servicenow

INNOVATION BRIEF

Physical AI

How AI and LLMs are supercharging robotics

WRITTEN BY

Eugene Chevyrov
Innovation Engineering

Nick Diaz
Innovation Engineering

Ian Krieger
APAC Innovation Officer

The Innovation Office

Physical AI

How AI and LLMs are supercharging robotics

Summary

- Generative AI is accelerating robotics development, leading to a new class of AI—physical AI—that is powering far more capable and versatile robots.
- AI is driving the emergence of robots both aware of their surroundings and capable of taking sequences of actions based on high-level instructions in natural language.
- In the enterprise, one of the most promising applications of physical AI will be combining it with digital AI, which will enable workflows to extend to the physical world.
- As physical AI advances, a platform approach will be vital to coordinating workflows across humans, digital AI agents, and physical AI agents to generate new business value.
- Physical AI will unlock new business possibilities by powering robots capable of handling complex tasks that will be a seamless part of operations, helping companies achieve new levels of efficiency, effectiveness, and profitability.
- Investing in physical AI and a platform approach to harnessing the technology today will allow businesses to be at the forefront of the next frontier in artificial intelligence.

We've always had big imaginations when it comes to what robots can do—from maids and personal assistants to brilliant colleagues and superheroes to future soldiers and Terminators. But the reality today is that robots have been long on promise but limited to applications such as factory lines, kitchens, warehouses, hospitality venues, and restaurants. While helpful and still emerging, today's robots are largely confined to single, repetitive tasks in controlled settings.

All of that is about to change. Advances in generative AI are leading to the development of a new class of AI known as physical AI. At its most basic, physical AI allows robots to interpret their surroundings and take actions based on that understanding.[1] Physical AI is helping overcome the historic barriers that have kept robots unaware, inflexible, and used only in very specific business applications.

Physical AI will unlock new capabilities that bring robotics into an entirely new class of applications such as surgical robots that adapt to variations during procedures, warehouse

fulfillment robots that could be deployed across sites, as well as agricultural robots that selectively harvest crops based on real-time quality assessment. More broadly, incorporating truly autonomous physical systems into the enterprise will revolutionize operations as businesses deploy robots to handle complex tasks, seamlessly coordinated with the rest of their operations, to drive new business value.

> **With physical AI, workflows are set to not just evolve, but get supercharged through robotics.**

The advancements in physical AI will help usher in a new era of possibilities for businesses. Generative AI is already generating massive productivity improvements by transforming workflows, including speeding up activities like code writing, automating processes like customer service requests, and generating content like digital marketing materials. Now with physical AI, workflows are set to not just evolve, but get supercharged through robotics.

In the future, digital and physical AI agents will work together to extend workflows to the physical world. Digital agents will help monitor a patient's vital signs while robots provide new levels of precision during complex surgeries. Robots will ensure store shelves are stocked following instructions given by their digital counterparts, who are analyzing real-time data and consumption patterns. And repair and maintenance robots will be activated by digital agents monitoring equipment 24/7 to ensure assembly lines keep running. These examples are just the tip of the iceberg of what's possible when digital and physical AI workflows join forces in tomorrow's enterprise. For businesses, investing in AI today will place you at the forefront of the physical AI revolution, ready to harness both digital and physical AI to drive future business value.

Industry leaders are making big bets on physical AI

Industry leaders already recognize the opportunity and as a result, there has been a surge of investment and increased strategic focus on physical AI recently. NVIDIA, Boston Dynamics, and Tesla are investing heavily and openly discussing the potential of physical AI and humanoid robotics.[2]

> **Estimates place the AI in robotics market at around $291.1 billion by 2035, up from an expected $55.1 billion in 2025.[7]**

OpenAI, after abandoning robotics in 2019, has restarted its physical AI program in 2024, presumably encouraged by the possibilities and advances in the field.[3] Tech entrepreneur and venture capitalist Marc Andreessen, in an interview, declared that, "the general purpose robotics thing is going to happen, and it is going to happen in the next decade, and it is going to happen on a giant scale, and I think there is a plausible argument…that robotics is going to be the biggest industry in the history of the planet."[4]

With a $40 billion valuation, Figure AI has recently become one of the hottest startups to watch in Silicon Valley.[5] The company makes general-purpose humanoid robots and the company is on a mission to create the world's first commercially viable autonomous bipedal robot. Recently, Figure AI made headlines when it placed several of its humanoid robots on the BMW factory floor and stated that it would put more than 200,000 robots on assembly lines and in homes by 2029.[6] Given their humanoid form and the promise of great dexterity, Figure AI robots are able to move around spaces designed for humans and use a wide array of tools designed for humans to use, with much higher precision and less frequent failure rates. This represents a breakthrough in robotics-assisted manufacturing, where the spaces and the tools will not need to be redesigned for physical AI to use, but instead, the physical AI will adapt to the constraints of the world around it and thrive within those constraints.

