

Extending robot minds through collective learning

Amanda Prorok

University of Cambridge, Cambridge, United Kingdom

Abstract

The current trend toward generalist robot behaviors with monolithic AI models is unsustainable. This article advocates for a paradigm shift that embraces distributed architectures for collective robotic intelligence. A modular ‘mixture-of-robots’ approach with specialized interdependent components can achieve superlinear gains, offering benefits in scalability, adaptability, and learning complex interactive skills.

Collective behaviors among robots are often treated as incidental rather than intrinsic to intelligence. Yet the classical approach to robot autonomy, which focuses on individual robots operating independently, is not well-suited for the complexities of real-world environments where interaction and collaboration are essential. Current state-of-the-art robot controllers rely on large, monolithic, centrally controlled models with millions or even billions of parameters. However, scaling laws demonstrate that achieving more sophisticated behaviors requires disproportionately high investments, making this approach to robot autonomy neither scalable nor sustainable—and ultimately, fundamentally misguided. Drawing on the principles of *active externalism* and the *extended mind*, this article argues for a paradigm shift towards designing robot collectives composed of diverse, specialized agents that function as part of larger systems.

Where does the mind end and the rest of the world begin? As roboticists, we are trained to abstract the environment as a distinct entity. For example, in navigation problems, the environment is typically modeled as a 2D or 3D map, with obstacles represented as static or dynamic regions to be avoided. Sensors such as laser rangefinders or cameras detect objects, which are then abstracted into a grid map, occupancy map, or a set of geometric shapes.

The historical roots of the endeavor to create abstractions date back to the 1956 Dartmouth workshop on AI, which sought, among other goals, to define formal models of the environment [1]. The workshop fostered a mindset that shaped the assumption that a robot’s cognitive processes are contained ‘within its own head’, requiring externalities to be either explicitly modeled or implicitly understood. Although undeniably *practical*, this abstraction is not necessarily *principled*. If we admit that external entities—such as bodies, morphologies, spaces, and other agents—play a notable role in cognitive processes, effectively functioning as a part of an extended mind, then the strict separation between robot mind and environment may no longer hold. Simple examples from our own daily lives illustrate this: we rely heavily on external aids to extend our cognitive capacity, in memory, computation, and experience—utilizing notebooks, calculators, and conversations with peers. Consider also examples from the animal kingdom that illustrate how organisms leverage their environment to survive and thrive, such as through stigmergy, tool use, indirect learning and exploratory processes [2, 3]. Research in the late 90s introduced the concept of *active externalism* [4], which explains how coupling with external elements enables us to overcome ‘on-board’ limitations by incorporating parts of the environment into the loop of cognitive processes. Recent popular literature has similarly argued that our thought processes are shaped by external resources, including the minds of others with whom we interact [5].


Intelligent beings do not learn in isolation. In humans, the physical limitations of individual brains have driven the evolution of capabilities that enable collective intelligence, as compellingly

demonstrated by neuroscience research [6]. Over generations, we have honed prosocial behaviors such as empathy, altruism, and effective communication. These behaviors stem largely from an intrinsic drive to reciprocate attention and the pleasure derived from communication and collaboration [7]. In addition, individuals are inherently vulnerable, and social interaction helps mitigate this vulnerability. When we raise our children, we do not teach them to be free of dependencies; instead, we guide them to function as part of groups—learning to cooperate and coordinate with others by assuming roles and responsibilities within communities. In essence, we do not prepare our offspring to survive through fully autonomous, dependency-free behavior. Research on the ‘social brain hypothesis’ highlights some physical consequences: data show a positive correlation between brain size and social group size, in both humans and primates [8]. Although still an active area of study, evidence suggests that social intelligence is linked to the ability to use basic decision-making functions in a more sophisticated way, due to the capacity to bring greater computing power to bear on the task at hand [7]. Does this not imply that inherently collectively intelligent agents possess superior cognitive abilities? Building on the idea of an extended mind, why not devise a new approach to understanding, modeling, and synthesizing robot intelligence — one that begins with the *collective*?


