Physical AI — the path to a beautiful tomorrow.



PRAVDA REPLIKANTA

THE REPLICANT TRUTH
A NEWSPAPER ABOUT PHYSICAL ARTIFICIAL INTELLIGENCE

BREUKELEN (Netherlands)
Forecast · 2026-11-03 · CET

MORNING · 09:00	9°	Cloudy Wind: NE 4.4 m/s · 11%
MIDDAY · 12:00	10.4°	Clear Wind: NE 4.3 m/s · 6%
EVENING · 18:00	8.1°	Clear Wind: NE 4.1 m/s · 0%

Daily: 5.2–11.1 °C
Precipitation probability: 21%

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

ISSUE № 6 · 2026-11-03 rev. 1 ·

Weather source: Open-Meteo · Retrieved: 2026-11-03 05:25 CET

Robots in practice: testing, safety and integration

Five studies examine deployment limits, from testing grippers to controlling permitted software commands.

News

Why 1,700 physical robot-hand trials matter more than attractive training scores

Putting a small network connector into its socket is harder than it looks. The plug must face the right way, line up and enter without slipping. A team led by Ammar Abbas compared ways to teach a five-fingered robot hand this task. They carried out roughly 1,700 attempts on real equipment, not just on a screen.

The team tested 281 versions of control software from seven groups. They changed what information the program received from cameras and how it described finger movements. This is called a data representation: the same task can be described using different pictures and numbers. Sometimes that choice mattered more than ordinary adjustments to the training process.

The researchers also examined 67 scores that could be calculated without performing the insertion. Only five, related to how smooth and direct the predicted motion was, closely matched real success. Even those worked only within one group of programs. A movement that looks excellent in a computer test may still miss the actual socket.

For point-cloud policies, cropping the relevant region, using robot-base coordinates and normalizing features improved success by 70 to 100 percentage points in individual controlled comparisons. For a vision-language-action policy, absolute rather than relative actions improved insertion success by up to 70 points. These figures come from different configurations and cannot be added together or generalized to every architecture. Seven policy families were included, exposing how representation choices interact with model design.

The findings concern particular hands, connectors and test conditions. They do not establish the best program for every object. They do explain why a robot must be tested with the things it will really handle.

For a customer automating assembly, this points to acceptance tests on actual components. Software may score well internally yet still miss a socket. A procurement specification should define connector geometry, tolerances, trial counts, cycle time and lighting. The paper offers no universal deployment price or guarantee for other connectors, and its most striking improvements arise from particular representation choices rather than a new robot purchase.

Only five of 67 offline measures reliably correlated with real insertion success, and only within one policy family. Correlations between minus 0.70 and minus 0.80 for predicted motion smoothness and directness do not turn these scores into substitutes for hardware trials. Selecting a controller from validation loss alone is particularly weakly supported here.

Only five of 67 offline measures predicted success, and only within one policy family. Replacing physical acceptance testing with a single model score is therefore not supported by these findings.

Sources

From Teleoperation to Long-Horizon Dexterity: Evaluating Visuomotor Skills and Offline Proxies for Dexterous Manipulation · 2026-10-29

CaPtain lets a small model control robots through software tools

Robot workcells normally execute commands through software interfaces that can restrict unsafe movement. Yijia Fan and colleagues introduced CaPtain, in which a small vision-language model chooses calls to such tools. A software harness surrounds the model with perception, geometry and guarded-motion functions. This differs from asking a neural network to output every joint movement directly, and it gives integrators a distinct place to define permitted operations.

The model was first trained on examples of tool-call sequences and then improved using multi-turn reinforcement learning. During training, a special evaluator inspected hidden simulator state to assign credit to individual commands. The deployed agent does not receive that privileged information. On held-out LIBERO tasks, the untrained model scored 0%, rising to 55.4% after the first stage and 73.8% after the second.

For a potential customer, the attraction is reducing manual programming of every action sequence. Integration still needs a strict list of allowed tools, action logs and tests for mistaken perception. The authors also report 78.5% on a different LIBERO-PRO benchmark and 83.3% on SimplerEnv WidowX without target-domain training. These are different evaluations, not guaranteed production throughput.

The paper reports a transfer to real hardware and the release of code, data and model weights. It does not independently establish industrial safety compliance or total deployment cost. A pilot must examine not only task success but also which commands the system can issue when its observations are wrong.