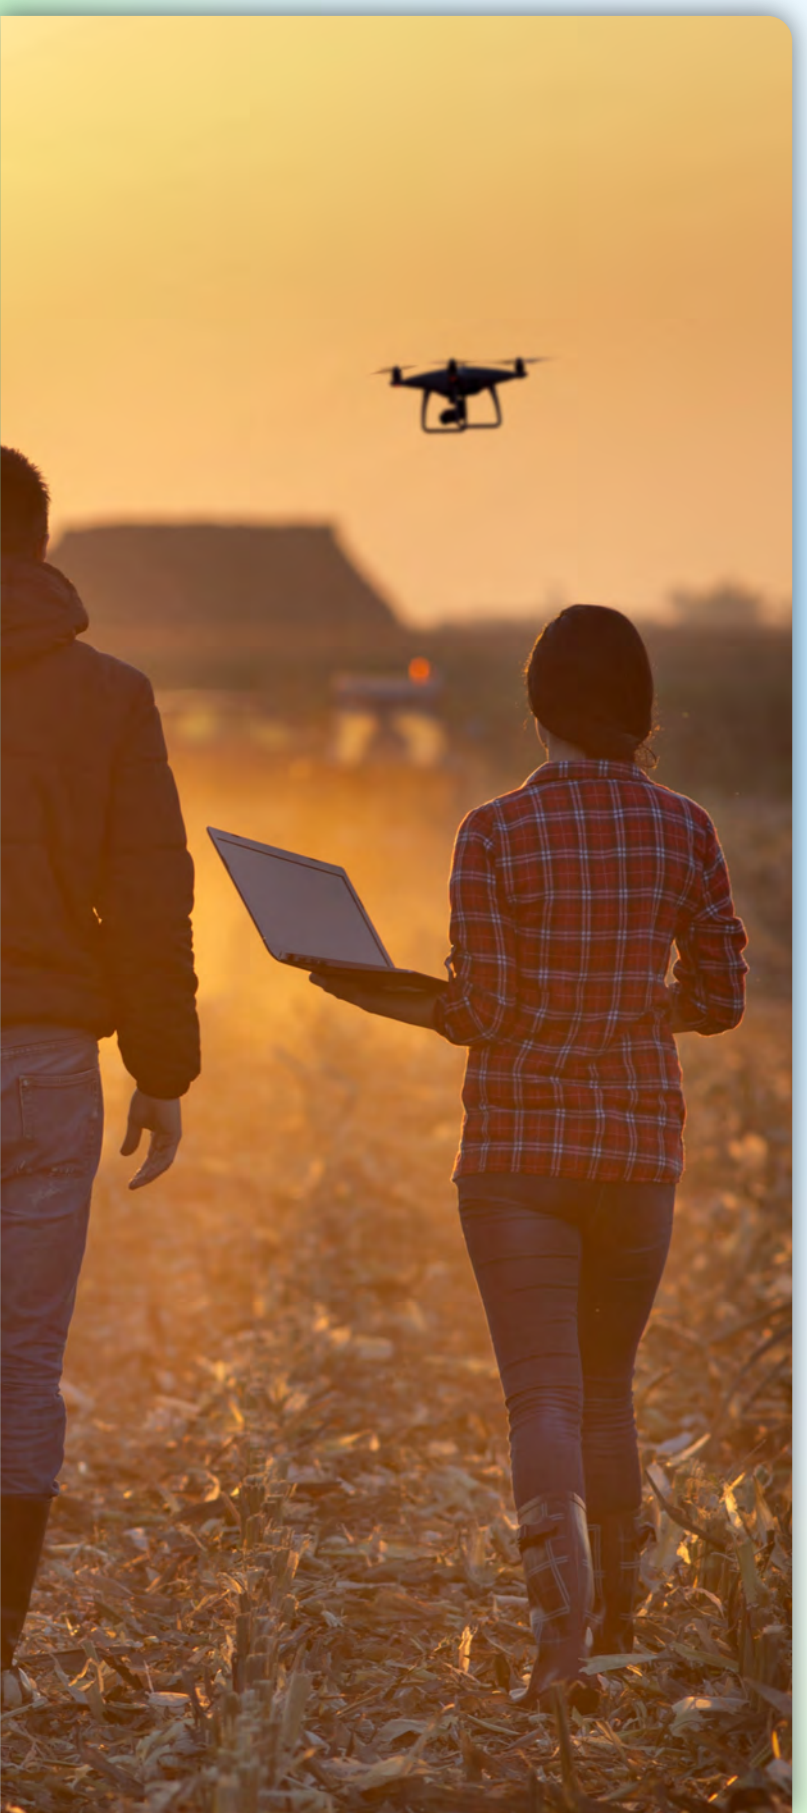

Optimistic estimates place the AI in robotics market at around $291.1 billion by 2035, up from an expected $55.1 billion in 2025—driven largely by increasing complexity and customization demands in manufacturing and logistics.[7] The business reasons driving this massive investment flow are increased confidence that robots will deliver measurable ROI through reduced labor costs, improved operational efficiency, and access to previously inaccessible capabilities. Beyond cost savings, robots will enable businesses to tackle challenges that have been impractical or dangerous, from handling hazardous materials in chemical processing to maintaining continuous 24/7 operations without human fatigue limitations. This combination of enhanced productivity and expanded operational boundaries will create new revenue streams while significantly reducing operational risks and insurance costs.

Today's investments and attention are one indication that physical AI will be the next frontier of AI, one that will continue to revolutionize work and workflows beyond AI agents and agentic AI.

AI: the robotics game changer

All of the hype and investment around physical AI is thanks to the breakthroughs that generative AI is driving in the robotics field. One of the major advantages of generative AI is that it processes multiple

modalities (text, image, voice) of input and transforms them into meaningful output. It empowers users with the ability to have an intelligent collaborator and even an autonomous counterpart, elevating and optimizing performance and productivity.

The same advances that let ChatGPT carry on a conversation with users toward productive outcomes or DALL·E create imagery are now evolving into physical embodiments that allow AI to interact with the physical world. GenAI is