Robotics’ quest for *autonomy*, in its classical sense, is misguided. Over its operational lifetime, a robot will inevitably interact with other agents. Yet, to this day, collective behaviors among agents are often viewed as incidental rather than intrinsic to intelligence [9]. As it stands, artificially intelligent agents—embodied or virtual—are generally not designed from the ground up to interact with other intelligent agents, whether machines or humans. Current understanding and foundational AI textbooks still frame the canonical AI problem as that of a solitary machine confronting a non-social environment. Similarly, the classical definition of robot autonomy positions an autonomous robot as one designed to engage with its environment independently, free of any dependencies. Although research domains such as multi-agent systems [10] and swarm intelligence [11] have existed for over two decades, groups of autonomous agents still tend to be modeled and dealt with as a self-contained problem that is not core to solving autonomy as a whole.

Today, the typical AI product consists of a monolithic, centrally controlled model with millions or even billions of parameters and requiring large amounts of memory, even for inference. Loading large models can require hundreds of gigabytes of working memory, excluding any additional memory needed to store data. At present, such models need to be run off-board and at low inference speeds due to forward pass latency, precluding their use in high-frequency control, and thus excluding many robotic applications. Figure 1a shows the runtime capability of four large models (DINOv2), run on two alternate platforms: a Raspberry Pi 5—an ideal board for small form-factor robots that need to be energy efficient—and an Orin NX—a larger board that can be run on mid-size mobile robots. The plot shows that even when running on the more powerful board, only the smallest of the four models achieves update rates that could be considered ‘real-time’ capable. Despite steady progress in chip design for on-device inference, commercial production cycles typically last numerous years, obliging near-horizon applications to fall back on methods such as model distillation and quantization, which come with compromises on model performance.

Robotics research has followed the trend of integrating off-the-shelf, centralized and monolithic AI models into autonomy pipelines [14] (with the creation of large, open datasets further spurring this approach [15]). This decision-making design is fueling a push to build ‘generalist robots’ [16]. In the worst case, an exclusive focus on this approach would result in expensive ‘single-minded’ models that exhibit myopic behaviors and develop internal representations that are incompatible with other worldviews or incentives. But intelligence is not a single super-organism, nor can all conceivable tasks be accomplished by a single super-robot. Instead of targeting independent robots that can do it all and that are run by the same general ‘brain’, we should focus on designing collectives of interdependent, diverse, and specialized robots—both in mind and body—that function as part of larger systems. A collective robotic intelligence would, hence, entail a more modular and compositional approach to both hardware and software design, in which multiple models can learn and interact within and across physical robot embodiments. These models wouldn’t just be simply combined; instead, they would not only learn individual skills, but would also learn to interact and ‘collaborate’ with other



(a)



(b)

Figure 1: (a) Inference frequency by model and board for DINOv2 [12], on two platforms: the Raspberry Pi 5 (8GB RAM) and the NVIDIA Orin NX (16GB RAM). DINOv2 includes four model variants: small (22.1M parameters), base (86.6M parameters), large (304.4M parameters), and huge (1136.5M parameters), for which the official float32 model weights (no quantization) are applied, ensuring consistent model parameters across both devices. Inference was performed using PyTorch in CPU mode on the Raspberry Pi and GPU mode on the Orin NX. Each data point represents the average performance over 50 inference runs. (b) A linear fit on model data from [13] shows a 4.68 year-on-year growth rate with respect to computation (measured in FLOPs) needed for training. Robotics-related models are labeled and shown in boxes.




Figure 2: Illustration of the scaling law for a collective of robot models. A ‘mixture-of-robots’ approach achieves super-linear performance gains with an increasing number of combinations of narrower ‘expert’ robot behaviors. A traditional approach with monolithic ‘do-it-all’ models is bounded by a power law.

components and skills from the outset, leading to a more defined understanding of individual and collective capabilities.

