Random robots are more reliable: New AI algorithm for robots consistently outperforms state-of-the-art systems

May 2 2024

Although the current study tested the AI algorithm only on simulated robots, the researchers have developed NoodleBot for future testing of the algorithm in the real world. Credit: Northwestern University

Northwestern University engineers have developed a new artificial intelligence (AI) algorithm designed specifically for smart robotics. By helping robots rapidly and reliably learn complex skills, the new method could significantly improve the practicality—and safety—of robots for a range of applications, including self-driving cars, delivery drones, household assistants and automation.

Called Maximum Diffusion Reinforcement Learning (MaxDiff RL), the algorithm's success lies in its ability to encourage robots to explore their environments as randomly as possible in order to gain a diverse set of experiences.

This "designed randomness" improves the quality of data that robots collect regarding their own surroundings. And, by using higher-quality data, simulated robots demonstrated faster and more efficient learning, improving their overall reliability and performance.

When tested against other AI platforms, simulated robots using Northwestern's new algorithm consistently outperformed state-of-the-art models. The new algorithm works so well, in fact, that robots learned new tasks and then successfully performed them within a single attempt—getting it right the first time. This starkly contrasts current AI models, which enable slower learning through trial and error.

The research, titled "Maximum diffusion reinforcement learning," is published in the journal *Nature Machine Intelligence*.

"Other AI frameworks can be somewhat unreliable," said Northwestern's Thomas Berrueta, who led the study. "Sometimes they will totally nail a task, but, other times, they will fail completely. With our framework, as long as the robot is capable of solving the task at all, every time you turn on your robot you can expect it to do exactly what it's been asked to do. This makes it easier to interpret robot successes and failures, which is

crucial in a world increasingly dependent on AI."

Berrueta is a Presidential Fellow at Northwestern and a Ph.D. candidate in mechanical engineering at the McCormick School of Engineering. Robotics expert Todd Murphey, a professor of [mechanical engineering](link) at McCormick and Berrueta's adviser, is the paper's senior author. Berrueta and Murphey co-authored the paper with Allison Pinosky, also a Ph.D. candidate in Murphey's lab.

The disembodied disconnect

To train machine-learning algorithms, researchers and developers use large quantities of big data, which humans carefully filter and curate. AI learns from this training data, using trial and error until it reaches optimal results.

While this process works well for disembodied systems, like ChatGPT and Google Gemini (formerly Bard), it does not work for embodied AI systems like robots. Robots, instead, collect data by themselves—without the luxury of human curators.

"Traditional algorithms are not compatible with robotics in two distinct ways," Murphey said.

"First, disembodied systems can take advantage of a world where physical laws do not apply. Second, individual failures have no consequences. For computer science applications, the only thing that matters is that it succeeds most of the time. In robotics, one failure could be catastrophic."

To solve this disconnect, Berrueta, Murphey and Pinosky aimed to develop a novel algorithm that ensures robots will collect high-quality data on-the-go.

At its core, MaxDiff RL commands robots to move more randomly in order to collect thorough, diverse data about their environments. By learning through self-curated random experiences, robots acquire necessary skills to accomplish useful tasks.

Getting it right the first time

To test the new algorithm, the researchers compared it against current, state-of-the-art models. Using [computer simulations](#), the researchers asked simulated robots to perform a series of standard tasks. Across the board, robots using MaxDiff RL learned faster than the other models. They also correctly performed tasks much more consistently and reliably than others.

Perhaps even more impressive: Robots using the MaxDiff RL method often succeeded at correctly performing a task in a single attempt. And that's even when they started with no knowledge.

"Our robots were faster and more agile—capable of effectively generalizing what they learned and applying it to new situations," Berrueta said. "For real-world applications where robots can't afford endless time for trial and error, this is a huge benefit."

Because MaxDiff RL is a general algorithm, it can be used for a variety of applications. The researchers hope it addresses foundational issues holding back the field, ultimately paving the way for reliable decision-making in smart robotics.

"This doesn't have to be used only for robotic vehicles that move around," Pinosky said. "It also could be used for stationary robots—such as a robotic arm in a kitchen that learns how to load the dishwasher. As tasks and physical environments become more complicated, the role of embodiment becomes even more crucial to consider during the [learning](#)

process. This is an important step toward real systems that do more complicated, more interesting tasks."

More information: Maximum diffusion reinforcement learning, *Nature Machine Intelligence* (2024). [DOI: 10.1038/s42256-024-00829-3](#)

Provided by Northwestern University