Sources

CaPtain: A Pipeline for Post-Training Code-as-Policy Robot Agents · 2026-10-29

Better person tracking can create a new collision risk

Imagine a robot following one person down a busy corridor. Someone wearing similar clothes crosses in front of it, and the robot may confuse the two people. A team led by Mingyang Lyu studied how improving a robot's ability to keep following someone affects the safety of its movements.

Training the software to keep track of its person made it better at following. But in some tests it also increased collisions. The researchers therefore added a separate danger check. It looks at clues about the target person's location and how confidently the robot can tell that person apart from others.

When the risk is high, the check shortens the next forward movement. It does not change the robot's turning command. In the experiments, more than 98% of steps remained unchanged. In a difficult tracking test, the collision rate fell by 3.96 percentage points compared with the version without this added check.

The authors fine-tuned only the planner head of a compact vision-language-action model while freezing the backbone. A separate risk predictor processes 84 cues, including candidate targets, identity margins, estimated range and bearing, and the proposed command. Above a calibrated risk threshold, it shortens the next translational waypoint without changing the turning component. This targets dangerous moments instead of applying a blanket speed reduction or replacing the learned controller with a wholly different planner.

That does not make the robot safe in every crowded shopping centre. The researchers compared programs under particular test conditions and did not claim protection from every possible mistake. Interestingly, simply slowing the robot all the time did not produce the same improvement.

In a service setting, the risk check should be assessed separately from the main tracking algorithm. It estimates signs that the robot may be following the wrong person and shortens forward movement when needed. That preserves ordinary behaviour most of the time, but it does not replace speed limits, emergency stops or building-specific trials. The paper does not independently establish compliance with industrial safety standards.

A concurrent evaluation found a 3.77-point collision reduction versus 0.29 points in matched null pairs. The approach is not a formal safety guarantee beyond the studied tracking scenarios. Constant slowing without target-state information did not reproduce the benefit, distinguishing the proposed filter from a simple speed governor.

Uniformly slowing the robot did not achieve the same benefit, and a repeat evaluation found a 3.77-point reduction. Neither figure is a promise of zero risk to visitors.

Sources

TRACE: Target-State Risk-Aligned Conditional Execution for Embodied Tracking · 2026-10-30

Robot gripping research shifts attention to where fingers touch

When you pick up a glass, bending your fingers is only part of the job. They must touch the glass in useful places or it may slip. A team led by Zhixuan Xu studied how to teach robot hands using the changing places where fingers contact an object. Their system includes pictures from cameras near the fingertips and information about contact surfaces.

One part records where human fingers touched an object. Another helps a robot choose contact points and wrist movements of its own. This matters because a human hand and a robot hand can be shaped quite differently. Simply copying the angles of human finger joints may not let a robot hold the same thing.

The researchers compared their approach with other methods on six combinations of robot hands and tasks. They report that average success in the final stage improved by at least 50 percentage points. A percentage point measures the difference between two success rates. For example, moving from 20% to 70% is an increase of 50 points.

The architecture separates demonstration capture from policy learning. Its demonstration interface combines shared contact geometry with colour fingertip views, aligning contact before retargeting the pose. The learning interface predicts relative contact-surface trajectories and then coordinates wrist and finger commands through a lower-level controller. The team has published software and hardware designs, allowing others to attempt experiments with different robotic hands. This is a contact-based representation rather than a claim that human and robot joint angles are interchangeable.

The tests used particular hands and tasks. They do not prove that a robot can now handle every cup or tool. But the experiment helps explain why similar-looking finger movements can produce very different results when the fingers touch an object differently.

A deployment would need to consider fingertip sensing, camera calibration and the surface of the customer's product, not just the shape of the robot hand. A contact surface is the region where a finger actually supports an object; changing it can determine whether a grip holds. Published hardware designs and code may reduce the cost of an early pilot, but integration into a particular production line remains separate work.

The at-least-50-point gain concerns the final stage of six hand-task combinations, not every complete manipulation sequence. Different materials, calibration errors and initial approach motions require further tests. The quantity preserved is the pattern of physical contact, not the visual resemblance of finger poses.

The 50-point result concerns the last stage of selected tasks. It says nothing directly about line uptime, gripper wear, complete cycle speed or the cost of dropping a product.

Sources

Contact Imitation Interface: Transferring Dexterous Skills by Showing How Contact Unfolds · 2026-10-29

Shared navigation may reduce adaptation work across mixed robot fleets

A small wheeled robot and a larger robot with a different steering system cannot move through a narrow corridor in exactly the same way. The command to turn right produces different motions. A team led by Zhidong Wang proposed teaching one program to plan paths for several kinds of robots while still respecting those differences.