What can we gain from collective robotic intelligence? First, consider the cost. Scaling laws from deep learning theory reveal that model size and data requirements scale exponentially in relation to model performance [17]. In other words, our investments in training—whether in energy, data, or time—must increase disproportionately as we aim for more sophisticated and general behaviors. This unsustainable trend is evident in the nearly 5x year-on-year increase in computing needs (see Figure 1b), which far outpaces Moore’s Law’s approximate 1.4x annual growth rate in transistor density. In contrast, a collective of agents employing narrow, specialized models could yield superlinear improvements in compositional skills and generalization—the ability to combine individual capabilities in novel, untrained ways. This modular, compositional approach to autonomy would enable our agents and robots to effectively emulate the behavior of a vast neural network, dynamically reconfiguring at runtime to suit the task. This could be achieved by employing only a fraction of model parameters when compared to monolithic homogeneous models [18]. Although some fundamental skills (e.g., ‘base’ behaviors) might be shared, individual robots would typically possess distinct model compositions, leveraging team-building and interaction strategies to exploit complementary behaviors, reflecting an *active externalism* approach. The performance of any generated compositional behavior on a specific task can then be used as feedback, informing and refining future compositional processes. The desired result is a *combinatorial expansion of the skill space* (as depicted in Figure 2). By training the specialist models to interact from the ground up, the robot collective would also become more resilient to unexpected disruptions and out-of-distribution events through dynamic reorganization of responsibilities and roles [19]. Recent research has demonstrated that encouraging diverse behaviors during the learning process also fosters the emergence of superior team strategies that give rise to collective intelligence that is not simply dependent on the maximum or sum intelligence of team members [20].

Secondly, there are skills that can only be acquired through social learning in collectives. Building upon the idea of training interaction from the ground up, a collaborative learning process also allows individual specialist models to develop a better intrinsic understanding of their own capabilities and

limitations. This awareness naturally fosters an understanding of when and where teaming up with other models becomes necessary to overcome individual shortcomings and achieve more complex goals. Moreover, skills such as ‘theory of mind’—the ability to represent and understand the mental states of others—and meta-cognition—the ability to reflect on and justify our behavior to others—are not innate, but must be learned in collective settings. These skills are essential for tasks involving interaction with humans or other agents and robots, where behavioral uniformity (policy homogeneity) cannot be assumed. The default state of the world is one in which heterogeneous agents coexist, making such skills essential for robust and adaptive robot operation. Additionally, the physical and morphological limitations of individual robots often necessitate collaboration across diverse embodiments to address complex tasks. For instance, collaborative object transport may require robots with varying sizes or strengths, whereas mapping and surveillance of intricate terrains might demand a combination of aerial and ground-based systems to achieve comprehensive coverage and efficiency.

Thirdly, collective learning stands to improve the performance of a single agent, even when tasked with a solitary problem, by boosting the rate of learning and improving long-term outcomes. These benefits arise from sharing experience and knowledge acquired by others, leading to substantial cost and time savings: in robotics, where execution is physical (rather than virtual), collecting experiential data is expensive, risky, and even dangerous. Leveraging *others’* experiences catalyzes the learning process [21], and crucially, prevents the repetition of negative or dangerous actions. Shared knowledge can also mitigate challenges like catastrophic forgetting by outsourcing and distributing certain learned skills, much akin to the distribution of wisdom and knowledge in human communities [22]. Furthermore, parallel policy executions allow us to collapse the search space, which is particularly beneficial in high-dimensional observation and action spaces. Techniques to achieve faster and more exhaustive learning for robotics are being advanced in fields such as distributed learning [23], federated learning [24], and multi-agent learning [25].

Towards future-proof robot collectives. Further research is needed to develop new paradigms for generating and training robot collectives and durable collective learning algorithms, resulting in a diversity of smaller models that are specialized yet adaptive, capable of running on-board, and facilitated by a universal (cross-embodiment) communication policy. Current methods for creating such collectives often fall short in scalability and adaptability. Although fixed monolithic model designs are relatively straightforward to come by, collective intelligence does not follow a centralized paradigm, nor is it rigid. Instead, it emerges from interactions across spatial and temporal scales. Nature provides compelling examples of this phenomenon: collections of genes cooperate to produce organisms, cells cooperate to produce multi-cellular organisms, and multi-cellular animals cooperate to form complex social groups [26]. In order to engineer an analogous outcome for groups of robots, we need to devise interaction methods—both explicit and implicit—to facilitate deployment and foster sustainable and durable collective learning. This includes considering whether training occurs primarily in simulation or directly on physical hardware, and exploring various methodologies such as model-based planning and control, reinforcement learning, or hybrid approaches.