helping robots understand the world, reason through tasks, and act—even in situations they've never seen before. This will result in robots, drones, autonomous vehicles, and other intelligent machines that can sense, move, and interact in real environments. For example, warehouse robots will adapt to reorganized shelves or unexpected obstacles without needing to be reprogrammed, using visual understanding and reasoning capabilities like those that power ChatGPT and other Large Language Models. Similarly, field service robots are beginning to assess infrastructure like bridges or pipelines, identifying potential issues and determining repair priorities by combining visual analysis with contextual reasoning, much like how AI models can examine complex documents and extract meaningful insights about structural patterns and relationships. These examples illustrate how far the robotics field is advancing thanks to AI.

> **AI is driving robotics beyond the previous limitations of pre-programmed instructions, evolving towards general-purpose intelligence that can adapt, interact, and learn. And most importantly, this form of AI can operate physically in the real world.**

Supercharging robot thinking and learning

AI is driving robotics beyond the previous limitations of pre-programmed instructions, evolving towards general-purpose intelligence that can adapt, interact, and learn. And most importantly, this form of AI can operate physically in the real world. We are at the very beginning of seeing physical AI agents become exponentially more useful and intelligent, both in businesses and at home. There are two major ways AI has transformed robotics development:

1. Vision + Language + Action (VLA) models: the new robot "brain"

One of the primary boosts robotics has received from generative AI is the development of Vision-Language-Action (VLA) models. These are multi-modal foundation models that allow robots to see, understand, reason and make plans within, and then act in their environment. VLAs are pretrained on both large-scale web data and real-world robotic demonstrations. They ingest a combination of visual input like camera images and natural spoken or written language, then translate that into robotic actions. Many VLAs work by transforming both the language and visual information into a shared representation, processing that representation, and outputting a sequence of robotic actions, movement trajectories, or even robotic code.

