8th Robot Learning Workshop: Is Physical AI Going Zero-Shot?

http://www.robot-learning.ml/2026

Anonymous Author(s)
Affiliation
Address
email

Abstract

The year 2025 has seen an unprecedented acceleration in the scale and diversity of foundation models for robotics with the trend continuing in 2026. With the 8th Robot Learning workshop, returning to its native venue at NeurIPS-2026, we explore a provocative question: Is Physical AI going zero-shot? Driven by the large-scale training of foundation models across diverse data sources and embodiments, we are witnessing the emergence of agentic approaches that allow robots to reason their way through complex, unseen tasks. We will critically examine how these paradigms shift the traditional boundaries of robot learning. Are we moving past narrow task-specific fine-tuning toward reasoning-based physical agents? What are the implications of scaling laws, diverse training datasets, and multi-modal models in real-world deployment? Can generalisation alone lead us to the strong performance levels required in robotics use cases? We seek diverse perspectives from the machine learning and robotics communities—both academia and industry—to map the trajectory of zero-shot physical AI. Capitalizing on our prior experience with robotics showcases, we will solicit several robotics researchers and companies to exhibit their latest platforms during the workshop's poster sessions to ground these models in physical reality.

Importance, freshness, and relevance to the NeurIPS community The Robot Learning workshop is now in its 8th edition and returns to NeurIPS with a strong foundation, but this proposal does not rely on past success alone. Previous NeurIPS editions have demonstrated sustained community interest, regularly attracting 100–200 attendees; the NeurIPS 2023 edition received 80 submissions, of which 45 were accepted. More importantly, the workshop has developed a distinctive format that connects machine learning advances with physical robotic systems through invited talks, contributed presentations, posters, panels, and robotics demonstrations.

For NeurIPS 2026, we propose a focused and timely theme: *Is Physical AI going zero-shot?* The rapid growth of robotics foundation models in 2025 has created a new set of questions for the NeurIPS community. Large-scale training across diverse embodiments, datasets, tasks, and modalities is beginning to challenge the traditional robot-learning pipeline based on narrow task-specific training and fine-tuning. At the same time, robotics remains a domain where generalization must be tested against physical constraints, safety requirements, uncertainty, embodiment mismatch, and real-world deployment failures. This makes zero-shot Physical AI not only an exciting frontier, but also a critical stress test for foundation models, world models, multimodal reasoning, reinforcement learning, planning, perception, and embodied decision-making.

The workshop will use the zero-shot Physical AI theme as a focused lens for open discussion among researchers in robot learning, foundation models, reinforcement learning, perception, planning, uncertainty, world models, autonomous systems, and real-world AI deployment. The program

Submitted to 40th Conference on Neural Information Processing Systems (NeurIPS 2026). Do not distribute.

structure will allocate substantial time to contributed talks, posters, panels, and interactive discussion, ensuring that new work and early-career researchers are visible alongside invited perspectives. Building on our prior experience with robotics showcases, we will also invite selected academic and industry groups to demonstrate physical platforms during the poster sessions, so that discussions of foundation models and agentic Physical AI remain grounded in real robotic systems.

By leveraging the Sydney venue, the workshop will further connect the global NeurIPS audience with the growing robotics and AI ecosystem in Australia and the broader Asia-Pacific region. Local robotics demonstrations from universities, startups, and industry labs, together with spotlight opportunities for Australian and regional early-career researchers, will help build a more active, internationally visible, and better connected robot learning community in the region. In this way, the workshop will serve both as a timely scientific forum on zero-shot Physical AI and as a community-building platform linking NeurIPS, robotics, and real-world embodied AI.

