

Industrials & Electronics Practice

# The turning point for real-world robotics

Daniela Rus, director of MIT's Computer Science and Artificial Intelligence Laboratory, shares a vision for robots that learn faster, respond in real time, and fit naturally into our lives.



**Robotics is entering** a new phase as breakthroughs in hardware, AI, and data begin to push machines out of controlled settings and into the real world. As robots become more capable and adaptable, the challenge is shifting from what they can do in theory to how reliably they can operate alongside people, especially in tasks that require dexterity and judgment. During a recent visit to the Massachusetts Institute of Technology's Computer Science and Artificial Intelligence Laboratory, McKinsey partner Ani Kelkar spoke with Daniela Rus, the lab's director, about trade-offs between humanoid and specialized robots, the limits of today's AI approaches, and the innovations still needed to make robots truly useful in everyday settings.

## **A step change in robotics capabilities**

**Ani Kelkar:** It's an exciting time to be looking at robotics. What's different today compared with five years ago, and what has stayed the same?

**Daniela Rus:** We are experiencing so much capability and so much possibility with robots because of extraordinary advances in hardware, data, and algorithms. We used to make robots out of metals and heavy plastics. Now we use a wide range of existing materials, or we create materials on demand.

We also have many kinds of miniaturized sensors and much more precise and powerful motors. If you take the hardware, perception systems, and the actuation systems, and you combine all the advances, you end up with hardware that is much more reliable, capable, and miniaturized than it was in past robots.

On the algorithmic side, we have invented many new solutions for robots, including those for navigation, map making, perception, and manipulation. All these advances benefit from the growing volume of data that is now available to our machines.

Together, this is providing the leap forward. What's important to realize is that a robot is a machine with two important aspects: the body and the brain.

The body is important because it determines the extent of a robot's capabilities. If you have a robot on wheels, it won't go up the stairs. The body has to be closely mapped to the desired task. The brain is important because it is a collection of software algorithms that get the body to do what it's meant to do. You need this tight connection and coupling between body and brain. Advances happen only when we have simultaneous advances on the body side and on the brain side.

**Ani Kelkar:** What about specialized robots? Will these advances in hardware and AI algorithms make industrial arms or mobile robots more capable than they are today?

**Daniela Rus:** Absolutely. We have always had this philosophical debate in the field of robotics: Is it better to create a universal machine that can address everyone's needs or optimized solutions

that can be very effective for a fixed set of tasks?

I think both sides of the coin are important. Right now, because [humanoids are so complex](#), I think we're more likely to see specialized robots that will do more and more for us in factories and in everyday life. But we will continue to study humanoids because understanding them gives us interesting insights into life itself and our understanding of intelligence. We need to expand both directions of work.

## The promise—and limits—of humanoid robots

**Ani Kelkar:** Many people have seen humanoid robots in videos—for instance, dancing for Lunar New Year performances in China—and they're getting excited, thinking, "This technology is here, and it's going to be deployed in our factories and hospitals soon." Do you subscribe to that? What challenges do we still need to address before that reality is close?

**Daniela Rus:** Those videos were pretty awesome, but it's not so difficult to get a robot to do something once for a video. Getting a humanoid to do a task in a robust and reliable way, with all the long-tail complications, is much harder.

Humanoid robots have been walking, falling, and flipping for decades. What's happening now is this extraordinary advancement on the body side of the machine. We have robots that are much closer to the human form than ever before.

We have seen interesting progress with respect to the capabilities of these robots, especially navigation. But the fact is that these multi-degree-of-freedom humanoid robots are very complex mechanisms that are very hard to control. Static stability requires one class of algorithms.<sup>1</sup> But if you want dynamic stability—for instance, if you want a robot to reliably pick up a heavy box and do something with it—that's a whole different ballgame.

Humanoid robots are great for human-centered environments—like factory floors and homes—because they essentially have the same dimensions as a human, which lets the machine be a "plug in" in those spaces. We don't have to change the environment to accommodate the machine. But the control side—the brain part of these robots—is not there yet. To get better humanoids, AI has to be better: It has to understand physics and common sense; it must give rapid responses. If the robot has to wait for the cloud to tell it what to do, even for a few seconds, the answer that comes may be inconsistent in the context of a dynamic task. There's a lot of opportunity, but we have a long way to go to get humanoid robots to achieve their promise. We have humanoid robots folding laundry in the lab, for instance, but at a price point that is not affordable for regular consumers.

