7TH ROBOT LEARNING WORKSHOP: TOWARDS ROBOTS WITH HUMAN-LEVEL ABILITIES

HTTP://WWW.ROBOT-LEARNING.ML/2025

ABSTRACT

The year 2024 has seen an explosion of interest in humanoid robots. In the 7th Robot Learning workshop, to be held at ICLR-2025, we will look beyond the humanoid *embodiment* and ask: how far are we from robots with human-level *abilities*? What do we need to improve about embodied learning, decision-making, perception, and data collection to train generally physically capable robots to robustly perform a wide range of activities such as cooking or tidying up the house – activities that people do without much thinking? We believe many of the weaknesses of the current robotic systems to be a reflection of the shortcomings of general AI methods and models. As such, in this workshop we will seek diverse perspectives from robotics-focused and robotics-orthogonal parts of the ICLR community alike, scientific contributions from academia and industry, as well as participants from a variety of backgrounds and career stages. Capitalizing on our prior experience with robotics showcases, in keeping with the spirit of the times we will solicit several humanoid robotics companies to exhibit their robots during the workshop's poster sessions.

Workshop summary Since late 2023, the public has witnessed an explosive growth of interest in humanoid robotics in the technology world. Announcement of new humanoid robots from companies including Unitree, 1X, and Sanctuary AI, ostensible progress on their capabilities demonstrated by, e.g., Tesla Optimus, as well as deals between smaller companies such as Figure AI and Sanctuary AI with large but non-robotics-focused players like OpenAI and Microsoft unmistakably point at the industry's expectations: the humanoid form factor and large models are about to make robots as skillful as people across a range of tasks. Yet, on the research side, the picture is very different. Even SOTA ML-based approaches are just beginning to enable humanoid robots to make first steps to autonomy (Fu et al., 2024), while impressive, sometimes human-level, performance is shown on non-humanoid robots, without the use of large models, or both (D'Ambrosio et al., 2024; Kaufmann et al., 2023; Zhao et al., 2024). However, some of these methods rely on ML in simulation being sufficiently powerful and others count on vast data quantities that are currently available only for specific kinds of activities, such as object manipulation with parallel grippers or navigation. In addition, a sizeable fraction of the research community at the intersection of ML and robotics holds that even if enough robotics data was available, we still lack a scientific understanding of ML to train models that perceive and act sufficiently reliably for controlling a robot in weakly structured environments like homes, where people operate with ease. In the meantime, the public expectations of robots' human-level abilities are so high that they can easily bring about "AI robotics winter" if the research community fails to bridge the gap between them and reality. This is an exciting chance for researchers and practitioners alike.

The theme of the Robot Learning workshop at ICLR-2025 will revolve around a group of inter-related questions about attainability of human-level performance in robots. **(1)** How far are we from robots with general human-level abilities? **(2)** What are we lacking from hardware, modeling, machine learning, and human-robot interaction perspectives to get there? **(3)** In which activities can robots *already* match humans, and what techniques enabled this? **(4)** What, besides data, prevents learning approaches for large multi-modal models from producing a artificial motor cortex that would be general-purpose for physical activities that don't require AGI-level reasoning capabilities, e.g., assembling a toy car out of a Lego kit? **(5)** What ML methods and data do we use to make robots behave as *physically intelligently* around people as people do, e.g., by yielding near doorways yet not blocking traffic through the doorway, etc?

To answer these questions, we will seek diverse perspectives from robotics-focused and robotics-orthogonal parts of the ICLR community alike, scientific contributions from academia and industry, as well as participants from a variety of backgrounds and career stages. **We caution that this workshop is not about debating the feasibility of AGI, measuring it, or pathways towards creating it.** It is

about fleshing out a joint ML+robotics research agenda whose outcome can plausibly enable robots to robustly perform a very wide range of activities like cleaning up in the house or cooking in a kitchen – activities that, most people would agree, don't require intense higher-level cognition. **The goal of the workshop is for the attendees to walk away with an end-to-end idea of what this agenda could be.**

Scope of contributions, outreach, attracting novel work, and submission review philosophy A call for contributions will be announced via the workshop website and to related research communities via social media, soliciting novel and unpublished work on a list of topics relevant to questions **(1)** – **(5)** above. Given this year's surge of general interest in robotics among the ML researchers and practitioners (in humanoid robotics in particular), given the organizing team members' outreach in the relevant robotics research subcommunities, and given ICLR-2025's location in Asia – a booming center of humanoid robotics – we expect the workshop to easily attract a large audience.