Yet we are still lacking robust techniques to support this vision. For instance, it is not clear how we would achieve the combinations of models referenced in Fig. 2. A key challenge lies in establishing how individual robots learn to effectively collaborate and understand the limits of their own capabilities, necessitating teaming up with others. Explicit local interaction among robots commonly relies on narrow-band communication networks. But the problem of devising methods that help robots decide ‘what, when, and with whom to communicate’ remains an open one, due to the difficulty of hand-designing effective communication strategies [27]. This is further complicated by the need to make assumptions about communication constraints and latencies within the collective. A recent line of work has demonstrated how this challenge can be addressed through automatic synthesis of inter-robot messages via differentiable communications channels—data-driven communication strategies as seen in [28, 29]. While promising, this line of research is still in its early stages, with ongoing efforts focused on developing suitable inductive biases and architectures, such as graph neural networks, deep sets, and cellular neural networks. More research is needed to understand the trade-offs and limitations of these emerging methods in the context of generating robust and adaptable robot collectives.

A mixture-of-robots approach. We need to relax our propensity towards centralized, monolithic models in robot learning, resisting the convenience of architectural and infrastructural simplicity, and the availability of large, centralized datasets. This does not rule out the development of reusable, multi-purpose models—on the contrary, positive transfer remains a desirable property that helps offset training cost [30]. Ultimately, the focus should not be on how large a robotics foundation model can be, but rather on ensuring that future models are capable of inter-robot interaction and inherently equipped for continuous improvement and refinement through these interactions. Unlike standard mixture-of-experts approaches, which primarily focus on combining sub-networks within a single model, the ‘mixture-of-robots’ paradigm encompasses not only the internal composition of more fundamental behaviors within individual robots, but also the dynamic assembly and interaction of distinct, physically embodied agents, each with potentially specialized model compositions.

In spite of rapid progress in the field, the concept of a ‘mixture-of-robots’ paradigm remains poorly defined. It is challenging to devise models that address distinct and potentially non-overlapping action domains, such as those involving different robot morphologies and actuators [30]. In a ‘mixture-of-robots’ approach, this challenge would be relegated to an algorithm or mechanism that connects the various models. Developing such a methodology is an open problem. One potential strategy could be to build a library of compositional, agent-specific foundation models alongside an API capable of ‘combining’ them. Initial forays could build on work on ensemble models, mixture-of-experts, hypernetworks, and hierarchical learning, but with the ultimate aim of assembling *robot collectives* as a function of underlying tasks—whether in offline or online settings. However, substantial gaps remain, particularly in achieving real-time compositionality of specialized robot skills and behaviors. Similarly, research on heterogeneity in robot systems is still in its infancy.

How will we measure success? Our current definition of performance is rudimentary—often reduced to a proxy for learning loss or an optimization objective for individual robot models, without considering more nuanced collective metrics that include team-level satisfiability and collective resilience as sub-objectives [19]. Goodhart’s law compounds this problem (the economist Goodhart is credited with the statement: “*When a measure becomes a target, it ceases to be a good measure*”). For instance, consider a team of robots tasked with collaborative object manipulation: each robot arm is trained to maximize its individual grasp success rate—whether the object is lifted off the surface or not. In unforeseen collective behavior, the team might begin to focus on scenarios involving only easily graspable objects or rely heavily on the most capable robot, neglecting the potential for more complex and transferable strategies involving the entire collective. The performance metric does not account for task diversity or the team’s ability to adapt to a wide range of collaborative scenarios, object distributions and object types. This results in a system that performs well in controlled individual robot tests but fails in environments requiring flexible team coordination or the precise allocation of roles. To address this, we might consider incorporating more comprehensive evaluation protocols that go beyond simple individual success rates. This could involve defining more holistic benchmarks with a diverse set of collective tasks involving coordinated navigation, distributed sensing, and multi-robot manipulation, potentially guided by domain experts to ensure real-world relevance. This could be accompanied by a more standardized approach towards automatically generating diverse and challenging training scenarios, perhaps through adversarial techniques or curriculum learning strategies that progressively increase task complexity and the need for inter-robot cooperation. Furthermore, incorporating metrics that explicitly reward collective robustness (e.g., tolerance to individual robot failures), collective adaptability (e.g., dynamic role assignment), and the ability to handle a wide range of team-based conditions, rather than just optimizing for individual robot performance and narrow objectives, could mitigate the effects of Goodhart’s law. Crucially, our evaluation of success must also extend beyond performance on specific benchmarks to consider the longevity of robot deployments in real-world settings and their potential for commercial viability. This necessitates evaluating factors such as the reliability of the system over extended periods, the ease of maintaining and repairing the entire team, user satisfaction with the collective’s output, and the ability of the collective to adapt to evolving user needs and environmental changes. Metrics related to the total cost of robot ownership, including downtime of the system and the need for frequent retraining or upgrades of the team’s skills,

