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 — Business edition

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

News

Japan's Forcesteed proposes robot programs from instructions, with human approval

On October 29, Japan-based Forcesteed Robotics announced a feature for its developing FOS platform. A person describes a desired operation in ordinary language, and the local Luqua software agent selects an approach for each step: predefined rules, taught movements or a model linking vision, language and action. It then proposes a node-based work program for the company's LEIVOR controller.

The company intends to reuse previously created motions, skills and complete tasks. A key boundary is the approval workflow: the agent is not supposed to execute its proposed sequence independently. Only programs reviewed and authorized by a human are run. Local processing is also intended for sites where internet access is restricted or unavailable.

This is a vendor announcement, not an independent performance evaluation. Forcesteed explicitly describes further internal and partner demonstrations as future work, and it provides no comparative figures for integration time, error rates or production throughput. The practical idea is not unlimited autonomy; it is combining several control techniques into a reviewable work program with human authorization.

For a customer, the central question is how the approach fits an existing workflow without hidden costs. Integration includes sensors, control interfaces, staff training, logs and a procedure for human intervention. A pilot should define acceptable risk, cost per successful operation and the consequences of downtime in advance. A laboratory score alone cannot support a credible return-on-investment calculation. A supplier should demonstrate the system on the customer's own objects and data and explain how it will be maintained after deployment. The key boundary is between proposing a program and authorizing physical execution. Decision logs and command checks are needed.

Sources

Forcesteed Robotics、AIエージェントがロボットの作業プログラムを組む「Agentic Robotics」を開発 ― Physical AI OS「FOS」の新機能 · 2026-10-29

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.

For a customer, the central question is how the approach fits an existing workflow without hidden costs. Integration includes sensors, control interfaces, staff training, logs and a procedure for human intervention. A pilot should define acceptable risk, cost per successful operation and the consequences of downtime in advance. A laboratory score alone cannot support a credible return-on-investment calculation. A supplier should demonstrate the system on the customer's own objects and data and explain how it will be maintained after deployment. 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

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.

For a customer, the central question is how the approach fits an existing workflow without hidden costs. Integration includes sensors, control interfaces, staff training, logs and a procedure for human intervention. A pilot should define acceptable risk, cost per successful operation and the consequences of downtime in advance. A laboratory score alone cannot support a credible return-on-investment calculation. A supplier should demonstrate the system on the customer's own objects and data and explain how it will be maintained after deployment. 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.

For a customer, the central question is how the approach fits an existing workflow without hidden costs. Integration includes sensors, control interfaces, staff training, logs and a procedure for human intervention. A pilot should define acceptable risk, cost per successful operation and the consequences of downtime in advance. A laboratory score alone cannot support a credible return-on-investment calculation. A supplier should demonstrate the system on the customer's own objects and data and explain how it will be maintained after deployment. 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

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.

For a customer, the central question is how the approach fits an existing workflow without hidden costs. Integration includes sensors, control interfaces, staff training, logs and a procedure for human intervention. A pilot should define acceptable risk, cost per successful operation and the consequences of downtime in advance. A laboratory score alone cannot support a credible return-on-investment calculation. A supplier should demonstrate the system on the customer's own objects and data and explain how it will be maintained after deployment. 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

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

Forcesteed Robotics、AIエージェントがロボットの作業プログラムを組む「Agentic Robotics」を開発 ― Physical AI OS「FOS」の新機能 · 2026-10-29

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

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

Issue glossary

Key terms in this issue.

01 **integration**

Connecting a new technology to existing workflows and equipment.

02 **pilot evaluation**

A limited trial using customer tasks before scaling deployment.

03 **risk control**

Defined conditions for allowing actions and stopping operation.

04 **annotation cost**

The labor and expense of labeling training data.

05 **local execution**

Running software on site without mandatory cloud access.

Sources

Forcesteed Robotics、AIエージェントがロボットの作業プログラムを組む「Agentic Robotics」を開発 ― Physical AI OS「FOS」の新機能 · 2026-10-29

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

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