What makes VLA models so revolutionary for robotics is that they generate behavior based on understanding the real world, instead of on predefined robotic rules. Robots can use VLA models like Google's RT-2[8] to directly interpret a command like "pick up the apple" and translate it into robotic control of "move arm three quarters of the foot to the left, then straight down," even if they've never seen that exact instruction in a given physical environment before. What that means for businesses is that programmers no longer have to pre-program every robotic action in minute detail.

Unlike prior iterations of robotics, VLAs such as 3D-VLA[9] have contextual and spatial awareness of their 3D surroundings in three dimensions, allowing robots to execute complex open-ended commands in dynamic environments. This marks a major shift in robotic intelligence. Robots aren't just reacting to their environment—they're reasoning. This leap in autonomy opens the door to more general-purpose, adaptable robots, and potentially low-code or natural language programming for robotics.

Picture this: your enterprise manages a fleet of delivery robots across multiple office campuses. One day, you notice delays on a specific route. After a quick investigation, you find that a robot is struggling to climb a narrow, winding staircase, dragging its legs awkwardly and slowing deliveries. In the past, this might have required rewriting movement code or redesigning the route altogether. Now, thanks to emerging physical AI capabilities, you simply issue a natural language prompt: "Lift legs higher when climbing stairs with sharp turns." The robot interprets the instruction, adjusts its gait using its onboard Vision-Language-Action model, and seamlessly improves its performance. No code. No downtime. Just results. Within hours, delivery times improve, customer satisfaction rises—and you've optimized robotic performance with nothing more than a few words.

2. Synthetic training data and simulation: supercharging learning

One of the biggest—if not the biggest—reason for the success of generative AI was the availability of vast amounts of training data, aka the open internet. Unfortunately, there is no equivalent trove of readily available training data for autonomous systems on which to train VLAs. Robots learn by experience and producing real training data for robots to perform certain tasks is painstakingly slow and unscalable, requiring hours of real-world observations and trial and error.

Today, thanks to progress in generative AI and simulation technologies, we can use online images and text as stand-ins for real-world experience. Additionally, we can virtually generate the environments for autonomous

systems to train in, simulate robotic systems in those environments, and then capture the data needed for training those systems at scale. We are beginning to see initial successes in deploying these strategies.

One example of this new robot training approach is the success of 1X's humanoid robot autonomously performing domestic tidying tasks.[10] By simulating the household environment and data and applying a state-of-the art robotic reasoning model (GR00T N1 from NVIDIA), 1X was able to adapt a robot to specific chores like vacuuming quickly. Another example of this new training approach is a team of Stanford researchers who extended the NVIDIA Omniverse platform to create an environment specifically dedicated to training for and evaluating the performance of a thousand household activities—picking up toys, setting the table, cleaning floors—by robots.[11] After training, the platform that Stanford researchers built allows them to deploy and test robots in the real world.

All of these new efforts address the classic shortcomings of robots. In the past, a robot might perform flawlessly in the lab and then fail in a new environment because it had never seen that exact combination of physical surroundings and instructions before. With massive synthetic training, the aim is to expose the AI to "everything that could happen" in simulation—varied object shapes, lighting, clutter, edge cases—so that when the physical robot encounters reality, it's not fazed by minor deviations. This all hints at a future where training a robot for a new task or environment could be as straightforward as gathering a few demonstrations or spinning up more simulated trials, instead of re-engineering the entire system.

Imagine a robot is deployed on an airport tarmac to transport cargo between terminals. This is a chaotic environment where obstacles appear unexpectedly, weather conditions shift by the hour, and workers and vehicles are constantly moving in and out of its path. Previously, robots in such environments had to be preprogrammed with rigid paths and behaviors, limiting their usefulness. But now, thanks to simulation-trained physical AI, the robot has learned to navigate uncertainty. It

adjusts its trajectory on the fly when new equipment blocks its route. It slows down in heavy rain. It reroutes in real time based on natural language prompts like: "Prioritize safety near Gate B—more foot traffic today." Instead of needing weeks of programming and manual retraining, the robot adapts in real time because it's already seen thousands of similar scenarios in simulation. It doesn't panic, freeze, or crash. It adjusts and keeps moving.