Scope of contributions and submission review philosophy We will solicit non-archival contributed papers on topics including, but not limited to:

- *zero-shot, few-shot, test-time, and continual adaptation for robot learning;*
- *vision-language-action models, multimodal foundation models, and generalist robot policies;*
- *world models, generative simulators, predictive models of dynamics;*
- *agentic reasoning, planning, memory, tool use, and long-horizon decision-making;*
- *cross-embodiment, cross-task, and cross-environment generalization;*
- *large-scale robot data, synthetic data, simulation, digital twins, data curation, and data governance;*
- *evaluation protocols, benchmarks, and stress tests for zero-shot physical autonomy;*
- *safety, uncertainty, interpretability, calibration, robustness, and failure analysis in deployed robots;*
- *real-world deployment studies, negative results, and lessons from industry-scale robotic systems;*
- *efficient adaptation, on-device deployment, and robust perception-action systems.*

Submissions will be managed through OpenReview and reviewed *double-blind*, with a sufficiently large program committee to provide three reviews per paper without overloading reviewers. Organizers will not review conflicted papers, and submissions from organizers, their PhD students, postdocs, or close collaborators will not be eligible. Accepted papers will be non-archival, with dissemination through the workshop website or OpenReview communicated clearly to authors.

Reviews will prioritize technical quality, relevance, clarity of claims, and evidence for generalization or deployment. We especially welcome papers that clarify limitations, report failures, propose evaluation protocols, or compare definitions of zero-shot Physical AI, so that the workshop helps establish shared standards and open problems rather than simply showcasing isolated demos.

Workshop format and interaction plan We propose a full-day in-person workshop with approximately 7–9 hours of session content, structured to encourage discussion rather than a talks-only format. The tentative program is:

- **Opening remarks and framing** (10 min): definition of the workshop questions and review of the zero-shot Physical AI debate.
- **Invited talks** (7 talks, approximately 25 min each including questions): covering foundation models for robotics, robot policy learning, embodied generalization, safe deployment, autonomous systems, and real-world robotics.
- **Contributed lightning talks** (40–50 min): short presentations from selected accepted papers.
- **Poster session and robotics demonstrations** (90 min): contributed posters alongside physical robot demonstrations or platform showcases where feasible.
- **Interactive panel/debate** (50 min): *"Is zero-shot Physical AI a realistic near-term goal, or a misleading framing?"*
- **Open discussion session** (30 min): community-driven discussion on benchmarks, datasets, safety practices, and future coordination.

- **Best paper and runner-up presentations** (20–30 min).
- **Breaks and informal interaction time**: distributed throughout the day to support community building.

The poster/demo session is an important part of the workshop design. It will ground abstract discussion of foundation models and zero-shot claims in physical systems, deployment constraints, hardware limitations, and real-world failure modes. We will invite robotics researchers and companies to present platforms or videos of physical deployment, subject to venue logistics and safety requirements.

Timeline If accepted, we will follow the NeurIPS 2026 workshop timeline and notify all contributing authors before the mandatory deadline. Our tentative timeline is:

- Workshop website and call for papers released: shortly after acceptance notification.
- Submission deadline for contributed papers: August 22, 2026 AoE.
- Review period: late August to late September 2026.
- Author notification: no later than September 29, 2026 AoE.
- Camera-ready / final version deadline: mid-October 2026.
- Final program, accepted paper titles, and speaker information posted: early November 2026.

Location preference Our location preference is: **Sydney only**. Sydney is our first preference because it is the main NeurIPS 2026 site, can host larger workshops, and is geographically well aligned with several confirmed or tentative speakers and organizers in Australia and the Asia-Pacific region. It also provides a strong opportunity to engage the growing robotics, autonomous systems, and embodied AI communities in Australia and the broader region.

Concurrent related proposals (YL: To be confirmed before submission:) all organizers will disclose whether they have submitted concurrently any related proposal to another venue and whether they are involved in any other workshop proposals for NeurIPS 2026. The final submission will include a complete statement. If there are no such submissions, we will state: *"The organizers are not concurrently submitting a related workshop proposal to another venue. Roberto Calandra is submitting another workshop proposal to NeurIPS 2026. No organizer is listed on more than two NeurIPS 2026 workshop proposals."*

Invited speakers The following invited speakers have been identified to provide complementary perspectives across robot learning, embodied intelligence, large-scale robot models, autonomous systems, and real-world deployment:

Confirmed *in-person* (Sydney)
- Weiming Zhi (University of Sydney, Australia)
- Huazhe Xu (Tsinghua University, China)
- Maria Bauza Villalonga (Google Deepmind, UK)
- Michael Milford (Queensland University of Technology, Australia)
- Yadan Luo (University of Queensland, Australia)

Tentatively Confirmed *in-person* (Sydney)
- Ingmar Posner (University of Oxford, UK)
- Mengdi Xu (Tsinghua University, China)

This speaker set is designed to cover multiple dimensions of the workshop theme: large-scale robotic learning, embodied generalization, robot perception and navigation, autonomous systems, academic and industry perspectives, and both optimistic and critical views on the zero-shot Physical AI framing. We will reconfirm speaker availability for the final assigned NeurIPS location and will prioritize in-person participation.

