ARTIFICIAL INTELLIGENCE

From autonomy to alliance: Robotic foundation models must learn with us, not just for us

Sharmita Dey[1]*, Robert Riener[1], Strahinja Dosen[2], Stefano V. Albrecht[3]

This Viewpoint urges reimagining of robotic foundation models, from treating the robot as a solitary, omnipotent agent to embracing a multiagent, alliance-aware paradigm. Alliance-aware models learn with humans and other robots, not merely for them, by embedding mechanisms that foster social interaction and generalization across heterogeneous partners. We outline six design pillars that cultivate such collaborative intelligence: interaction priors, partner modeling (machine theory of mind), modular and composable policies, norm adaptation, trust-aware memory, and communication. Together, these pillars empower robots to fluidly switch social roles, adapt to unfamiliar collaborators, and coordinate robustly within dynamic multiagent ecologies spanning homes, factories, clinics, and field operations.

INTRODUCTION

Robotic foundation models, large-scale pretrained models adaptable to many tasks, are poised to redefine autonomy in the physical world. Models such as robotics transformer (RT-2), Gato, and Octo have demonstrated broad generalization in perception, decision-making, and control by leveraging internet-scale data and transfer learning (*1*–*3*). However, most robotic foundation-model designs still cast the robot as a lone, hypercompetent agent that intermittently receives high-level commands from a static world. The assumption is that if a model transfers across many tasks, it will be broadly useful.

We argue that this view is too narrow. A vast range of real-world robotics scenarios are inherently interactive and multiagent in nature. Consider rehabilitation robots that support patient-cooperative training, semiautonomous cars that arbitrate shared control with human drivers, and warehouse collaborative robots that negotiate etiquette with human co-workers and other environmental agents. In all of these settings, the robot is one node in a tightly coupled human-machine web whose interactions evolve dynamically. This reframes learning in robotics from autonomy to alliance.

To thrive on this interactive frontier, robotic foundation models must learn with us, not merely for us. We argue that the next generation of foundation models in robotics should be explicitly designed for alliance: collaborative, multiagent intelligence in which robots and humans learn and adapt together as partners. We propose an "ecological" perspective on robotic learning and transfer. By ecology, we mean an open-ended environment with multiple agents (robots and humans) interacting and coevolving. In such settings, other decision-makers' goals, behaviors, and norms may change over time. Robotic foundation models should be designed to thrive in these multiagent ecologies, adapting to new partners and team structures. Beyond task and embodiment transfer, we highlight the need for social transfer: generalizing across new partners, roles, and social interactions.

This shift is not just conceptual but practical. Robots in homes, clinics, factories, or disaster sites will encounter unfamiliar collaborators and must join ad hoc teams. Humans can coordinate with new teammates by using social common sense, a capability current robot models lack. Bridging this gap requires algorithms and models for interactive learning, partner modeling, and role flexibility. It also requires new success metrics: A robot should be judged not only on solo task performance but also on how robustly it adapts in human-robot and robot-robot teams.

In this Viewpoint, we focus on the computational and learning dimensions of alliance-aware robotic foundation models. We scope our discussion to these algorithmic aspects of collaborative intelligence, emphasizing how robots can learn and adapt with humans and other robots. However, achieving truly functional collaborative robots will also require progress in other interdependent domains, beyond our present remit. For example, factors such as energy-efficient compliant actuators (*4*), high-fidelity tactile (*5*) and multimodal sensing, and embodied cognition (how a robot's physical form and sensorimotor experience shape its behavior) (*6*) are all crucial to the success of human-robot alliances. Although these areas lie outside the focus of this article, we acknowledge their importance and briefly consider their implications later in the "Other Considerations" section.

TRANSFER ACROSS AGENTS: A NEW FRONTIER IN ALLIANCE-AWARE GENERALIZATION

Achieving the above vision requires rethinking how robots learn and generalize. The key insight is to enable transfer of knowledge across agents, essentially allowing a robot to learn not only from its own trials, where it performs tasks alone, but also through its interactions with humans or other robots. This adds a new dimension to generalization. Traditionally, we judge a robot model by how well it transfers skills to new tasks or new domains (for example, adapting a policy from simulation to a physical robot). Now, we must also consider transfer from shared knowledge, policies, and partners: Can insights gained by one agent, human, robot, or other agents in an environment be absorbed or complemented by others so that they coordinate better together? We argue that future robotic foundation models should include this agent-axis transfer as a fundamental goal, alongside task and domain transfer, to achieve true interactive intelligence.

