FIXTHEROBOT

Weekly Robotics Digest

Week ending 25 Oct 2026

Self-contained briefings for beginners, engineers and researchers. Numbered references point to each story's sources; they are there for verification, not required for understanding.

Editorial introduction

Covered in depth this week: MARE (research result, simulation) and Eka (funding, robotics industry). 2 shorter items follow.

At a glance

1 **MARE: An Open-Source Lunar Rover Simulator with Shadow Ground Truth for Perception and Autonomy** TECHNICAL BRIEFING

The authors released MARE, an open-source lunar rover simulation built on Unreal Engine 5 and ROS 2.

2 **Eka is worth $1.45B thanks to a physical AI model that has experts raving** TECHNICAL BRIEFING

Eka Robotics announced a $150 million Series A round that values the company at $1.45 billion post-money, led by a16z and F-Prime.

3 **PanoGraphGS: Unifying Panoptic Segmentation and Hierarchical 3D Scene Graphs with Gaussian Splatting** SHORT NEWS

The authors introduce PanoGraphGS, a framework that learns a shared 3D Gaussian representation of a scene and simultaneously produces panoptic segmentation and a hierarchical 3D scene graph.

4 **Generative Bionics to launch Physical AI Data Factory in Italy** SHORT NEWS

The first Data Factory hub is set to become one of Europe's largest physical AI training facilities for humanoid robots.

The week in perspective

Editorial interpretation, built only from the stories in this issue.

THEMES

- **Simulation as training ground.** Both the lunar rover simulator and the robotic arm's AI model depend on high-fidelity virtual environments to generate data for perception and control. This suggests a broader trend toward reducing real-world data collection by leveraging detailed physics-based simulations. MARE: An Open-Source Lunar Rover Simulator with...; Eka is worth $1.45B thanks to a physical AI model...

READ WITH CAUTION

- The shadow prediction performance in the rover simulator is reported without independent validation, so its real-world applicability remains uncertain. MARE: An Open-Source Lunar Rover Simulator with...
- Eka's claimed pick rate and durability are based on company-provided figures without disclosed test conditions or third-party verification. Eka is worth $1.45B thanks to a physical AI model...

WHAT TO WATCH NEXT

- Practitioners should monitor any real-hardware tests that compare the simulator's shadow ground truth to actual lunar terrain data. MARE: An Open-Source Lunar Rover Simulator with...
- The rollout of the Physical AI Data Factory may set new standards for large-scale humanoid robot training, warranting attention to its impact on dataset availability. Generative Bionics to launch Physical AI Data...
- Developments in unified Gaussian scene representations could influence future perception pipelines across domains. PanoGraphGS: Unifying Panoptic Segmentation and...

Technical briefings

`TECHNICAL BRIEFING · SIMULATION · RESEARCH RESULT · 20 OCT 2026 · PREPRINT (ARXIV 2610.28465): NO PUBLISHED VERSION FOUND · ARXIV`

MARE: An Open-Source Lunar Rover Simulator with Shadow Ground Truth for Perception and Autonomy

Evidence status: Reported by the paper's authors (preprint): no peer review or independent replication

The authors released MARE, an open-source lunar rover simulation built on Unreal Engine 5 and ROS 2. It provides terrain generation, physics-based rover models with multiple sensors, and a web-based guidance, navigation and control stack. The paper reports a mean Intersection over Union of 0.79 for shadow prediction at low sun angles and shows that stereo depth coverage drops in shadowed regions. [1]

THE PROBLEM

Lunar poles experience very low sun elevations, creating long, high-contrast shadows that degrade stereo vision, terrain segmentation and localization. Developing robust perception and navigation for these conditions needs large, accurately labelled sensor datasets, which are expensive and scarce from field campaigns. [1]

HOW IT WORKS

MARE comprises three parts: Terrain Studio creates cratered and rocky Digital Elevation Maps (DEMs) for scene building; an Unreal Engine project runs two physics-based rover models equipped with stereo, lidar, event and Single-Photon Avalanche Diode (SPAD) cameras, outputting time-stamped sensor streams and pixel-perfect ground-truth shadows; and a ROS 2-based GNC stack offers a web interface for mission control. The simulator renders realistic illumination, allowing shadow-ground-truth comparison for perception research. [1]