The enterprise advantage

While generative AI, agentic AI, and now physical AI are advancing rapidly, the real breakthrough for the enterprise will come from the integration of both smarter software and more capable, intelligent robots. Physical and digital AI agents will work together to provide value for enterprise stakeholders. Digital AI like chatbots, workflow engines, and AI agents will continue to improve and excel at coordinating data, systems, and logic, which means enterprises will have faster decision-making, smarter workflows, and more streamlined operations. At the same time, physical AI—robots, drones, autonomous devices—will bring that digital intelligence into the real world, boosting productivity in warehouses, factories, hospitals, offices, and other places of work.

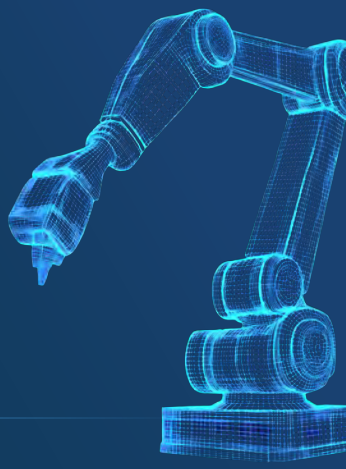

We are already starting to see examples of these hybrid AI workflows today, with many more possibilities to be enabled as the technology matures:

Healthcare

Today

Digital agents analyze data for proper care and assist doctors in making diagnostic decisions.

Tomorrow

Robots could distribute medications and help patients regain mobility using special exoskeletons that execute digitally prescribed rehabilitation routines.

Agriculture

Today

Digital agents analyze weather patterns and monitor crops, making suggestions for proper irrigation and care.

Tomorrow

Robots could perform both irrigation and maintenance based on optimized schedules created by their digital counterparts, resulting in a higher crop yield while improving operational efficiency of agricultural enterprises.

Smart cities

Today

Digital twins of urban centers provide insights into infrastructure, energy consumption, and traffic patterns, and help make data-driven decisions to improve urban living.

Tomorrow

These digital twins could sync with physical AI agents such as traffic drones and waste management robots to autonomously optimize traffic flow, energy usage, and emergency response, making cities more efficient and sustainable.

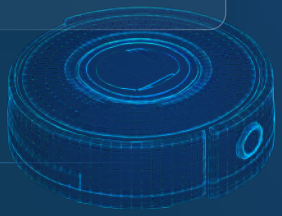

Field service management

Today

Human technicians leverage digital tools and remote diagnostics to perform routine maintenance and address complex field tasks, often operating in hazardous conditions.

Tomorrow

Robot technicians could work alongside human counterparts —or independently—executing precise maintenance and handling dangerous duties, significantly enhancing safety, operational efficiency, and minimizing downtime across industries.

These are just a few industry examples to illustrate the powerful potential of linking digital and physical AI. For businesses, connecting digital and physical AI can turn isolated automation into intelligent, enterprise-wide orchestration. Companies that adopt both digital and physical AI early will lead the next wave of productivity and innovation.

A platform approach unites agentic, digital, and physical AI

To unlock this kind of coordination at scale, enterprises will need more than siloed tools—they'll need a unified platform that orchestrates people, digital agents, and physical AI in harmony.

That's where a platform approach becomes essential. Today, a platform approach acts as the AI control tower for the enterprise, enabling the creation and coordination of digital agents across IT, HR, customer service, and beyond. As intelligence extends beyond screens into the physical world, the same level of orchestration will be needed for robots, drones, and autonomous devices, ensuring they work in sync with human teams and digital agents.

In the future, a single platform could dispatch a field technician or a repair robot based on context, track progress, log the outcome, and trigger the next step—all with the same visibility and control the platform provides today.

> **As robotics becomes more capable and connected, enterprises will need governance that spans both the digital and physical layers.**

As robotics becomes more capable and connected, enterprises will need governance that spans both the digital and physical layers. A platform approach will ensure that every agent, whether software or robotic, contributes to a smarter, safer, and more automated future.

Tomorrow's opportunity: digital to physical workflows

Over the next three to five years, the rise of physical AI stands out as one of the most important technological developments with

material impacts on enterprise operational efficiency and effectiveness and, hence, profitability. Although we are in the early stages of this rapid evolution in robotic capabilities, now is the ideal time to explore opportunities that integrate today's digital AI agents with the physical AI solutions of tomorrow.

