equipment or workflows. Such use lets them simulate scenarios and assess the impact of changing conditions and capabilities.

Digital twins also help organizations decide whether to proceed with their physical AI investments. Through simulations, they can identify flawed assumptions or engineering mistakes before deployment, improve risk assessments, and build stakeholder trust at every stage. The net effect is that errors surface earlier, when they cost less to fix.

“Digital twins are one of our strategies for scaling responsible AI safely before it reaches clinical practice,” says Kass-Hout. “In the real world, you cannot assume that you can address every edge case. Digital twins allow us to simulate many of those scenarios using synthetic data. We can generate images that look like they were produced by mammography, test whether the technology detects the relevant features, and continue to improve it.”

Digital twins are increasingly useful tools for integrating physical AI and automating processes, according to Shikany. “It’s also much better for the users,” he says, “because they can get a picture of how this [AI-powered device] might help them.” Shikany contrasts this with the more traditional deployment model. “It used to be you’d have to build the [robotic or physical AI] system in your facility, bring the customers in and walk them through it, [and] make sure it all worked properly, and then you’d have to disassemble it, ship it, and set it up again” at the customer’s location.

Physical AI can encounter failures in clinical and laboratory settings when real-world conditions differ from those anticipated during development. If equipment or environmental changes disrupt workflows, the application may miss the mark. A high-fidelity digital twin can narrow that “simulation-to-reality” gap by letting organizations test physical AI systems across a range of clinical and laboratory scenarios before deployment.

In health care, that application could involve evaluating how an AI-enabled robotic or imaging system responds to variations in patient anatomy, positioning, imaging conditions, or procedural steps. In laboratory settings, it could include simulating different sample characteristics, instrument configurations, or workflow interruptions. By identifying potential failure modes and edge cases in a simulated environment, organizations adopting physical AI can safely and efficiently validate perception, movement, decision making, and human handoffs before the system interacts with patients, manages valuable samples, or enters a regulated production environment.

As AI-based health care solutions approach field testing, digital twins can provide evidence that a device is ready for patient or clinician use. In 2018, FDA researchers reported on the Virtual Imaging Clinical Trial for Regulatory Evaluation, which simulated a conventional clinical trial involving human patients and radiologists.[7] This study created 2,986 simulated patients using computer-generated breast anatomy, lesions, imaging techniques, and computational image readers to compare the efficacy of digital mammography with digital breast tomosynthesis. The results matched those of a human clinical trial, leading FDA researchers to conclude that computer-simulated imaging trials could sometimes provide viable evidence for regulatory evaluation. By enabling simulations to test and validate the performance of complex systems in development, digital twins can reduce the cost and risk associated with real-world deployment.

Shikany expects greater use of digital twins and synthetic data to develop and challenge applications before they enter regulated production environments. But he offers one note of caution: “Simulation can accelerate development, but it doesn’t replace appropriate real-world testing, controlled pilots, site-acceptance testing, and validation.”

Conclusion

The future of physical AI in biotechnology and medical devices will depend less on frontier advancements in algorithms and robotics than on the harder work of making existing capabilities hold up in practice: achieving reliable real-world performance, validating systems within regulatory frameworks, building foundational data for training and monitoring, and earning the trust of those responsible for patient safety, product quality, and operational outcomes.

Physical AI is turning medical devices, wearables, and robots into adaptive systems that sense patients and their surroundings, read conditions in real time, and respond with precise physical action. In clinical care and surgery, these capabilities could enhance monitoring, personalize rehabilitation, improve procedural precision, and help clinicians detect and respond to complications earlier.

> **“The future is not about AI replacing people. It is about creating systems in which humans and AI continuously learn from each other.”**
>
> Lee of the University of Maryland

While fully autonomous surgery is still years away, many hospitals have been using robot-assisted surgery for more than two decades. These include systems that aid in radical prostatectomy procedures as well as those that guide surgeons, provide anatomical details, detect potential errors, and assist with instrument positioning. In such cases, the surgeon directs the procedure while the robot offers enhanced visualization, improved dexterity, and instrument control.

Computer vision lets a system sense a patient’s movement or physiological state and adjust robotic assistance in real time. The use of autonomous mobile robots in bedside logistics and hospital transport is rising, driven partly by the Covid-19 pandemic and technological advancements, according to a February 2026 study in the journal *Healthcare*, “Embodied Artificial Intelligence in Healthcare,” which examined the clinical impact of robotic perception and decision making.[8]

Teleoperated surgery, which lets surgeons perform procedures remotely, is growing. The constraints are latency, regulation, and liability, not feasibility. Physical AI can transform assistive technologies from programmable tools into adaptive systems. Advanced prosthetics can interpret a user’s intent and adjust movements in real time; exoskeletons can tailor assistance to an individual’s movements, environment, and physical condition; and sensor-driven proprioceptive systems can establish feedback loops between the human body and machines. In each instance, physical AI combines sensing, interpretation, and physical action to create a more responsive human-machine interface.