Diversity, inclusion, and community-building plan Diversity, inclusion, and community building are central to this workshop. Building on the continuity of the workshop series, we have assembled an organizing team and speaker list that reflect diversity across geography, institution type, career seniority, gender, research background, and scientific perspective, with participants from academia, industry research labs, and robotics-focused institutions across Australia, Europe, and Asia. The team combines experienced organizers of previous robotics and machine learning workshops with first-time organizers, and the program includes perspectives ranging from PhD students and postdoctoral researchers to senior academics and industry research leaders. Where appropriate, we will encourage senior speakers to share their presentations with PhD or postdoctoral researchers, and we will use contributed talks, posters, and panels to provide visibility for early-career researchers and emerging directions. The workshop also has a concrete track record of supporting broader participation. In recent editions, with financial support from sponsors including Frodobots, Microsoft, Naver Labs Europe, Google, DeepMind, NVIDIA, and Toyota Research Institute, we covered NeurIPS registration fees for a number of researchers from underrepresented groups. These opportunities were advertised through communities such as Women in ML, Black in AI, Queer in AI, and LatinX in AI. For NeurIPS 2026, we plan to continue this practice, subject to sponsorship, and to advertise support broadly through relevant community networks.

Expected attendance and previous workshop statistics We estimate **300–400 attendees** and **70–90 submissions**, with approximately **30–50 accepted papers** depending on submission quality and capacity. This estimate is based on the history of the Robot Learning Workshop (last year we accepted 49 of 74 submissions) series and the broader interest in Physical AI, robotics foundation models, and embodied generalization. The Robot Learning Workshop held last year at ICLR in Singapore filled a 300 person room where the growing interest in physical AI supports the lifted submissions and attendance estimates. Prior NeurIPS editions (2017, 2019–2023) regularly attracted approximately 100–200 attendees. The NeurIPS 2023 installment received 80 submissions and accepted 45 papers.

Special requirements and technical needs The workshop will use the standard NeurIPS workshop room setup, including projector, podium microphone, Q&A microphone, panel table, and poster boards. We also request: sufficient poster-board space for contributed posters; standard live-streaming and recording support; one panel table with microphones for the interactive debate; space near the poster area for small robotics demonstrations or table-top platform showcases, subject to venue safety policies; access to power strips or nearby power outlets for demos, if permitted by the venue. Any physical demonstrations will be planned conservatively and will comply with NeurIPS venue rules. If physical robots cannot be accommodated, we will use video demonstrations and interactive poster materials instead.

Difference from similar workshops and previous iterations Many recent workshops have addressed robot learning, foundation models, embodied AI, reinforcement learning, simulation, or robotics deployment. This workshop differs in its specific framing around **zero-shot Physical AI**. Rather than focusing broadly on scaling robot learning or showcasing foundation-model demonstrations, the workshop asks what zero-shot capability means in physical systems, what evidence is required to support such claims, and what scientific bottlenecks prevent reliable deployment.

Compared with previous Robot Learning Workshop editions, this proposal introduces a sharper and more provocative theme. Prior editions examined scalability, safety, lifelong learning, self-supervised learning, and large models in robotics. The 2026 edition centers on the current tension between rapid progress in large-scale embodied models and the persistent limitations of real-world robot deployment. It will emphasize discussion, failure analysis, evaluation standards, and community coordination, rather than relying primarily on invited talks or past workshop success.

Website Workshop website: `http://www.robot-learning.ml/2026`

Tagline **Is Physical AI truly going zero-shot? Join the debate on foundation models, robotics, deployment, and embodied generalization.**

Organizing team's experience and short bios

This workshop is being managed by a highly experienced team of 8 organizers where most have organized a workshop before. It is the 8th robot learning workshop in a series that was held at NeurIPS in 2017 and 2019 – 2023 as well as at ICLR in 2025. Half of the organizers (Jen Jen Chung, Jiaxu Xing, Moritz Reuss, Johannes Busch) are new to the organizing committee with the remaining have organized workshops of this series in previous years.