We also expect it to attract a lot of novel work, simply because of the tremendous progress in robotics over the past year and because addressing the key questions considered in this workshop *requires* novel research. We recognize that a lot of work in robot learning, large model training for robotics, RL, imitation learning, etc. *can be* connected to our workshop's theme questions. However, in many works this connection is non-apparent or tenuous, and we will be evaluating the submissions partly on how clearly they articulate it as well as on technical novelty per se.

The review process for contributed material will be double-blind to reduce institutional and author bias, and the program will be selected with an eye to establishing broad coverage of research areas linked to the workshop's theme while preserving merit awarded by the double-blind reviewing process. The workshop will, as a result, benefit from a more diverse set of participants, since contributed-work speakers will be invited to participate along with the invited ones.

Platforming contributed work and Tiny Papers Our workshop will have a Full Papers and a Tiny Papers track. The former will invite papers with a recommended length of 4 to 10 pages but without a strict page limit, as we don't want the authors to waste their time on formatting issues, and will otherwise have expectations that are standard for full-length workshop papers. The latter will invite submissions as per the ICLR-2024 Call for Tiny Papers on topics that align with the workshop's theme.

Every accepted submission will be presented as a lightning talk and as a poster. For more details, see the **Workshop format and tentative schedule** section. To promote the accepted Tiny Papers, help their authors integrate into the research community, and give them an opportunity to gain more experience with presenting their work, we will have **the Tiny Papers' authors present their findings in the same lightning-talk and poster sessions as the accepted full papers**. The slides from the lightning talks with highlights of the paper's content will be posted on the workshop webpage. We will encourage the authors to upload thumbnails of their posters to the webpage as well.

Note that **our workshop is non-archival, which is stated in the Call for Papers on the workshop webpage**, with the caveat that the Tiny Papers are subject to the non-workshop-specific rules in the ICLR-2024 Call for Tiny Papers. However, the pdfs of their accepted papers available on the workshop webpage.

In summary, the accepted work will be publicized via 3 channels: oral lightning talks, poster presentations, and the workshop webpage (using the paper pdf as well as lightning talk slides and, optionally, poster thumbnails).

In-person vs. virtual attendance The workshop will be an in-person event. In particular, we expect the presenters of accepted papers and invited talks, as well as the panelists, to be at the workshop in-person, as we believe that the opportunity for spurious in-place interactions with them is crucial for a vibrant scientific discourse.

Exceptions will allowed in 3 cases: (1) at most 1 invited talk, (2) for the Tiny Papers' authors who may be unable to travel due to the lack of funding; they will still be expected to present their work as a lightning talk via Teams or Zoom, (3) presenters facing clearly unforeseen issues; we will consider them and ways to handle them on a case-by-case basis.

Opportunities to engage for those who are unable to attend + access to workshop materials. While the presenters are expected to attend in-person as described above, we will broadcast the invited talks, panel/debate, and lightning talks on the Web, conditioned on the appropriate equipment being available in the workshop room. Subject to equipment availability and speaker agreement, we will also record the workshop talks and post the recordings online later.

As mentioned previously, the contributed papers will also be available on the workshop webpage along with lightning talk slides and poster thumbnails.

Workshop format and tentative schedule The proposed format reflects a full-day **8-hour** in-person workshop:

- **3.5 hours of invited talks** (7 or 8 invited talks of $\sim$ 30 min, including 5 min for questions for each)
- **45 min of interactive panel debate**, which will include the audience as well as the panelists
- **35 min of contributed lightning talks**, where each accepted submission, including Tiny Papers, will be presented; see the **Attendance** section for an estimate of the number of accepted papers
- **1.5 hours of contributed poster sessions + *humanoid* robot demos** (2 poster sessions of 45 min each.) We are planning to solicit several humanoid robot companies to bring their robots to our workshop. While we recognize the companies' potential challenges of transporting these robots and their batteries, we predict that several of them will be willing to do it due to Singapore being close many of the robotics companies' centers, which are in Asia.
- **1 hour of lunch**
- **30 min for Best Paper and Runner-up Award-winning paper presentations** (15 min for each)
- **10 min of opening and closing remarks**

