Physical AI — the path to a beautiful tomorrow.



PRAVDA REPLIKANTA

THE REPLICANT TRUTH
A NEWSPAPER ABOUT PHYSICAL ARTIFICIAL INTELLIGENCE

BREUKELEN (Netherlands)
Forecast · 2026-11-02 · CET

MORNING · 09:00	9.8°	Cloudy Wind: N 2.9 m/s · 54%
MIDDAY · 12:00	12.4°	Clear Wind: N 5.3 m/s · 11%
EVENING · 18:00	11.1°	Cloudy Wind: N 4.1 m/s · 21%

Daily: 8.6–12.5 °C
Precipitation probability: 94%

All we can do is learn to be human while the rain falls.

ISSUE № 5 · 2026-11-02 rev. 1 ·

Weather source: Open-Meteo · Retrieved: 2026-11-01 20:51 CET

Pravda Replikanta — For young readers

Robotics research, from training and touch to evaluating skills and practical limitations.

News

Drawing from instructions becomes a test of robot generalization

A robot repeating a familiar motion does not necessarily understand a new request. Xiangyu Xu and collaborators introduced a benchmark where a robot arm receives a verbal description of a drawing and chooses a sequence of strokes. Starting from the same blank canvas, multiple drawings and stroke orders can be valid. This reveals more than reproducing one memorized trajectory.

The researchers generated 28,200 simulated demonstrations spanning 282 drawing tasks and evaluated five models connecting visual input, language and robot actions. They examined the effects of training scale, task diversity, spatial composition and the ability to combine familiar objects in unseen arrangements. The paper reports distinct failure patterns among models and examples of transfer to new compositions.

The team also deployed a trained policy on a physical robot arm and demonstrated drawing that reacts to human intervention. Yet drawing is a deliberately constrained task, not a complete proxy for kitchen or warehouse manipulation. A successful drawing does not demonstrate unrestricted autonomy, safety or reliability outside the tested setup.

A useful distinction is between doing something correctly once and doing it reliably in many situations. Imagine finding your way through a new school: memorizing one landmark may not be enough when the route changes. Robot learning also needs tests with unfamiliar objects and slightly different surroundings. The numbers in a paper describe its particular experiment. Asking where the robot fails, when a person helps and what was changed during testing makes it easier to understand what the invention can really do today. The signal is not artistic beauty but whether familiar elements can form an unseen composition. Line accuracy and language understanding are separate abilities. For reference, the developers named this system BURGER.

How we know

BURGER: Benchmarking and Understanding Robotic Generalization via Embodied Drawing · 2026-10-29

A robot map estimates how doors and drawers respond to force

A conventional room map can show a robot where a door is without revealing which way it moves or how much force opening it requires. Dibyendu Ghosh proposes a map that stores probabilistic beliefs about articulated components. After a measured pull, the system updates its estimate, and a planner can ask whether an action is likely to succeed within a stated risk limit.

Watching a part move does not uniquely reveal its dynamics. Even a measured pull directly informs only the tested component and direction. The method therefore maintains a distribution of possible properties, shares relevant evidence with visually similar untested parts and can abstain when the evidence is inadequate.

In real doors and drawers across 18 homes, the method identified joint types in 86% of held-out pulls versus 48% for a single-view learned comparison. Its prediction error measured by CRPS was 0.323 versus 0.544 for differentiable simulation, while computation took 13 milliseconds versus 230 seconds per case. These are

results for the studied objects, not a universal safety guarantee.

A useful distinction is between doing something correctly once and doing it reliably in many situations. Imagine finding your way through a new school: memorizing one landmark may not be enough when the route changes. Robot learning also needs tests with unfamiliar objects and slightly different surroundings. The numbers in a paper describe its particular experiment. Asking where the robot fails, when a person helps and what was changed during testing makes it easier to understand what the invention can really do today. Force, motion direction and uncertainty should remain measurable. False-safe predictions and their consequences need separate testing. For reference, the developers named this system PhysVLMM.

How we know

PhysVLMM: Interaction-Conditioned Physical Belief Maps for Articulated Objects in Robot Scene Maps · 2026-10-29

Human fingertip contact becomes training data for robot hands

Human fingers and robotic fingers have different geometry, so copying a hand trajectory does not guarantee the same contact with an object. Haozhe Zhang and collaborators propose a combined hardware-and-software system with a wearable human demonstration interface and a corresponding robotic interface. Both capture visual and fingertip tactile information to establish correspondence across different hand designs.

Its the contact-alignment method component aligns fingertip contact patches and their spatial relationships using measured touch. The authors report that this conversion does not require an object mesh, an object-pose estimate or a physics simulator. With visual synthesis and camera-anchored action conversion, the transformed examples are used to fine-tune a pretrained model that maps visual observations and instructions to robot actions.

Across four physical contact-rich tasks, policies trained solely on converted human demonstrations reached 72.5% average success. An equal number of robot demonstrations collected through direct teleoperation yielded 77.5%, while human data collection was seven times faster. These results are promising for collecting examples efficiently, but cover specific tasks, hands and training conditions, not every object or robotic platform.

A useful distinction is between doing something correctly once and doing it reliably in many situations. Imagine finding your way through a new school: memorizing one landmark may not be enough when the route changes. Robot learning also needs tests with unfamiliar objects and slightly different surroundings. The numbers in a paper describe its particular experiment. Asking where the robot fails, when a person helps and what was changed during testing makes it easier to understand what the invention can really do today. Sevenfold faster collection does not remove the need to test contact on target objects. The 72.5% and 77.5% results matter alongside failure costs. For reference, the developers named this system TouchBridge.