Barriers to entry in robotics are declining, thanks in large part to cloud-based simulation platforms that eliminate the need for upfront hardware investment. The potential value of early evaluation and prototyping is substantial, offering significant benefits across the enterprise landscape.

Physical AI is set to take enterprise workflows into the physical world in exciting new ways, and businesses that act today will have a competitive advantage. We might not have robot superheroes or spaceship pilots in the immediate future, but physical AI is poised to accelerate and optimize how every business is run. It might not make for great movies or TV shows, but it will make for great ROI for those who take the leap on physical AI.

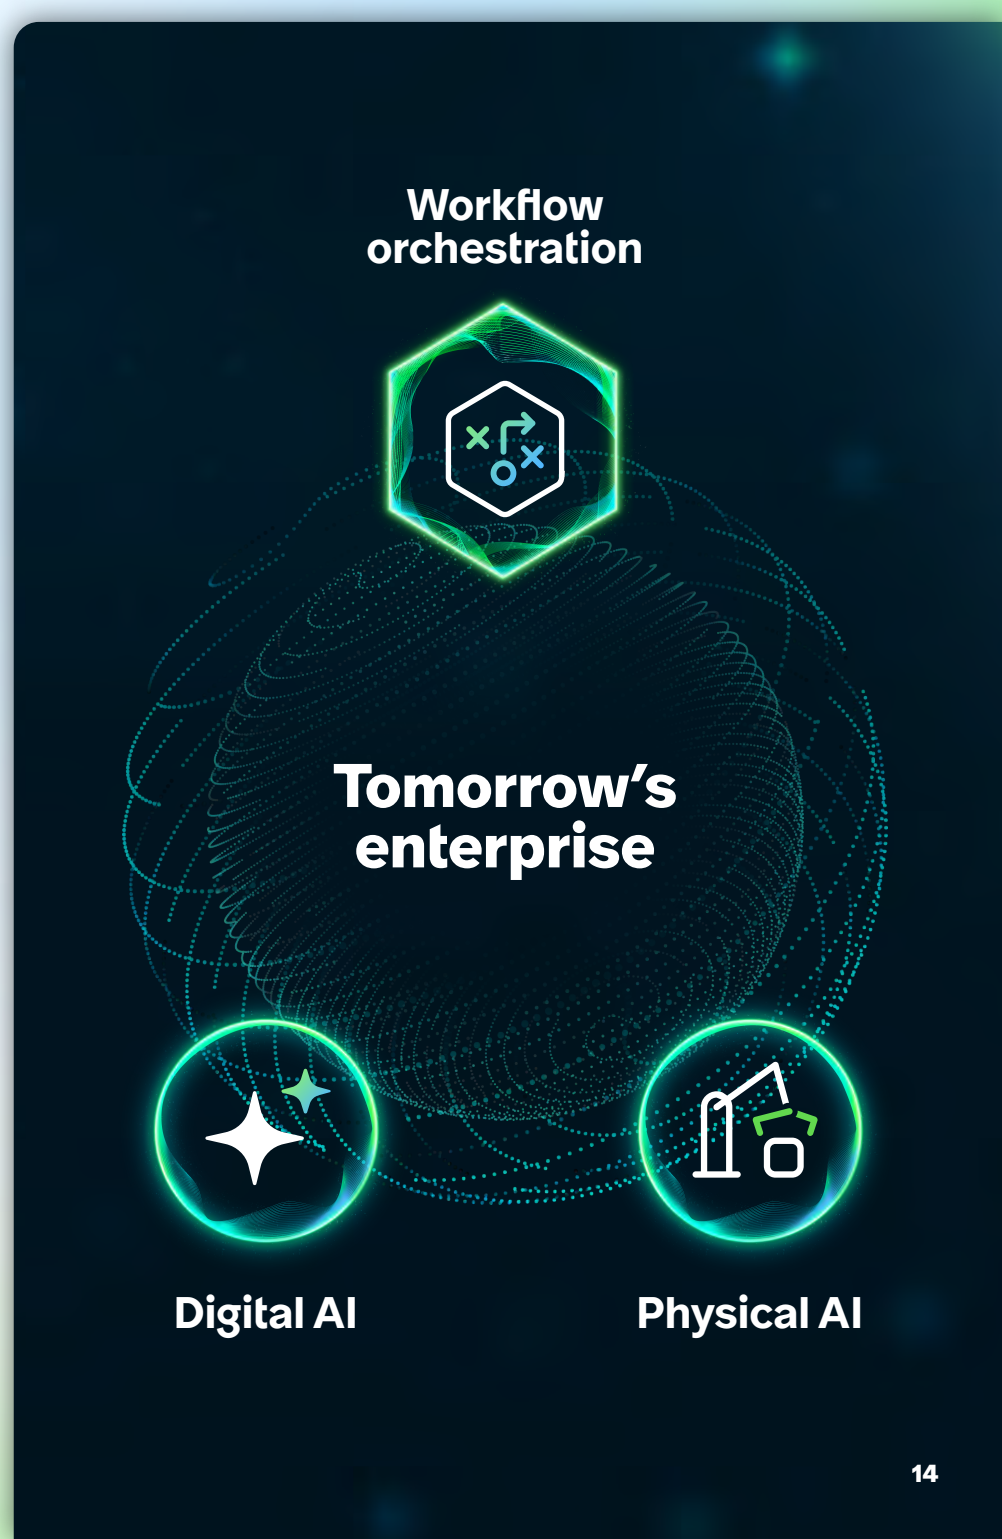

Endnotes

[1] *What is Physical AI?* (n.d.). NVIDIA. https://www.nvidia.com/en-us/glossary/generative-physical-ai/

[2] Spatacco, A. (2025, June 25). *Nvidia CEO Jensen Huang thinks Tesla has a Multitrillion-Dollar artificial intelligence (AI) opportunity up its sleeve (Hint: It's not robotaxi).* Yahoo Finance. https://finance.yahoo.com/news/nvidia-ceo-jensen-huang-thinks-104500210.html

[3] Oitzman, M. (2024, May 31). *OpenAI is restarting its robotics research group.* The Robot Report. https://www.therobotreport.com/openai-is-restarting-its-robotics-research-group/

[4] Mind Intellect News. (2025, June 3). *Marc Andreessen: "The next AI revolution is physical"* [Video]. YouTube. https://www.youtube.com/watch?v=ZQoWSqDqEwk

[5] Glazer, E., Jin, B., & Saeedy, A. (2025, April 9). The $40 billion startup mystery shaking up Silicon Valley. *The Wall Street Journal.* https://www.wsj.com/tech/the-hottest-pre-ipo-stock-an-ai-robotics-startup-with-bold-claims-little-revenue-b0c1f03b