Interaction opportunities. The lunch and two poster sessions will break up the day into 4 parts and serve as natural opportunities to discuss workshop content in an informal setting. Depending on the availability of lunch at the venue (we will find out about this in advance), we will also attempt to find a sponsor for a catered/box lunch to be delivered to the workshop room, in order to encourage the participants to stay there around lunch time and thereby increase the interaction opportunities. **Thus, we are aiming for 2.5 hours of unstructured interaction and 45 min of interaction during the panel/debate.**

Panel/debate topics. The choice of panelists and exact topics will depend on the final invited speaker lineup as well as which prominent researchers working in ML and robotics other than the invited speakers will attend ICLR in person. It will be finalized closer to the workshop. It may include questions **(1)** — **(5)** from the **Workshop summary** as well as - How can insights from fields outside of robotics, such as cognitive science, contribute to developing robots with human-level abilities? - What are the metrics and benchmarks that really help us evaluate how close we are to human-level intelligence and abilities?

Timeline and organization We will have a general contributed paper **submission deadline on February 3, 2025 AOE**. We will release **submission decisions on February 21, 2025 AOE**, to give the authors ample time for travel arrangements. We will import the workshop program and **accepted paper titles to iclr.cc and workshop webpage by March 25, 2025 AOE**. We will have **invited talk titles and the full workshop schedule on the workshop webpage by April 1, 2025**. The tentative deadline for camera-ready versions of the accepted papers is April 11, 2025 AOE.

Previous and related events ICLR, ICML, and NeurIPS have hosted various workshops on ML application areas (e.g. NeurIPS'17, '19, '20, '21, '22, '23 Robot Learning, ICML'17-'19 autonomous driving workshop), application of a specific ML tool to real-world problems (e.g. ICML'19, '21 RL for real-life, NeurIPS'19, '20, '21, '22 Deep RL workshop), and the recent advancements in foundation models for decision making (NeurIPS '22, '23). Inspired by the vision of embodied intelligence, robotics has played an increasingly important role within the scope of machine learning applications.

Our workshop is a continuation of the Robot Learning workshop series that so far has been held only at NeurIPS (2017 and 2019 – 2023. **It will not be held at NeurIPS in 2024**). Previous iterations

of the workshop focused on scalability of robot learning, developments of embodied intelligence in real-world robotics, ensuring safety, reliability, and trustworthiness in human-robot interactions, and specific training settings, lifelong and self-supervised, that might be required. Our Robot Learning workshop, the 7th in this series, differs from all of them as we describe next.

Why is this robot learning workshop different, timely, and exciting? The past robot learning workshops have been driven by progress in ML research subfields, from Deep Learning to Deep RL to foundation models. In contrast, ours is ultimately motivated by the recent rapid progress robot in *hardware*, especially but not limited to humanoid robot hardware. On the one hand, this rapid progress has drastically increased *physical* capabilites of the robot hardware and raised the industry's and the public's expectations of robots' behavior. On the other hand, in the case of humanoid robots, these advances left even the SOTA ML research techniques barely able to take advantage of these physical capabilities. Resolving this mismatch is critical both for avoiding an "Embodied AI winter" and for helping the ML robotics research community understand how to proceed forward in this climate. Our workshop is a platform for doing so.

We also emphasize opportunity presented by ICLR-2025 to more closely involve the Asia ML robotics community in this discourse. Asia is home to a most of robotics companies, including likely the majority of humanoid robotics ones. However, major ML conferences are rarely in Asia, hampering research exchange between the local robotics community and the ML community from other parts of the world. At ICLR-2025 in Singapore, which makes logistics easier for participants from Asia, we hope to bring the two communities closer, among other things **by inviting leading robotics companies, many of which are Asia-based, to showcase their robots at our workshop**.

Attendance The past (NeurIPS) workshops of this series (http://www.robot-learning.ml/) have regularly attracted 100-200 attendees. At the NeurIPS-2023 installment of this workshop, the latest one and the first predominantly in-person one since 2019, we received 80 submissions, out of which 45 were accepted. ICLR has been smaller than NeurIPS-2023, but the size gap has been shrinking, and ICLR-2025 will be close to the vibrant robotics community in Asia that often has difficulties attending NeurIPS. Considering all these factors as well as our workshop's topic novelty, we expect **100-150 attendees** and **60-70 submissions**, out of which we will **accept 30-35**.