**Ani Kelkar:** On locomotion, what types of problems are you seeing with robots today?

**Daniela Rus:** It's easy for a robot to do static stability locomotion. It's much harder to be stable

---

<sup>1</sup>Static stability refers to a robot's ability to stay balanced without needing to move. For example, if it's pushed slightly, it can settle back into place instead of tipping over.

in dynamic settings. Imagine the robot playing tennis, or running, or doing something much more vigorous than walking around. These are challenging research problems.

In our lab, we're using data to [teach our robots to learn from humans](#), especially for tasks that are difficult to model from first principles. The philosophy is this: If it's possible to describe, using equations, what you want the machine to do, then that's the best path forward. But with a complex mechanism like a humanoid robot, many tasks are too complicated to be modeled this way, so we use human data to teach the robots how to do those tasks. This is a delicate problem because executing a task involves not just movement but also the forces and torques that come into play when interacting with the world. It's easier to train a robot on tasks that don't require much physical interaction.

**Ani Kelkar:** Is that why there are so many dancing robot videos?

**Daniela Rus:** Yes. Everyone makes dancing robots because that data is much easier to get. If you want the robot to be an important presence in a physical space, and interact with objects in the space, you cannot model those tasks from video alone. If I bump into the fridge while I'm walking, there is a force there that pushes against me that is not modeled in videos.

## The missing piece: AI that understands the physical world

**Ani Kelkar:** You talked about how AI has to improve to enhance the humanoid "brain." Help us calibrate that. People see large language models [LLMs] proliferating in their everyday experiences. How much of the AI in LLMs shares the same architecture or logic with the AI in robotics? Or is it a completely different paradigm?

**Daniela Rus:** Most of the LLMs that have achieved such impressive performance rely on a statistical technique called transformer architecture. It's a technique that allows us to be very effective about predictions, within some context. But these techniques do not have physics baked inside them. They also do not have common sense. Everything that you get from an LLM is statistically driven. For a machine to understand the physical world, you need physics and context.

I'll tell you a funny story. I was at a conference where a humanoid was being demoed, and I asked the humanoid, "Hey, what can you do?" It said, "I help people in the home." Well, the home was a very simple environment: a shelf with a plant and a watering can. I said to the robot, "Can you water that plant?" The robot clunkily moved, and it took a while, but it successfully watered the plant.

Then I said to the robot, "Now, can you water my friend here?" The robot said, "Sure, no problem," took the watering can, and was ready to dump the water on my friend's expensive Italian shoes. That's because the robot does not have the common sense required for the watering task.

The machines that we use in everyday life must understand physics, have a better understanding of our world, and be more responsive and faster to address our needs than what is now possible via the cloud. That means we need AI models that understand physics and that can run “on device”—directly on the robot’s body.

**Ani Kelkar:** How far do you think we are from solving these technical problems—and how hard will it be?

**Daniela Rus:** We already have powerful physics-based AI solutions that are alternatives to transformers that can run on device and power this next generation of robots.<sup>2</sup> We need to make similar progress on the perception side—the robots’ sensors. Then we need to improve and miniaturize the motors.

When I was a student, we did not have laser scanners, so we could not use light for measuring distance. We used sonar, and none of the algorithms worked because sonar is imprecise. When the laser scanner was invented, those same algorithms started to work. We have made huge progress in navigation as a result of this powerful sensor.

But we have not had the same kind of advancement with respect to sensors that are needed for manipulation, so we’re still quite far from powerful machines that can be seamlessly integrated with people in everyday manipulation tasks. We need better bodies—better fingers and generally better hand configurations. We need more compliance in our robot hands, and we need sensors that give us tactile information. We don’t have high-enough resolution “skin” for our robot hands. We also don’t have cameras that are fast enough. We don’t have methods that allow a robot to manipulate an object when the camera is not in the palm and is instead external to the robot.

## The next era of robot training

**Ani Kelkar:** Are we ever going to have something similar to an app ecosystem, where you have robots that learn a specific task and then do all sorts of variations of it—for instance, with picking and placing—or do you think you will need to collect precise data for individual tasks and then train the robot?