The extended workshop team includes 2 advisors, who aren't formal organizers of this installment of the Robot Learning workshop, but participated in running its previous installments at NeurIPS and are providing feedback to the official organizing team on the best practices and speaker selection.

The organizers' bios and the list of advisors are below:

Andrey Kolobov | akolobov@microsoft.com | G. S. | Microsoft Research, USA
Andrey Kolobov is a Principal Research Manager at Microsoft Research (MSR), the head of MSR's **S**pecial **P**hysically **E**mbodied **AI** and **R**obotics team (SPEAR), and the person responsible for MSR's internal and external partnerships in the physical AI space. His research focuses on training physical AI models for robotic manipulation, with a special focus on making them easily adaptable to new tasks and environments. His prior work on glider drones was featured in Nature Communications and New York Times, and his RL algorithms for Web crawling power Microsoft Bing search engine. Andrey holds 7 Outstanding PC Member awards from NeurIPS, ICML, ICLR, IJCAI, and AAAI. He received his Ph.D. in CS from the University of Washington, Seattle and a double B.A. in CS and Applied Math from the University of California, Berkeley. Andrey co-organized the 7th Robot Learning Workshop @ ICLR-2025, which attracted 300+ in-person attendees.

Alex Bewley | bewley@google.com | G. S. | Google DeepMind, Zurich, Switzerland
Alex is a Senior Research Scientist at Google DeepMind leading the post-training of vision language action models with prior work in related fields from robot-perception to human-level competitive robot table tennis. His research interests cover multi-task learning, policy gradient, visual attention and sequence modeling, model distillation and interpretability for efficient and reliable robotics. Alex has previously co-organized several workshops at NeurIPS'20 '21 '22, CVPR'21 and ICLR'25.

Hamidreza Kasaei | hamidreza.kasaei@rug.nl | G. S. | University of Groningen, Netherlands
Hamidreza is an Associate Professor in the Department of Artificial Intelligence at the University of Groningen, the Netherlands. He runs the IRL-Lab focusing on Lifelong Interactive Robot Learning in the areas of 3D Object Perception, Grasp Affordance, and Object Manipulation. He holds the Google Research Award 2023 and the Robotics: Science and System Pioneers Award 2019. Hamidreza has organized workshops at IROS'19, RSS'19, NeurIPS'21, '22, '23, ICLR'25, ICML'25, and also served as an associate editor for the RA-L journal and ICRA and IROS since 2020, as well as area chair for RSS '26. He was selected as an outstanding associate editor for the RA-L journal in 2023. Prior to joining the University of Groningen, Hamidreza was a visiting scholar with the Imperial College London, Computer Vision and Machine Learning lab.

Roberto Calandra | rcalandra@lasr.org | G. S. | TU Dresden, Germany
Roberto Calandra is a Full (W3) Professor at the Technische Universität Dresden where he leads the Chair of Machine Learning for Robotics. Previously, he founded at Meta AI (formerly Facebook AI Research) the Robotic Lab in Menlo Park. Prior to that, he was a Postdoctoral Scholar at the University of California, Berkeley (US) in the Berkeley Artificial Intelligence Research (BAIR) Lab. His education includes a Ph.D. from TU Darmstadt (Germany), a M.Sc. in Machine Learning and Data Mining from the Aalto university (Finland), and a B.Sc. in Computer Science from the Università degli Studi di Palermo (Italy). His scientific interests are broadly at the conjunction of Machine Learning and Robotics, with the goal of making robots more intelligent and useful in the real world. Roberto served as Program Chair for AISTATS 2020, as Guest Editor for the JMLR Special Issue on Bayesian Optimization, and has previously co-organized over 19 international workshops (including at NeurIPS, ICML, ICLR, ICRA, IROS, RSS). In 2024, he received the IEEE Early Academic Career Award in Robotics and Automation.