will also be critical indicators of true success in the long run.

Defining learning objectives. Continual learning frameworks must also incentivize interactions and reward shared intentionality—i.e., the ability to share attention and intention—and ensure alignment with human-designed higher-order objectives. Defining a singular, global objective function that effectively captures the desired collective behavior and accounts for the contributions of individual agents is substantially more complex than in single-agent systems, due to the emergent nature of group interactions and the difficulty in attributing collective outcomes to specific actions. To achieve this, we need methods that are capable of translating such objectives into global reward signals that adjust existing base behaviors. These methods should also incorporate credit assignment techniques to ensure sample efficiency across the robot collectives. In parallel, we would need to further our research on methods that robustify robot-to-robot interactions and connectivity. Adversarial manipulation, communication failures, and faulty robots may compromise the techniques. Similarly, challenges such as rumor propagation and the difficulty of ‘unlearning’ undesirable or misaligned behaviors within robot collectives require careful consideration. Privacy and accountability during learning also remain critical concerns, as we have only just begun to explore the potential threats. Addressing these challenges requires the concurrent development of robust networking infrastructure [27] and on-board AI acceleration to support these systems effectively.

Concluding thoughts. Robotics research is navigating a delicate balance between the transformative potential of machine learning and the unique complexities inherent to robotics. Although harnessing advances in AI is undoubtedly crucial, real progress requires resisting the allure of short-term gains that sidestep deeper foundational challenges. In particular, we must avoid relying exclusively on large, monolithic models that circumvent systems issues rather than addressing them directly. Instead, collaborative and collective intelligence should form the cornerstone of robotics research—not only because these approaches offer scalable and sustainable solutions, but also because the inherently interactive, distributed, and collective nature of real-world problems demands it. Moreover, this research agenda holds the promise of yielding theoretical advances with implications that extend far beyond robotics, fostering innovation in fields ranging from distributed systems to artificial intelligence at large. By embracing this vision, we can ensure that robotics develops as a discipline capable of addressing the complex, interconnected challenges of the present and future.

Acknowledgements

I am very grateful to Carl Henrik Ek and Sven Goyal for their thoughtful comments on draft versions of this article. My thanks also go to Heedo Woo for collecting the data for the DINOv2 runtime plot. I am indebted to Oliver Brock, Gaurav Sukhatme, and Matt Mason for lively discussions on the science of robotics, which seeded some of the thoughts presented herein. Finally, I thank the anonymous reviewers for their constructive feedback on this manuscript. A Prorok was supported in part by the European Research Council (ERC) Project 949940 (gAIa).

References

- [1] John McCarthy, Marvin L Minsky, Nathaniel Rochester, and Claude E Shannon. A proposal for the Dartmouth summer research project on artificial intelligence, august 31, 1955. *AI magazine*, 27(4):12–12, 2006.
- [2] Marco Dorigo, Eric Bonabeau, and Guy Theraulaz. Ant algorithms and stigmergy. *Future generation computer systems*, 16(8):851–871, 2000.
- [3] Jeannette Bohg, Karol Hausman, Bharath Sankaran, Oliver Brock, Danica Kragic, Stefan Schaal, and Gaurav S Sukhatme. Interactive perception: Leveraging action in perception and perception in action. *IEEE Transactions on Robotics*, 33(6):1273–1291, 2017.
- [4] Andy Clark and David Chalmers. The extended mind. *Analysis*, 58(1):7–19, 1998.
- [5] Annie Murphy Paul. *The extended mind: The power of thinking outside the brain*. Eamon Dolan Books, 2021.
- [6] Hannah Critchlow. *Joined-up thinking: the science of collective intelligence and its power to change our lives*. Hachette UK, 2022.
- [7] Robin IM Dunbar and Susanne Shultz. Evolution in the social brain. *science*, 317(5843):1344–1347, 2007.