Shikany and other experts believe physical AI will become more widely deployed and multimodal, with simulation playing a larger role in development and deployment within five years. This trend is already evident in drug development and is expected to grow in that sector.

One of the most significant commercial developments may be the creation of pre-engineered platforms with reusable supplier documentation and qualifications, which can simplify site-specific processes. Shikany expects deployment to get easier, less through custom software engineering than through proven off-the-shelf solutions, reusable robot skills, and eventually natural language interfaces. Those interfaces could lower the engineering burden, he says, but they will not replace safety controls.

Kass-Hout says his organization is exploring autonomous imaging in X-ray and ultrasound. With about 4.5 billion imaging scans performed industry-wide each year, he says, improving consistency and reducing reliance on scarce expertise could have a significant impact. Using AI-driven cameras, also called computer vision, the goal is for future versions of these devices to position patients, take X-rays, perform quality checks, determine when to escalate to a human, and generate reports for radiologists to review and finalize remotely, he says.

Major transformations in drug development, health care delivery, and regulation may require coordinated changes across organizations and institutions. “A bottom-up approach can produce faster, incremental progress,” Lee says. “A top-down approach can enable transformation at much greater scale because it requires coordination across organizations, regulators, standards, and systems. Ultimately we need both.”

Looking further ahead, digital twins could evolve from models of machines and manufacturing processes into dynamic representations of an individual’s health. Such a digital twin could potentially make health care more proactive rather than just reactive. “The future is not about AI replacing people,” Lee says. “It is about creating systems in which humans and AI continuously learn from each other. If we can reduce uncertainty, provide earlier and better decision windows, and continuously improve from real-world evidence, physical AI can fundamentally change how we develop therapies, deliver health care, and improve human health.”

Endnotes

1 “The CEO’s Guide to Physical AI,” Boston Consulting Group, Tilman Buchner, Mei-Jung Chen, and Daniel Küpper, May 2026. https://www.bcg.com/publications/2026/the-ceos-guide-to-physical-ai

2 “Technology Trends Outlook 2026,” McKinsey & Company, September 2026. https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/the-top-trends-in-tech

3 “2025 Year in Review: AI/ML Medical Device 510(k) Clearances,” Innolitics, December 2025. https://innolitics.com/articles/year-in-review-ai-ml-medical-device-k-clearances/

4 “Request For Public Comment: Measuring and Evaluating Artificial Intelligence-enabled Medical Device Performance in the Real-World,” U.S. Food & Drug Administration, September 2025. https://www.fda.gov/medical-devices/digital-health-center-excellence/request-public-comment-measuring-and-evaluating-artificial-intelligence-enabled-medical-device

5 “Robotic Surgery Market Is No Longer a Technology Story. It Is a Discipline Test for Hospitals,” MMR Statistics, February 2026. https://www.prnewswire.co.uk/news-releases/robotic-surgery-market-is-no-longer-a-technology-story-it-is-a-discipline-test-for-hospitals-mmr-statistics-302679157.html

6 “Machine Vision 2026,” Interact Analysis, May 2026. https://interactanalysis.com/research-products/vision-technologies/

7 “Evaluation of Digital Breast Tomosynthesis as Replacement of Full-Field Digital Mammography Using an In Silico Imaging Trial,” *JAMA Network Open*, Aldo Badano, PhD; Christian G. Graff, PhD; Andreu Badal, PhD, November 2018. https://jamanetwork.com/journals/jamanetworkopen/fullarticle/2717000

8 “Embodied Artificial Intelligence in Healthcare: A Systematic Review of Robotic Perception, Decision-Making, and Clinical Impact,” *Healthcare*, Bilal Ahmad Mir, Dur E. Nishwa, and Seung Won Lee, February 2026. https://www.mdpi.com/2227-9032/14/5/572



VISIT US ONLINE

hbr.org/hbr-analytic-services

Harvard Business Review Analytic Services is an independent commercial research unit within Harvard Business Review Group, conducting research and comparative analysis on important management challenges and emerging business opportunities. Seeking to provide business intelligence and peer-group insight, each report is published based on the findings of original quantitative and/or qualitative research and analysis. Quantitative surveys are conducted with the HBR Advisory Council, HBR's global research panel, and qualitative research is conducted with senior business executives and subject-matter experts from within and beyond the *Harvard Business Review* author community. Email us at hbranalyticservices@hbr.org.

CRE81641026

Copyright © 2026 Harvard Business School Publishing.