How we know

TouchBridge: Tactile-Conditioned Human-to-Robot Transfer for Dexterous Manipulation · 2026-10-29

Millions of labeled short actions expose the structure of long robot tasks

An instruction such as clearing a table hides many steps: approach, grasp, move and release. Yet large robot datasets often store only the overall task instruction. Xueyuan Wei and collaborators propose a dataset that assigns consistent time-aligned labels to short, or atomic, actions so learning systems can recover the structure of longer behavior.

The authors describe annotations spanning 1.45 million trajectories, 6.88 million skill-labeled segments and 16,882.1 hours. Their pipeline combines human annotation, propagation of labels across repeated skill sequences and multiple quality-control stages. They also introduce a dedicated benchmark to test both skill recognition and the localization of skill boundaries in time.

According to the paper, an annotation model trained on these labels outperformed the evaluated open and closed models on both a dedicated benchmark and an external WGO-Bench. In control experiments, atomic-action supervision combined with hierarchical execution improved task success on LIBERO-Long, LIBERO-Plus and RoboTwin. The abstract does not specify every numerical gain, and the promised release of data and tools is not equivalent to an already verified public download.

A useful distinction is between doing something correctly once and doing it reliably in many situations. Imagine finding your way through a new school: memorizing one landmark may not be enough when the route changes. Robot learning also needs tests with unfamiliar objects and slightly different surroundings. The numbers in a paper describe its particular experiment. Asking where the robot fails, when a person helps and what was changed during testing makes it easier to understand what the invention can really do today. Consistent time-aligned labels help combine datasets, but incorrect skill boundaries can distort

later learning. For reference, the developers named this system RoboAtom-Data.

How we know

RoboAtom: Large-scale Atomic-Action Supervision for Robot Learning · 2026-10-29

One learning pipeline adapts coordinated manipulation to different robot hands

Robotic hands can have very different numbers of joints yet face similar problems: hold an object, turn it and coordinate both wrists and fingers. Xiaoyu Yang and Nan Wu propose a pipeline combining human-motion pretraining, simulation-based expert learning, robot demonstration collection and visual feedback during execution.

During pretraining, the system predicts one second of bimanual motion from five images and thirty annotated hand states recorded over the preceding second. Shared wrist representations are retained when adapting to robots, while finger control accounts for the mechanics of each hand. The researchers tested six-channel the first hand hands and sixteen-joint the second hand hands, with control predictions revised from new camera observations.

In their evaluations, direct regression achieved 40.6% success on opening and returning a lid, 100% on the first hand object reorientation and 93.3% on the second hand reorientation. Flow matching achieved 90% on the second hand reorientation. These figures describe different tasks and conditions, not an overall ranking of robotic hands. Constraining new plans with previous plans can improve continuity.

A useful distinction is between doing something correctly once and doing it reliably in many situations. Imagine finding your way through a new school: memorizing one landmark may not be enough when the route changes. Robot learning also needs tests with unfamiliar objects and slightly different surroundings. The numbers in a paper describe its particular experiment. Asking where the robot fails, when a person helps and what was changed during testing makes it easier to understand what the invention can really do today. Shared wrist representations and hand-specific finger control support transfer studies. Results still depend on the object and control method. For reference, the developers named this system AcrossDEX 1.0.

How we know

AcrossDEX 1.0: From Human Motion Pretraining to Bimanual Dexterous Manipulation across Hand Embodiments · 2026-10-29

A closer look

Testable skills matter more than striking demonstrations

New research emphasizes structured data, measurements and testable results.

The next step is independent testing on different robots and real tasks.

How we know

BURGER: Benchmarking and Understanding Robotic Generalization via Embodied Drawing · 2026-10-29

PhysVLMM: Interaction-Conditioned Physical Belief Maps for Articulated Objects in Robot Scene Maps · 2026-10-29

TouchBridge: Tactile-Conditioned Human-to-Robot Transfer for Dexterous Manipulation · 2026-10-29

RoboAtom: Large-scale Atomic-Action Supervision for Robot Learning · 2026-10-29

AcrossDEX 1.0: From Human Motion Pretraining to Bimanual Dexterous Manipulation across Hand Embodiments · 2026-10-29

Words to know

What the terms in these articles mean.

01 **training example**

A recorded action used to learn a similar task.

02 **test task**

A defined challenge used to check a learned skill.

03 **touch sensor**

A device measuring contact between a finger and an object.

04 **action sequence**

Several short steps that make up one task.

05 **probability**

A way to express how likely an outcome is thought to be.

How we know

BURGER: Benchmarking and Understanding Robotic Generalization via Embodied Drawing · 2026-10-29

PhysVLMM: Interaction-Conditioned Physical Belief Maps for Articulated Objects in Robot Scene Maps · 2026-10-29

TouchBridge: Tactile-Conditioned Human-to-Robot Transfer for Dexterous Manipulation · 2026-10-29

RoboAtom: Large-scale Atomic-Action Supervision for Robot Learning · 2026-10-29

AcrossDEX 1.0: From Human Motion Pretraining to Bimanual Dexterous Manipulation across Hand Embodiments · 2026-10-29