Johannes V. S. Busch | jbusch@lasr.org | G. S. | TU Dresden, Germany
Johannes is a PhD student supervised by Prof. Roberto Calandra at the Learning, Adaptive Systems, and Robotics Lab, TU Dresden, Germany. He previously obtained his Dipl.-Ing. in Electrical

Engineering at TU Dresden and held a research position at the Telekom Chair of Communication Networks. His primary research focus is generative modeling for hierarchical decision making.

Moritz Reuss | mreuss@nvidia.com | G. S. | NVIDIA, Switzerland
Moritz Reuss is a Researcher at NVIDIA's Seattle Robotics Lab. He completed his Ph.D. at the Karlsruhe Institute of Technology (KIT) under the supervision of Prof. Rudolf Lioutikov, focusing on generative imitation learning for robotic manipulation. His research interests include diffusion- and flow-matching-based robot policies, vision-language-action models, and scalable learning methods for generalizable manipulation. During his Ph.D., Moritz was a research intern at Apple and was supported by the Apple Scholar in AI/ML fellowship.

Jiaxu Xing | jixing@ifi.uzh.ch | G. S. | University of Zurich, Switzerland Jiaxu Xing is a PhD student at the Robotics and Perception Group, University of Zurich. He is also an associated doctoral Student at ETH AI Center. Previously, he completed his Master in Robotics at ETH Zurich. His research focuses on developing high-performance, generalizable control policies for agile robots that can be effectively deployed in real-world environments. He has a publication record at ICLR, ICRA, IROS, CoRL, and RSS, and was invited to present at several ICLR and CoRL workshops.

Jen Jen Chung | jenjen.chung@uq.edu.au | G. S. | The University of Queensland, Australia
Jen Jen Chung is an Associate Professor in Mechatronics within the School of Electrical Engineering and Computer Science at The University of Queensland, and she is also a Visiting Scientist at CSIRO. Her current research interests include perception, planning and learning for robotic mobile manipulation, algorithms for robot navigation through human crowds, informative path planning and adaptive sampling. Prior to working at UQ, Jen Jen was a Senior Researcher in the Autonomous Systems Lab (ASL) at ETH Zürich from 2018-2022 and was a Postdoctoral Scholar at Oregon State University researching multiagent learning methods from 2014-2017. She completed her Ph.D. on information-based exploration-exploitation strategies for autonomous soaring platforms at the Australian Centre for Field Robotics in the University of Sydney. She received her Ph.D. (2014) and B.E. (2010) from the University of Sydney.

<u>**Advisors:**</u> The advisors are past organizers and senior members of the community who help the organizing committee by providing feedback and proposing speakers. Additionally, they also contribute to the program committee below.

- **Markus Wulfmeier** (Nomagic, Switzerland/Poland)
- **Masha Itkina** (Research Lead and Manager, TRI, USA)

Program Committee The workshop will be supported by a Program Committee composed of researchers from academia and industry with expertise spanning robot learning, reinforcement learning, embodied AI, computer vision, and foundation models. We expect between 70 and 100 submissions based on participation in recent editions of the workshop and growing interest of this field.

To ensure a rigorous yet scalable review process, we will follow a community-reviewing model similar to that used successfully in previous years. Authors submitting to the workshop will be asked to indicate their willingness to review other workshop submissions. The requirement that authors of submissions will be required to review at least three submissions with a maximum of three reviews per author will be included in the call for submissions, the workshop website and in the OpenReview submission page. Reviews and conflict management will be handled through OpenReview, which automatically identifies conflicts of interest based on institutional affiliations and collaboration history. Organizers will adhere to the same conflict of interest checks, oversee reviewer assignment and monitor review quality throughout the process.

Our goal is to obtain at least three reviews per submission while maintaining a reviewer load of no more than three papers per reviewer. This approach allows us to provide thorough feedback despite the relatively short reviewing period between the suggested submission deadline (August 29, 2026) and the mandatory notification deadline (September 29, 2026). Additional invited Program Committee members will supplement author-reviewers where necessary to ensure adequate expertise and coverage across all submissions. For these additional reviewers we can draw upon prior year's program committees on a case-by-case basis where some papers may not find relevant non-conflicting reviewers in the existing pool of reviewers.