BraintankAI ASIA

EXPLAINER · VLM ROBOT CONTROL · 5 MIN

BEIJING · Tsinghua University

Off-the-shelf AI models steer robot arms with simple move-and-grip commands

A 2026 preprint gives a frozen vision-language model a small command vocabulary and at most one example. In simulation it rivals robots trained on task data.

19 November 2026 · arXiv preprint, not yet peer reviewed



Photo: Kinsey Wang

Vision-language models have learned a great deal from the digital world. Trained on vast collections of images and text, they can describe a scene, follow instructions and reason through a problem. Robots need many of the same skills: they must understand what they are looking at, work out what a request means and plan a sequence of steps. So it is natural to ask whether the intelligence these models picked up online can be carried over to controlling a physical machine.

The usual answer has been to retrain the model on robot data until it speaks the language of motors. A 2026 paper by Meng-Hao Guo, Shi-Min Hu and colleagues, posted to arXiv as a preprint that has not yet been peer reviewed, tries something different. It leaves the model untouched and changes only how the model is asked to act.

The idea

A vision-language model, or VLM, takes images and text as input and produces text. The dominant way to turn one into a robot controller is to train it further on recordings of robots doing tasks, so that it outputs motor actions directly. The result is called a vision-language-action model, or VLA. A close relative, the world-action model, learns both to predict how a scene will change and what action to take. In both cases the model must be trained on paired robot data.

The authors argue that this data is a bottleneck. It is costly to collect, tied to a particular robot body and hard to scale. Robot data is also orders of magnitude scarcer than the image and text data VLMs learn from. They also cite evidence that training a VLM to predict actions can wear away the general abilities that made it attractive, such as following instructions and reasoning. Earlier agent-style systems avoided retraining by using large models only for high-level planning, and handed the actual movement to libraries of pre-built skills or to code. Most more recent agents still rely on a trained VLA underneath for low-level control.

The authors pose a narrower question. Can a frozen VLM, with no changes to its parameters, control a robot itself, using only a lightweight interface and a few examples placed in its prompt? Teaching a model through prompt examples, without altering its weights, is known as in-context learning.

How this paper puts it to work

The system, called RoboDawn, is a harness: software that feeds the model observations, reads its commands, executes them on the robot and reports back. It works in a closed loop. The model looks, acts, sees what happened and decides again, rather than planning everything up front. In each round the model receives the task instruction, annotated camera images from several viewpoints, the robot's measured state, feedback from the last round and a running memory of the interaction. A fixed profile describes the workspace limits, the cameras, a grid used to help locate things in the images, and the properties of the gripper.

The model replies with a short progress estimate, its current plan and some working notes. It then issues commands from a small, game-like vocabulary. It can move the left or right arm a set distance along an axis, rotate it, point it in a preset direction such as down or forward, open, close or partly open the gripper, return home, wait, or ask for the task to be checked as done. Every movement is measured from the gripper interaction point, the midpoint between the two fingertips. Each command is capped at 20 centimetres of travel or 90 degrees of rotation. Software then works out the joint angles needed to reach each target, a calculation called inverse kinematics, and carries out the motion. The model never receives the exact positions of objects; it has to work from what it sees.

The examples come in two layers. A shared command primer shows what each basic command does and is always present. On top of that, the model can be given zero, one or a few worked task demonstrations: the zero-shot, one-shot and few-shot settings. These demonstrations come from scripted experts in simulation or from people remotely operating real robots, recorded in scenes kept separate from testing. They are translated into the same command vocabulary and annotated with explanations written afterwards by a VLM.

What it shows, and what it does not

On RoboTwin 2.0, a 50-task simulation benchmark for two-armed robots tested in visually randomised scenes, the authors compared RoboDawn with policies trained on 50 demonstrations per task. Each task was run 10 times. Driven by GPT-6 Astra with no task demonstration, RoboDawn reached 53.2 per cent success. That beat trained policies such as pi_0.5 at 46.0 per cent and LingBot-VLA at 50.4 per cent, but fell short of the agentic HarnessVLA at about 58 per cent. With one demonstration, drawn from the same pool the baselines trained on, success rose to 73.6 per cent. On a second benchmark, RoboDojo, with 42 tasks run five times each, the one-shot system reached 47.17 per cent against 19.34 per cent for the best trained baseline. Eight of those tasks are open-ended ones the baselines may lack training data for. Excluding them, the one-shot figure was 43.33 per cent.

The results depend heavily on which model sits inside the harness. With one demonstration, success on RoboTwin ranged from 14.4 per cent with the weakest model tested to 73.6 per cent with GPT-6 Astra. Removing explicit reasoning or the localisation grids cut zero-shot success sharply, from 47.0 per cent to 34.8 and 32.4 per cent respectively. Adding more demonstrations helped up to four, then slipped back at eight. The demonstration and component ablations were run with a single model, Gemini-3.8-Flash. The trial counts are small and the results are reported without error bars. The 'zero-shot' setting also still includes the command primer and robot profile.

Speed is a clear weakness. In the authors' timing test, which used yet another model, each decision took 9.74 seconds, compared with about a tenth of a second for pi_0.5. The model spent more than four times as long thinking as moving, far above the ratio the authors say would let thinking overlap with motion. Physical tests were few: 10 trials each of three tasks, with no robot baselines. A Franka arm put a block in a basket nine times out of 10 and stacked blocks five times out of 10. A Piper arm folded cloth successfully in none of its 10 trials. The authors suspect that failure stems mainly from VLMs handling rotations less reliably than straight-line moves. They also report that the commands are too coarse for fine corrections near contact. Sensible-looking commands sometimes produce motions that collide with nearby objects, and the model can declare a task finished too early. They flag that errors could damage objects or endanger people. The work shows that a top proprietary model can handle a meaningful share of two-armed manipulation in simulation without robot training. It does not yet show fast, precise or robust control in the physical world.

PAPER	*Transferring the Intelligence of VLMs to Robotic Control*
AUTHORS	Meng-Hao Guo, Zhe-Han Mo, Jia-Jun Wang, Yi Zhang, Kejin Wang, Yi-Xuan Deng, Jia-Peng Zhang, Yongming Rao, Shi-Min Hu et al.
VENUE	arXiv preprint — not yet peer reviewed · 19 September 2026
LINK	arxiv.org/abs/2609.22966
INDEX	stacks.braintank.ai