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 — Engineering edition

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

News

Coding agents build robot training worlds while frontier models struggle on new tests

Constructing a realistic virtual training environment can require extensive expert labor before a robot starts learning. Pranjal Aggarwal and Sean Welleck propose Gym-Anything, where coding agents build simulations and an adversarial audit checks whether their physics and tasks satisfy explicit criteria. The selection of physical work draws on US economic activity instead of focusing solely on familiar toy demonstrations.

RoboWorld contains 25 simulations, 175 training environments and 25 held-out test environments. According to the authors, even the strongest evaluated frontier models failed every test task, averaging 0.06 reward. Reward is a numerical simulator score, not a real-world robot success percentage. The paper also explores learning from repeated attempts, privileged teachers that can inspect simulator state, and demonstrations produced by code policies.

Several approaches improved the measured results, but transfer to physical hardware remains unproven. Automatically generated worlds can misrepresent friction, contact dynamics and unexpected events in actual industrial settings.

The engineering value lies in reproducibility rather than a headline number alone. A replication needs the hardware configuration, training data, evaluation protocol, control frequency and stopping conditions. Comparisons should use consistent compute budgets and attempt counts. Rare failures deserve separate analysis rather than disappearing into an average. Until an independent team reproduces the result on another setup, an improvement is evidence within the authors' experimental conditions, not a universal property of the method. Contact, friction and task conditions should be checked against real observations before relying on simulated scores.

Sources

Gym-Anything for Robotics: Scaling physical simulations with digital agents · 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 PhysVLMM, 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.

The engineering value lies in reproducibility rather than a headline number alone. A replication needs the hardware configuration, training data, evaluation protocol, control frequency and stopping conditions. Comparisons should use consistent compute budgets and attempt counts. Rare failures deserve separate analysis rather than disappearing into an average. Until an independent team reproduces the result on another setup, an improvement is evidence within the authors' experimental conditions, not a universal property of the method. Force, motion direction and uncertainty should remain measurable. False-safe predictions and their consequences need separate testing.

Sources

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 TouchBridge, 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 Tac-Retarget 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.

The engineering value lies in reproducibility rather than a headline number alone. A replication needs the hardware configuration, training data, evaluation protocol, control frequency and stopping conditions. Comparisons should use consistent compute budgets and attempt counts. Rare failures deserve separate analysis rather than disappearing into an average. Until an independent team reproduces the result on another setup, an improvement is evidence within the authors' experimental conditions, not a universal property of the method. 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.

Sources

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 RoboAtom-Data, which 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 RoboAtom-Bench 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 RoboAtom-Bench 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.

The engineering value lies in reproducibility rather than a headline number alone. A replication needs the hardware configuration, training data, evaluation protocol, control frequency and stopping conditions. Comparisons should use consistent compute budgets and attempt counts. Rare failures deserve separate analysis rather than disappearing into an average. Until an independent team reproduces the result on another setup, an improvement is evidence within the authors' experimental conditions, not a universal property of the method. Consistent time-aligned labels help combine datasets, but incorrect skill boundaries can distort later learning.

Sources

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 AcrossDEX 1.0, 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 Inspire hands and sixteen-joint Allegro 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 Inspire object reorientation and 93.3% on Allegro reorientation. Flow matching achieved 90% on Allegro reorientation. These figures describe different tasks and conditions, not an overall ranking of robotic hands. Constraining new plans with previous plans can improve continuity.

The engineering value lies in reproducibility rather than a headline number alone. A replication needs the hardware configuration, training data, evaluation protocol, control frequency and stopping conditions. Comparisons should use consistent compute budgets and attempt counts. Rare failures deserve separate analysis rather than disappearing into an average. Until an independent team reproduces the result on another setup, an improvement is evidence within the authors' experimental conditions, not a universal property of the method. Shared wrist representations and hand-specific finger control support transfer studies. Results still depend on the object and control method.

Sources

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

Analysis

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.

Sources

Gym-Anything for Robotics: Scaling physical simulations with digital agents · 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

Issue glossary

Key terms in this issue.

01 **simulation**

A computational model reproducing physical processes for experiments.

02 **tactile signal**

A contact measurement used when transferring hand movements.

03 **probabilistic model**

A representation retaining a distribution of possible values.

04 **time-aligned annotation**

A skill label linked to an interval in a motion record.

05 **feedback**

Adjustment of control commands using newly observed state.

Sources

Gym-Anything for Robotics: Scaling physical simulations with digital agents · 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