Generate lunar DEMs in Terrain Studio → Import DEMs into Unreal Engine scene → Run physics-based rover with sensor plugins → Stream sensor data and pixel-accurate shadow ground truth → Consume data in ROS 2 for algorithm training/evaluation

EVIDENCE

PAPER RESULT Stereo depth analysis shows reduced pixel coverage in shadowed regions compared to sunlit terrain. [1]

LIMITATIONS AND OPEN QUESTIONS

- The sources describe simulation; they report no validation on real hardware.
- The sources report no independent test or third-party validation of the performance claims.
- Only the paper's abstract and metadata were read, not its tables or experimental setup, so results are reported as the authors state them and their conditions are not checked here.

WHO SHOULD CARE, AND WHY

WHY AN ENGINEER SHOULD CARE

Robotics engineers can plug MARE into their ROS 2 pipelines to generate labeled sensor streams for training perception models under lunar lighting; they should verify sensor noise models and consider adding real-world validation before deployment. [1]

RESEARCHER / PROFESSIONAL

The paper presents a novel simulation pipeline with quantified shadow rendering accuracy, but the evidence is limited to a preprint and lacks independent or hardware validation, so further replication and real-world testing are needed to assess transferability. [1]

GO FURTHER

- Clone the MARE repository, run Terrain Studio to create a DEM, and launch the Unreal Engine project to record a stereo camera sequence with ground-truth shadows. [1]
- Use the provided ROS 2 GNC stack to send simple waypoint commands to the simulated rover and observe the web-based mission control interface. [1]

SOURCES AND CONFIDENCE

Medium confidence. Based on a paper; a preprint with no stated peer review or replication. Only the abstract of the paper was read; its tables and experimental conditions were not checked. One further page could not be used.

1. Paper: MARE: An Open-Source Lunar Rover Simulator with Shadow Ground Truth… 20 Oct 2026
   arxiv.org/abs/2610.28465

Sources read 25 Oct 2026. Links are for verification; this entry is meant to stand on its own.

TECHNICAL BRIEFING · ROBOTICS INDUSTRY · FUNDING · 20 OCT 2026 · TECHCRUNCH

Eka is worth $1.45B thanks to a physical AI model that has experts raving

Evidence status: Reported by the company: no independent verification

Eka Robotics announced a $150 million Series A round that values the company at $1.45 billion post-money, led by a16z and F-Prime. The founders, MIT professor Pulkit Agrawal and former DeepMind researcher Tuomas Haarnoja, said the funding will accelerate development of their robotic arm that can pick about 1,500 objects per hour. The company also highlighted that its gripper costs under $1,000 and is rated for 5 million duty cycles. [1]

HOW IT WORKS

Eka's system combines a force-sensitive gripper—equipped with extra force sensors that react to contact—to avoid crushing items, with a deep-learning model trained in a physics simulator that models multiple contact forces. The simulator generates thousands of hours of robot experience, which feeds a reinforcement-learning environment that learns a vision-force-action policy, i.e., the robot uses vision for coarse pose, force feedback for the actual grasp, and then executes the action. [1]

Estimate object pose → **Apply force-sensitive gripper contact** → **Feedback forces to model** → **Execute grasp action**

WHAT IS GENUINELY NEW

Unlike vision-language-action models that map images and text to actions, Eka's vision-force-action approach explicitly incorporates tactile information, allowing the robot to adapt to many objects without retraining. The company also claims to achieve this at scale by training in simulation rather than collecting massive video datasets. [1]

EVIDENCE

REPORTED FIGURE An anonymous frontier-lab roboticist described seeing the robot in action as an emotional experience comparable to early legged-locomotion breakthroughs. [1]

LIMITATIONS AND OPEN QUESTIONS

- The performance figure here comes from the company; the sources do not describe the test conditions, the objects or sites involved, or any independent measurement.
- The company has not disclosed how the approach is implemented; the source says details are withheld.
- The $150 million round measures investor interest, not performance: the pilots and deployments mentioned are announcements, and the sources report no results from them.