- [8] Simon M Reader and Kevin N Laland. Social intelligence, innovation, and enhanced brain size in primates. *Proceedings of the National Academy of Sciences*, 99(7):4436–4441, 2002.
- [9] Guillaume Dumas, Julien Laroche, and Alexandre Lehmann. Your body, my body, our coupling moves our bodies. *Frontiers in Human Neuroscience*, 8:1004, 2014.
- [10] Michael Wooldridge. *An Introduction to Multi-Agent Systems*. John Wiley and Sons Limited, 2002.
- [11] Eric Bonabeau, Marco Dorigo, and Guy Theraulaz. *Swarm Intelligence from Natural to Artificial Systems*. Santa Fe Institute Studies in the Sciences of Complexity. Oxford University Press, 1999.
- [12] Maxime Oquab, Timothée Darcet, Theo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Russell Howes, Po-Yao Huang, Hu Xu, Vasu Sharma, Shang-Wen Li, Wojciech Galuba, Mike Rabbat, Mido Assran, Nicolas Ballas, Gabriel Synnaeve, Ishan Misra, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. Dinov2: Learning robust visual features without supervision, 2023.
- [13] Epoch AI. Data on AI models, 2025. Available online at <https://epoch.ai/data/ai-models>. Accessed: 2025-07-28.
- [14] Danny Driess, Fei Xia, Mehdi S. M. Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, Wenlong Huang, Yevgen Chebotar, Pierre Sermanet, Daniel Duckworth, Sergey Levine, Vincent Vanhoucke, Karol Hausman, Marc Toussaint, Klaus Greff, Andy Zeng, Igor Mordatch, and Pete Florence. PaLM-E: An Embodied Multimodal Language Model, March 2023. arXiv:2303.03378 [cs].
- [15] Open X-Embodiment Collaboration, Abby O’Neill, Abdul Rehman, Abhinav Gupta, Abhiram Maddukuri, Abhishek Gupta, Abhishek Padalkar, Abraham Lee, Acorn Pooley, Agrim Gupta, Ajay Mandlekar, Ajinkya Jain, Albert Tung, Alex Bewley, Alex Herzog, Alex Irpan, Alexander Khazatsky, Anant Rai, Anchit Gupta, Andrew Wang, Andrey Kolobov, Anikait Singh, Animesh Garg, Aniruddha Kembhavi, Annie Xie, Anthony Brohan, Antonin Raffin, Archit Sharma, Arefeh Yavary, Arhan Jain, Ashwin Balakrishna, Ayzaan Wahid, Ben Burgess-Limerick, Beomjoon Kim, Bernhard Schölkopf, Blake Wulfe, Brian Ichter, Cewu Lu, Charles Xu, Charlotte Le, Chelsea Finn, Chen Wang, Chenfeng Xu, Cheng Chi, Chenguang Huang, Christine Chan, Christopher Agia, Chuer Pan, Chuyuan Fu, Coline Devin, Danfei Xu, Daniel Morton, Danny Driess, Daphne Chen, Deepak Pathak, Dhruv Shah, Dieter Büchler, Dinesh Jayaraman, Dmitry Kalashnikov, Dorsa Sadigh, Edward Johns, Ethan Foster, Fangchen Liu, Federico Ceola, Fei Xia, Feiyu Zhao, Felipe Vieira Frujeri, Freek Stulp, Gaoyue Zhou, Gaurav S. Sukhatme, Gautam Salhotra, Ge Yan, Gilbert Feng, Giulio Schiavi, Glen Berseth, Gregory Kahn, Guangwen Yang, Guanzhi Wang, Hao Su, Hao-Shu Fang, Haochen Shi, Henghui Bao, Heni Ben Amor, Henrik I Christensen, Hiroki Furuta, Homanga Bharadhwaj, Homer Walke, Hongjie Fang, Huy Ha, Igor Mordatch, Ilija Radosavovic, Isabel Leal, Jacky Liang, Jad Abou-Chakra, Jaehyung Kim, Jaimyn Drake, Jan Peters, Jan Schneider, Jasmine Hsu, Jay Vakil, Jeannette Bohg, Jeffrey Bingham, Jeffrey Wu, Jensen Gao, Jiaheng Hu, Jiajun