We expect this workshop be especially conducive to attracting ML robotics students from Asia, due to the vibrancy of this specific research community in this part of the world, limited travel budgets and visa issues that often prevent them from attending robotics workshops at major ML conferences taking place outside Asia, and ICLR-2025's location in Singapore.

List of Invited Speakers Below is a list of speakers:

Confirmed *in-person*
- Sandy Huang (Google DeepMind)
- Joseph Lim (KAIST)
- Kate Tsui (Toyota Research Institute (TRI))
- Davide Scaramuzza (University of Zurich)
- David Hsu (Nat'l University of Singapore)

Tentatively confirmed *in-person*:
- Chelsea Finn (Stanford / Physical Intelligence)
- Chris Paxton (Hello Robot)
- Eric Jang (1X Technologies)

Speaker selection philosophy. We carefully constructed the list of speakers to represent the gamut of perspectives on attaining human-level robot abilities as described in the **Workshop summary** and to be diverse according to the criteria mentioned in the **Diversity statement** below. In particular, **Chris Paxton** works on robots capable of considerable fluency despite being non-humanoid, **Davide Scaramuzza**'s research has enabled drones – a type of robots the robot community typically doesn't have in mind – to perform at super-human level at drone racing, **Kate Tsui** and **Sandy Huang** bring the human-robot interaction point of view, **Joseph Lim** can speak about the strengths and weaknesses of learning for robot decision-making, and **Chelsea Finn** represents the data collection and large

model training angle, and **Eric Jang** can talk about training robots that are actually humanoid in form factor.

Panelists. We will recruit the panelists at a later date, partly from the invited speakers and partly from among prominent community members who will decide to attend ICLR in person, which we will find out closer to ICLR itself.

Contingencies. We are confident that at least 2 of the speakers form the **Tentatively confirmed *in-person*** list will be able to come to ICLR in person after finalizing their plans. Nonetheless, the organizing committee is able to recruit other in-person speakers in their place *without sacrificing the diversity of the speaker set and without reducing the range of perspectives on the workshop theme.* We aim to have at most 1 invited talk delivered in-person, i.e., to have almost all the invited speakers attend in person and be available for interaction with other attendees.

Diversity statement We actively solicited – **and largely succeeded in recruiting, as can be verified by inspecting the corresponding lists** – organizers, speakers, and advisory board members from a range of diverse backgrounds, encompassing individuals with diverse gender, race, affiliation, nationality, geographic location, scientific background, and views on the theme of this year's workshop. Among organizers and speakers, the full scale of scientific seniority is covered, from PhD candidates to assistant and full professors. The organizing team includes both experienced organizers of previous robotics and ML workshops and several first-time organizers. Additionally, we encourage our senior presenters to share their talk with a PhD or postdoctoral researcher to provide a platform for early-career researchers. For the past three years, we were able to cover NeurIPS registration fees for many researchers in underrepresented minorities thanks to financial support from our corporate sponsors (Naver Labs Europe, Google, DeepMind (then a separate organization from Google), NVIDIA, and Toyota Research Institute). We advertised this resource to groups including Women in ML, Black in AI, Queer in AI, and LatinX in AI, and plan to do the same with this iteration of the workshop.

We have assembled a group of organizers and speakers that reflects diverse regional affiliations, research areas, career seniority, and perspectives from both academia and industry. Several of the workshop organizers and speakers belong to underrepresented minorities in ML and STEM at large as female, POC, and LGBTQ+ researchers.

Managing conflicts of interest We acknowledge that we have read the requirements for managing the conflicts of interest described in Guidance for ICLR Workshop Proposals 2025 and promise to abide by them.

Organizing team's experience and short bios

This workshop is being managed by a highly experienced team of 9: every organizer has co-organized at least one workshop before, and over 66% of the organizers have previously organized at least 3 workshops at top-tier ML and robotics conferences.

It is the 7th Robot Learning workshop in a series that was held at NeurIPS in 2017 and 2019 – 2023. Out of the 9 organizers this time, 3 (Anqi Li, Ted Xiao, Mahi Shafiullah) haven't organized workshops of this particular series before. Several of the remaining 6 have organized workshops of this series more than once.

The extended workshop team includes advisors – they 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 – the corresponding author.
Andrey Kolobov is a Principal Researcher and Research Manager at Microsoft Research (MSR). He leads MSR's efforts in large models for robotics, including collaborations with Microsoft's internal and external partners in this space. His research focuses on learning for decision making in multi-modal models. 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.