The method uses a world model: software that predicts how the scene around a robot may change after it moves. But pictures alone are not enough. The researchers also describe how each particular robot accelerates and turns. The planner can therefore consider what the machine can actually do, not merely which space looks empty.

In computer experiments, the researchers used seven robots representing five kinds of movement. The shared program succeeded 3.8 to 10.6 percentage points more often than some comparison methods. They then drove a real robot manually for three minutes to learn how it moved, before testing navigation in an unfamiliar indoor space.

The method couples a shared world model with a differentiable reduced-order model of each robot's motion through the change in camera viewpoint. The world predictor learns from multiple platforms, while kinematic parameters are identified separately. During imagined rollouts, a proposed command first passes through the robot template and then the visual world predictor. This makes the planner sensitive to turning and speed limits without training an entirely separate policy for each platform.

That real-world test matters, but it does not show that the same program works for every vehicle or on staircases. The study specifically discusses movement on flat ground and particular robot designs.

For operators of mixed fleets, the attraction is avoiding full retraining for every machine design. Each robot still needs

its motion parameters identified. The researchers used three minutes of manual driving on a new physical robot for that step. This is an experimental parameter-identification time, not a promise that installation, mapping, networking, charging and safety setup take three minutes altogether.

In the physical experiment, the researchers identified a new robot's parameters from three minutes of teleoperated driving and then tested unfamiliar indoor navigation. The flat-ground assumption does not cover stairs, steep slopes or contact-heavy manoeuvres. The numerical baseline gains were measured in simulation; the physical test supports a narrower transfer claim.

The approach is primarily designed for flat ground. Fleets using ramps, lifts or crowded pedestrian crossings must test those cases independently of the published result.

Sources

Cross-Embodiment Motion Planning by Coupling Body-Agnostic World Model with Differentiable Ego-Dynamics · 2026-10-30

Analysis

Physical acceptance tests remain essential

High benchmark success cannot replace testing components, motion limits and perception failures.

Deployments must separately assess integration cost, test protocols and consequences of failed actions.

Sources

From Teleoperation to Long-Horizon Dexterity: Evaluating Visuomotor Skills and Offline Proxies for Dexterous Manipulation · 2026-10-29

Contact Imitation Interface: Transferring Dexterous Skills by Showing How Contact Unfolds · 2026-10-29

TRACE: Target-State Risk-Aligned Conditional Execution for Embodied Tracking · 2026-10-30

Cross-Embodiment Motion Planning by Coupling Body-Agnostic World Model with Differentiable Ego-Dynamics · 2026-10-30

CaPtain: A Pipeline for Post-Training Code-as-Policy Robot Agents · 2026-10-29

Issue glossary

Key terms in this issue.

01 **CaPtain**

An agent that calls software tools to control a robot.

02 **LIBERO and LIBERO-PRO**

Two robot-manipulation benchmark suites used to evaluate task success.

03 **software harness**

Tools, memory and constraints around a trained model.

04 **reinforcement learning**

Improving decisions through feedback on action outcomes.

05 **data representation**

The format of observations and commands that affects success.

06 **RJ45 connector**

A network connector used in the physical insertion trials.

07 **contact surface**

The part of an object where a finger applies force.

08 **Contact Imitation Interface**

A skill-transfer method based on contact geometry.

09 **TRACE**

A danger filter for person-following robot movement.

10 **Hourglass**

Shared navigation with a separate motion model for each robot.

11 **world model**

Software predicting how the scene changes after an action.

12 **SimplerEnv WidowX**

A separate robot-arm evaluation environment used to assess transfer without target-domain training.

Sources

From Teleoperation to Long-Horizon Dexterity: Evaluating Visuomotor Skills and Offline Proxies for Dexterous Manipulation · 2026-10-29

Contact Imitation Interface: Transferring Dexterous Skills by Showing How Contact Unfolds · 2026-10-29

TRACE: Target-State Risk-Aligned Conditional Execution for Embodied Tracking · 2026-10-30

Cross-Embodiment Motion Planning by Coupling Body-Agnostic World Model with Differentiable Ego-Dynamics · 2026-10-30

CaPtain: A Pipeline for Post-Training Code-as-Policy Robot Agents · 2026-10-29