Wu, Jialin Wu, Jiankai Sun, Jianlan Luo, Jiayuan Gu, Jie Tan, Jihoon Oh, Jimmy Wu, Jingpei Lu, Jingyun Yang, Jitendra Malik, João Silvério, Joey Hejna, Jonathan Booher, Jonathan Tompson, Jonathan Yang, Jordi Salvador, Joseph J. Lim, Junhyek Han, Kaiyuan Wang, Kanishka Rao, Karl Pertsch, Karol Hausman, Keegan Go, Keerthana Gopalakrishnan, Ken Goldberg, Kendra Byrne, Kenneth Oslund, Kento Kawaharazuka, Kevin Black, Kevin Lin, Kevin Zhang, Kiana Ehsani, Kiran Lekala, Kirsty Ellis, Krishan Rana, Krishnan Srinivasan, Kuan Fang, Kunal Pratap Singh, Kuo-Hao Zeng, Kyle Hatch, Kyle Hsu, Laurent Itti, Lawrence Yunliang Chen, Lerrel Pinto, Li Fei-Fei, Liam Tan, Linxi ”Jim” Fan, Lionel Ott, Lisa Lee, Luca Weihs, Magnum Chen, Marion Lepert, Marius Memmel, Masayoshi Tomizuka, Masha Itkina, Mateo Guaman Castro, Max Spero, Maximilian Du, Michael Ahn, Michael C. Yip, Mingtong Zhang, Mingyu Ding, Minh Heo, Mohan Kumar Srirama, Mohit Sharma, Moo Jin Kim, Naoaki Kanazawa, Nicklas Hansen, Nicolas Heess, Nikhil J Joshi, Niko Suenderhauf, Ning Liu, Norman Di Palo, Nur Muhammad Mahi Shafiullah, Oier Mees, Oliver Kroemer, Osbert Bastani, Pannag R Sanketi, Patrick ”Tree” Miller, Patrick Yin, Paul Wohlhart, Peng Xu, Peter David Fagan, Peter Mitrano, Pierre Sermanet, Pieter Abbeel, Priya Sundareshan, Qiuyu Chen, Quan Vuong, Rafael Rafailov, Ran Tian, Ria Doshi, Roberto Mart’in-Mart’in, Rohan Bajjal, Rosario Scalise, Rose Hendrix, Roy Lin, Runjia Qian, Ruohan Zhang, Russell Mendonca, Rutav Shah, Ryan Hoque, Ryan Julian, Samuel Bustamante, Sean Kirmani, Sergey Levine, Shan Lin, Sherry Moore, Shikhar Bahl, Shiva Dassy, Shubham Sonawani, Shubham Tulsiani, Shuran Song, Sichun Xu, Siddhant Haldar, Siddharth Karamcheti, Simeon Adebola, Simon Guist, Soroush Nasiriany, Stefan Schaal, Stefan Welker, Stephen Tian, Subramanian Ramamoorthy, Sudeep Dasari, Suneel Belkhale, Sungjae Park, Suraj Nair, Suvir Mirchandani, Takayuki Osa, Tanmay Gupta, Tatsuya Harada, Tatsuya Matsushima, Ted Xiao, Thomas Kollar, Tianhe Yu, Tianli Ding, Todor Davchev, Tony Z. Zhao, Travis Armstrong, Trevor Darrell, Trinity Chung, Vidhi Jain, Vikash Kumar, Vincent Vanhoucke, Wei Zhan, Wenxuan Zhou, Wolfram Burgard, Xi Chen, Xiangyu Chen, Xiaolong Wang, Xinghao Zhu, Xinyang Geng, Xiyuan Liu, Xu Liangwei, Xuanlin Li, Yansong Pang, Yao Lu, Yecheng Jason Ma, Yejin Kim, Yevgen Chebotar, Yifan Zhou, Yifeng Zhu, Yilin Wu, Ying Xu, Yixuan Wang, Yonatan Bisk, Yongqiang Dou, Yoonyoung Cho, Youngwoon Lee, Yuchen Cui, Yue Cao, Yueh-Hua Wu, Yujin Tang, Yuke Zhu, Yunchu Zhang, Yunfan Jiang, Yunshuang Li, Yunzhu Li, Yusuke Iwasawa, Yutaka Matsuo, Zehan Ma, Zhuo Xu, Zichen Jeff Cui, Zichen Zhang, Zipeng Fu, and Zipeng Lin. Open X-Embodiment: Robotic learning datasets and RT-X models. <https://arxiv.org/abs/2310.08864>, 2023.
- [16] Sergey Levine and Karol Hausman. The global project to make a general robotic brain, January 2024. Available online at <https://spectrum.ieee.org/global-robotic-brain>. Accessed 2025-08-27.
- [17] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.