Dhruv Shah | dhruvshah@google.com | G. S. | Google DeepMind & Princeton University, USA
Dhruv Shah is a Senior Research Scientist at Google DeepMind and an Incoming Assistant Professor of Robotics at Princeton University. His work focuses on robot learning from large-scale datasets, studying generalization across embodiments, and long-horizon planning for mobile robots. He has been a core organizer of workshops at NeurIPS, CoRL, ICRA, RSS, and IROS, on topics around large-scale machine learning and open-world robotic deployment. Previously, he obtained his PhD at UC Berkeley, where he was a Berkeley Fellow and Microsoft Future Leader; his research has been nominated and won several Best Paper Awards at ICRA and RSS.

Anqi Li | anqil@nvidia.com | G. S. | NVIDIA, Seattle, USA
Anqi Li is a Research Scientist at NVIDIA. Her research focuses on bringing formal performance guarantees and sample efficiency to learning. Specific research topics include offline reinforcement learning (RL), safe RL, learning stable policies, learning from human demonstrations, and planning & control with guarantees. Anqi organized the RSS pioneer workshop at RSS 2023 and is organizing an upcoming RSS 2024 workshop on geometric and algebraic structure in robot learning. She received her PhD degree from the University of Washington and her Master's degree from Carnegie Mellon University. Her graduate research was partly supported by the NVIDIA Graduate Fellowship and the Siebel Scholarship. She was selected as an EECS Rising Star and an RSS Pioneer.

Feras Dayoub | feras.dayoub@adelaide.edu.au | G. S. | University of Adelaide, Australia
Feras is a Senior Lecturer at the Australian Institute for Machine Learning (AIML) at the University of Adelaide, Australia, where he leads a research group on Embodied AI. He served as a Chief Investigator (CI) in the ARC Centre of Excellence for Robotic Vision. His research interests cover autonomous perception in real-world environments, focusing on topics such as semantic scene understanding, out-of-distribution detection, and open-set recognition in object detection. Feras has organized workshops at CVPR in 2021, 2019, and 2018, RSS in 2018, IROS in 2020 and 2019, and ICRA in 2020 and 2018. He has also contributed as an Associate Editor for the IEEE Robotics and Automation Letters (RA-L) and ICRA.

Roberto Calandra | roberto.calandra@tu-dresden.de | G. S. | TU Dresden, Germany
Roberto Calandra is Full (W3) Professor at TU Dresden where he leads the Learning, Adaptive Systems and Robotics (LASR) lab. Previously, he was a Research Scientist at Meta AI (formerly Facebook AI Research), where he founded the Robotic Lab in Menlo Park, California. He was also a Postdoctoral Scholar at the University of California, Berkeley (US) in the Berkeley Artificial Intelligence Research Laboratory (BAIR) working with Sergey Levine. His education includes a Ph.D. from TU Darmstadt (Germany), a M.Sc. in ML 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 Robotics and ML, 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 previously co-organized over 16 international workshops (including at NeurIPS, ICML, ICLR, ICRA, IROS, RSS). In 2024, he received the IEEE Early Academic Career Award in Robotics and Automation.

Ted Xiao | tedxiao@google.com | G. S. | Google DeepMind, USA
Ted Xiao is a Senior Research Scientist at Google DeepMind working on robot learning. His research agenda focuses on scaling robot learning in the real world, with a particular focus on approaches that can leverage internet-scale foundation models and methods that improve with more experience. Prior to joining Google DeepMind, Ted founded Machine Learning at Berkeley and worked at Adobe Research. Ted received his B.S. and M.S. in Electrical Engineering and Computer Science from UC Berkeley, where he was advised by Professor Claire Tomlin. He has been a core organizer of workshops at NeurIPS 2022, CoRL 2022, ICLR 2023, ICRA 2023, and CoRL 2023 on topics such as robotics, language, and reinforcement learning.