[6] See endnote 5.

[7] Future Market Insights Global and Consulting Pvt. Ltd. (2025, March 23). Industrial Robotics Market to Surge from USD 55.1 Billion in 2025 to USD 291.1 Billion by 2035 Future Market Insights, Inc. *GlobeNewswire.* https://www.globenewswire.com/news-release/2025/03/24/3047512/0/en/Industrial-Robotics-Market-to-Surge-from-USD-55-1-Billion-in-2025-to-USD-291-1-Billion-by-2035-Future-Market-Insights-Inc.html

[8] Vanhoucke, V. (2023, December 5). Speaking robot: Our new AI model translates vision and language into robotic actions. *Google.* https://blog.google/technology/ai/google-deepmind-rt2-robotics-vla-model/

[9] Zhen, H., Qiu, X., Chen, P., Yang, J., Yan, X., Du, Y., Hong, Y., & Gan, C. (2024, March 14). *3D-VLA: A 3D Vision-Language-Action Generative World model.* arXiv.org. https://arxiv.org/abs/2403.09631

[10] *NVIDIA announces Isaac GR00T N1—the world's first open humanoid robot foundation model—and simulation frameworks to speed robot development.* (2025, March 18). NVIDIA Newsroom. https://nvidianews.nvidia.com/news/nvidia-isaac-gr00t-n1-open-humanoid-robot-foundation-model-simulation-frameworks

[11] Stanford Vision and Learning Lab. (n.d.). https://behavior.stanford.edu/

The Innovation Office

The Innovation Office is the voice of future-focused innovation in ServiceNow. Our research, thought leadership, and co-innovation solutions serve as a blueprint for AI-first business transformation. We are a strategic partner with our colleagues and customers to ignite the "spark of the possible."

servicenow The Innovation Office

Publication date: July 2025