- [18] Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarczyk, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*, 2017.
- [19] Amanda Prorok, Matthew Malencia, Luca Carlone, Gaurav S Sukhatme, Brian M Sadler, and Vijay Kumar. Beyond robustness: A taxonomy of approaches towards resilient multi-robot systems. *arXiv preprint arXiv:2109.12343*, 2021.
- [20] Matteo Bettini, Ryan Kortvelesy, and Amanda Prorok. Controlling behavioral diversity in multi-agent reinforcement learning. *International Conference on Machine Learning (ICML)*, 2024.
- [21] Sudeep Dasari, Frederik Ebert, Stephen Tian, Suraj Nair, Bernadette Bucher, Karl Schmeckpeper, Siddharth Singh, Sergey Levine, and Chelsea Finn. Robonet: Large-scale multi-robot learning. *arXiv preprint arXiv:1910.11215*, 2019.
- [22] Amanda J Barnier, John Sutton, Celia B Harris, and Robert A Wilson. A conceptual and empirical framework for the social distribution of cognition: The case of memory. *Cognitive Systems Research*, 9(1-2):33–51, 2008.
- [23] Javier Yu, Joseph A Vincent, and Mac Schwager. Dinno: Distributed neural network optimization for multi-robot collaborative learning. *IEEE Robotics and Automation Letters*, 7(2):1896–1903, 2022.
- [24] Stefano Savazzi, Monica Nicoli, Mehdi Bennis, Sanaz Kianoush, and Luca Barbieri. Opportunities of federated learning in connected, cooperative, and automated industrial systems. *IEEE Communications Magazine*, 59(2):16–21, 2021.
- [25] Stefano V. Albrecht, Filippas Christianos, and Lukas Schäfer. *Multi-Agent Reinforcement Learning: Foundations and Modern Approaches*. MIT Press, 2024.
- [26] Simon Levin. Crossing scales, crossing disciplines: collective motion and collective action in the global commons. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1537):13–18, 2010.
- [27] Jennifer Gielis, Ajay Shankar, and Amanda Prorok. A critical review of communications in multi-robot systems. *Current robotics reports*, 3(4):213–225, 2022.
- [28] Alvaro Serra-Gómez, Bruno Brito, Hai Zhu, Jen Jen Chung, and Javier Alonso-Mora. With whom to communicate: learning efficient communication for multi-robot collision avoidance. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 11770–11776. IEEE, 2020.
- [29] Qingbiao Li, Fernando Gama, Alejandro Ribeiro, and Amanda Prorok. Graph neural networks for decentralized multi-robot path planning. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 11785–11792. IEEE, 2020.
- [30] Ria Doshi, Homer Walke, Oier Mees, Sudeep Dasari, and Sergey Levine. Scaling cross-embodied learning: One policy for manipulation, navigation, locomotion and aviation. *Conference on Robot Learning (CoRL)*, 2024.