Rika Antonova | rika.antonova@cst.cam.ac.uk | G. S. | University of Cambridge, UK and Stanford University, USA
Rika is an Associate Professor at the University of Cambridge, Department of Computer Science and Technology. Earlier, Rika was a postdoctoral scholar at Stanford University upon receiving the NSF/CRA Computing Innovation Fellowship from the US National Science Foundation (for research

on transfer learning for robotics). Rika completed her Ph.D. in CS at KTH, Stockholm. Earlier, she obtained a research Master's degree from the Robotics Institute at Carnegie Mellon University, where she developed Bayesian optimization methods for robotics and personalized tutoring. Before that, Rika was a senior software engineer at Google in the Search Personalization team and then in the Character Recognition team (developing open-source OCR engine Tesseract). Rika previously organized workshops at NeurIPS, ICRA, and RSS on physical reasoning, mobile and dexterous manipulation, and RL.

Nur Muhammad "Mahi" Shafiullah | mahi@cs.nyu.edu | G. S. | New York University, USA
Nur Muhammad "Mahi" Shafiullah is a Ph.D. student at the Courant School of Mathematical Sciences at New York University. His research focuses on household robots and the generalizable machine learning algorithms that are required to bring those robots into homes. During his Ph.D. Mahi has also spent some time at Fundamental AI Research (FAIR) at Meta as a visiting scientist, and has been supported by the AI/ML fellowship by Apple. Mahi has been a core organizer for workshops and tutorials at ICML, RSS, and ICRA on dexterous manipulation, supervised learning for robotics, and vision-language models for embodied intelligence. Prior to his Ph.D., Mahi completed his undergraduate and Masters at Massachusetts Institute of Technology.

Masha Itkina | masha.itkina@tri.global | G. S. | Toyota Research Institute, Los Altos, USA
Masha is a Research Scientist at TRI in the Robotics division. She completed her PhD at the Stanford Intelligent Systems Laboratory (SISL) advised by Prof. Mykel Kochenderfer. Her research focuses on trustworthy learning under uncertainty in large behavior models for robotics applications. Masha has a publication record at NeurIPS, ICRA, IROS, CoRL, and RSS, was invited to present at NeurIPS, WiML, ICRA, and RSS workshops, and co-organized the Robot Learning Workshop at NeurIPS in 2020, 2021, and 2022. She is also co-organizing several workshops at RSS this year. Masha has led efforts to increase the presence of women in academia. She served as Co-President of Women in Aeronautics and Astronautics at Stanford, mentored students in the AI4ALL program, and now serves as Co-Chair for the Women and Allies Employee Resource Group at TRI.

Advisors The advisors are past organizers and senior members of the community who help the organizing committee by providing feedback and proposing speakers.

- **Alex Bewley** (Google DeepMind, Swizerland)
- **Georgia Chalvatzaki** (TU Darmstadt, Germany)
- **Hamidreza Kasaei** (University of Groningen, Netherlands)
- **Jonathan Tompson** (Google DeepMind, USA)
- **Markus Wulfmeier** (Google DeepMind, UK)

REFERENCES

David B. D'Ambrosio, Saminda Abeyruwan, Laura Graesser, Atil Iscen, Heni Ben Amor, Alex Bewley, Barney J. Reed, Krista Reymann, Leila Takayama, Yuval Tassa, Krzysztof Choromanski, Erwin Coumans, Deepali Jain, Navdeep Jaitly, Natasha Jaques, Satoshi Kataoka, Yuheng Kuang, Nevena Lazic, Reza Mahjourian, Sherry Moore, Kenneth Oslund, Anish Shankar, Vikas Sindhwani, Vincent Vanhoucke, Grace Vesom, Peng Xu, and Pannag R. Sanketi. Achieving human level competitive robot table tennis. In *arXiv*, 2024. URL https://arxiv.org/abs/2408.03906.

Zipeng Fu, Qingqing Zhao, Qi Wu, Gordon Wetzstein, and Chelsea Finn. Humanplus: Humanoid shadowing and imitation from humans. In *arXiv*, 2024. URL https://arxiv.org/abs/2406.10454.

Elia Kaufmann, Leonard Bauersfeld, Antonio Loquercio, Matthias Müller, Vladlen Koltun, and Davide Scaramuzza. Champion-level drone racing using deep reinforcement learning. *Nature*, 620: 982–987, 2023. URL https://doi.org/10.1038/s41586-023-06419-4.

Tony Z. Zhao, Jonathan Tompson, Danny Driess, Pete Florence, Kamyar Ghasemipour, Chelsea Finn, and Ayzaan Wahid. Aloha unleashed: A simple recipe for robot dexterity. In *arXiv*, 2024. URL https://arxiv.org/abs/2410.13126.