WHO SHOULD CARE, AND WHY

> BEGINNER
>
> Imagine a robot hand that not only sees a cup but also feels how hard to squeeze it, like how you gently pick up a fragile egg. By training in a computer game that mimics real-world forces, the robot learns quickly without needing thousands of real-world tries.

WHY AN ENGINEER SHOULD CARE

Eka's approach shows that adding force sensors to a gripper and training a vision-force-action policy in a physics simulator can dramatically boost pick-and-place speed. When designing manipulation systems, consider integrating tactile feedback and using simulated reinforcement-learning loops to generate massive experience data. [1]

RESEARCHER / PROFESSIONAL

The claims are company-sourced and lack independent benchmarks; the architecture remains opaque, and the funding does not prove real-world deployment success. Nonetheless, the shift to a vision-force-action paradigm and simulation-scale training could influence future research on dexterous manipulation, provided the gaps are addressed. [1]

GO FURTHER

- Read about force-sensitive grippers and how they use tactile sensors to modulate grip force. [1]
- Explore reinforcement-learning environments that use physics simulators for robotic grasping. [1]

SOURCES AND CONFIDENCE

Medium confidence. Based on a press report; figures that come from the company alone, with no independent verification.

1. Press report: TechCrunch: Eka is worth $1.45B thanks to a physical AI model that has… 20 Oct 2026
   techcrunch.com/2026/10/20/eka-is-worth-1-45-billion-thanks-to-…

Sources read 25 Oct 2026. Links are for verification; this entry is meant to stand on its own.

Short news

PanoGraphGS: Unifying Panoptic Segmentation and Hierarchical 3D Scene Graphs with Gaussian Splatting

Evidence status: Reported by the paper's authors (preprint): no peer review or independent replication

The authors introduce PanoGraphGS, a framework that learns a shared 3D Gaussian representation of a scene and simultaneously produces panoptic segmentation and a hierarchical 3D scene graph. They evaluate it on the Replica benchmark and a new real-world dataset called NaviGS, reporting more than 50 percentage-point gains in object-node accuracy over the strongest baseline. The code and dataset are released publicly. [1]

What is new. The contribution is a single model that jointly learns panoptic segmentation and hierarchical scene graph construction from a unified Gaussian representation, rather than treating them as separate pipelines. This integration enables direct Gaussian-level decoding and improves object-node identification accuracy. [1]

Limits. This is a preprint: the sources we could read show no peer review and no independent replication of the results.

[1] Paper: PanoGraphGS: Unifying Panoptic Segmentation and Hierarchical 3D Scene...
arxiv.org/abs/2610.27696 · [2] Code / documentation: gaojh0333/PanoGraphGS
github.com/gaojh0333/PanoGraphGS

Generative Bionics to launch Physical AI Data Factory in Italy

> **Evidence status:** Reported by the company: no independent verification

The first Data Factory hub is set to become one of Europe's largest physical AI training facilities for humanoid robots. The partnership builds on AMD Ventures' participation in Generative Bionics' €70 million seed funding round, while AMD computing technology is also used on board the Gene.01 robot. [1] [2]

[1] Press report: The Robot Report: Generative Bionics to launch Physical AI Data Factory in Italy...
therobotreport.com/generative-bionics-to-launch-physical-ai-da... · [2] Press report: Robotics & Automation News: Generative Bionics to build European 'physical AI data factory'...
roboticsandautomationnews.com/2026/10/23/generative-bionics-to...

How this digest is made

Stories are collected and ranked automatically, then researched: for each one we read the article and any paper, code repository or company page it links to, keep only the sentences that are about that story, and record who says what. The briefing is written from those facts and every statement is checked against the facts it cites. A story appears only if the reader learns something the headline did not say; thin or unverifiable stories are left out rather than padded. Company and founder statements stay attributed to them, preprints are labelled as preprints, and interpretation is marked as editorial. Source links are for verification; each story is meant to stand on its own. Generated 25 Oct 2026.

Headlines and images belong to their publishers. All digests · FixTheRobot News