**Daniela Rus:** There is a difference between whether you are getting the robot ready to execute a task for a roboticist, or whether you have a preprogrammed robot that you want to interact with as a user. On the skill side, we have to make much more progress.

A lot of people are working on vision–language–action [VLA] models that allow us to connect the perceptual, language, and action spaces in a robot. The language space is very important because people are used to reasoning and communicating in concepts and abstractions. But robots don’t understand words. They understand pixels and numbers. In the past, when we worked with robots, all the algorithms were about XYZ and pixel interactions.

---

<sup>2</sup> Some of these solutions are available through Liquid AI, a company cofounded by Daniela Rus that specializes in AI for processors outside data centers.

People are not so good at reasoning about XYZ coordinates and pixels. We're much better at reasoning about chairs and papers and robot arms. Language has an important role in elevating the level at which humans can interact with machines and get them to reason. We still have a way to go to make the whole thing work.

The good news is that there is a lot of effort and enthusiasm for solving these problems. Eventually, I think we will have robotic companions that we can use for physical tasks in our homes and workplaces. They will adapt to us, rather than the other way around. Right now, we still have to adapt to the robots. In my ideal world in the future, we will have friendly robotic assistants that will understand us and work with us, much like a friend would.

## The next decade in robotics

**Ani Kelkar:** What's the greatest challenge that the robotics community should solve today?

**Daniela Rus:** There isn't one greatest challenge. There are big challenges on both the hardware and the software sides. For example, for hardware, we need to invent skin-like sensors that will enable better manipulation and faster cameras. There is a challenge with getting fast, precise, miniaturized motors. There is a challenge with the perception algorithms that are needed for robots to make sense of the world. There is a challenge with controlling complex mechanisms like humanoids. In order to get humanoids to improve, AI has to be better.

**Ani Kelkar:** I imagine there are safety implications as well?

**Daniela Rus:** The safety of these machines is very, very important, and we need to put in place mechanisms to ensure the machines are safe and can be trusted. AI models that control robots are typically not closed-form solutions. There is always a chance that they will make a mistake. How will the robot respond if the AI brain tells it to do something wrong?

In our lab, we developed a technique we call BarrierNet, where we take an AI model and add a layer to it.<sup>3</sup> This layer understands mathematically what safety means in the context of a task. Through this layer, we force all the outputs of the robot to remain within a safe region, and that ensures that our models are safe. There are other techniques that people are developing. We often talk about trust; one important aspect of trust is to verify that the solution is correct. There are also hardware solutions to increase safety. We can make robots out of soft materials, which are intrinsically much safer to be around than heavy plastics or hard materials.

Our Rainbow Robotics RB-Y1 commercial platform is a humanoid on wheels. It doesn't have legs. So as a mechanism, it's simpler to control than a legged robot. This allows us to abstract out the legged navigation of the robot and really focus on what bimanual manipulation with a humanoid-sized and -shaped robot looks like. We're very excited about focusing on this kind of platform as a way of making advancements on the manipulation side.

---

<sup>3</sup>Wei Xiao, Ramin Hasani, Xiao Li, and Daniela Rus, *BarrierNet: A safety-guaranteed layer for neural networks*, Cornell University, November 22, 2021.

Find more content like this on the  
**McKinsey Insights App**



Scan • Download • Personalize



**Ani Kelkar:** Let's think further ahead, to 2040. Where do you think we will be in terms of state-of-the-art robots? Will we have *The Jetsons* era of companion robots?

**Daniela Rus:** Twenty years ago, computing was a task reserved for the expert few because computers were large and expensive, and you needed to know what to do with them. All that changed with the miniaturization of the computer and the introduction of the smartphone. Today, everyone computes.

I am expecting a similar [kind of transformation](#) with respect to physical work. The apps, LLMs, and computers have democratized cognitive work. I would like to have the same level of democratization for physical work—and that means we should imagine a wide range of machines for our future. Some of the machines will definitely be humanoids. Some will be optimized and targeted for specific tasks. The choice depends on what the machine has to do.

**Daniela Rus** is the director of the Computer Science and Artificial Intelligence Laboratory at the Massachusetts Institute of Technology. [Ani Kelkar](#) is a partner in McKinsey's Boston office.

---

This interview was edited by Eileen Hannigan, an executive editor in McKinsey's Boston office.

Copyright © 2026 McKinsey & Company. All rights reserved.