Implementing transfer across agents entails changes in both how we train models and what capabilities those models have. On the training side, robots will need exposure to diverse partners and multiagent experiences

[1]Institute of Robotics and Intelligent Systems (IRIS) - Sensory-Motor Systems Lab, ETH Zurich, Zurich, Switzerland. [2]Department of Health Science and Technology, Aalborg University, Aalborg, Denmark. [3]NTU Singapore, Singapore, Singapore.
*Corresponding author. Email: dey.sharmita01@gmail.com

so that they learn general patterns of interaction, a kind of social curriculum. For example, a robot trained on data from many different partner behaviors (different human personalities or different robot teammates) can learn to generalize its behavior to accommodate a brand-new partner in a zero-shot manner. Indeed, recent benchmarks like zero-shot coordination (ZSC)-Eval (*7*) and best-response diversity objectives (*5*) systematically generate a spectrum of partner policies to evaluate an agent's adaptability to unseen collaborators. However, simply training on diverse multiagent data is not enough; the model must also be equipped with the right architectural mechanisms and skills to adapt quickly (few-shot transfer with minimal experience of the new partner) (*8*, *9*). In practice, this means building certain principles or inductive biases into the robot's artificial intelligence (AI), enabling it to understand other agents, share knowledge, adapt roles, and communicate. Below, we outline several key design pillars to achieve scalable transfer across agents. These pillars work in tandem to "activate" the robot's ability to learn with others, allowing a foundation model to recombine its knowledge across tasks, partners, and social contexts.

DESIGNING FOUNDATION MODELS FOR ALLIANCE-AWARE ADAPTIVE INTELLIGENCE

What would it take for a robotic foundation model to achieve the kind of social compositionality and role-flexible transfer described in the previous section? In other words, how can we design and build learning systems that readily recompose their behaviors to interact in new social contexts and team configurations? We suggest that several design pillars are needed to complement the usual ingredients of large-scale robot learning. These pillars (Fig. 1), described in the following section, work in tandem to enable a robot to fluidly adapt in a multiagent interaction to transfer not only skills across tasks but also its learned knowledge across social roles, partners, and group structures.

Learned interaction priors

Adapting on the fly to new partners also calls for robots to have strong priors about interaction (*10*), in essence, some built-in expectations of how agents typically behave and how multiagent cooperation works. Humans bring rich priors into any joint activity (gleaned from lifelong social experience), which is why we can coordinate, even with strangers, relatively effectively. Robots may need analogous interaction priors derived from large-scale data or simulations of multiagent behavior. One example is to train foundation models on multiagent trajectories (real or synthetic) so that they implicitly learn patterns of coordination (for example, if one agent moves to pick up an object, others often make room). Such learned priors could help a robot predict and react to partners' actions in ways that smooth coordination.

Partner modeling

A fundamental step toward agent transfer is enabling robots to model their partners (*11*), to internally represent what another agent might do, want, or believe. In humans, this faculty is often attributed to a theory of mind (ToM), the ability to attribute mental states to others (*12*). Equipping robot foundation models with analogous capabilities would allow them to anticipate and adapt to new agents. A growing literature frames this capability in terms of a machine ToM (*13*). Neural ToMnets treat a partner's objective as a latent variable: After observing a short behavior window, they output a vector embedding that both predicts the partner's next action and conditions the robot's own policy, enabling on-the-fly adaptation (*13*). When such ToM-style inference heads are integrated into a large foundational transformer pretrained on thousands of multiagent trajectories, the network can discover that certain short behavior snippets (for example, rapid direction changes and frequent gaze shifts) consistently precede particular future actions; it compresses those regularities into embeddings that capture intuitive social traits such as "explores versus follows routine" or "assertive versus timid." At run time, the robot can feed a few recent observations of an unfamiliar partner into the inference head, receive the corresponding trait vector, and condition its motion planner on that vector.

Another approach for generalizing to new partners is through ad hoc teamwork in which an autonomous agent must cooperate immediately with teammates whose policies, conventions, and even action spaces were not shared during training (*14*). Because success is evaluated from the first joint episode, algorithms that perform well in this setting are forced to discover partner-agnostic coordination strategies and to update internal partner models online, a capability directly aligned with transfer across agents. Recent benchmarks introduce dynamic team composition: The N-Agent Ad hoc Teamwork suite replaces or adds teammates mid-task, and policies that infer a teammate embedding in real time recover team performance more rapidly than baselines without such modeling (*15*). Training foundation models under these ad hoc protocols teaches the network a rich prior over interaction patterns. In effect, ad hoc teaming converts the abstract goal of "working with anyone" into a concrete learning signal, pushing robotic policies toward the rapid inference and adaptation abilities that real-world collaborative intelligence demands.

Human-robot interaction (HRI) often demands an additional capability: reasoning about what the human believes concerning the robot itself. By maintaining a model of the human's belief state, the robot can decide when to supply clarifying information, for example, verbally announcing its grasp target when it infers that the user holds an incorrect assumption, thereby improving trust and task efficiency (*16*). Equipping robots with such social reasoning and planning can substantially improve interagent adaptation, especially in HRI (*17*), and convert an unfamiliar collaborator into a predictable element of the control loop by enabling the rapid adaptation that is essential for agent-axis generalization.

Policy modularization and compositionality

A promising route to adaptable teamwork is to factor a robot's policy representations into reusable submodules that can be recombined on demand. One module might encode a task skill, for example, "insert peg," whereas another captures an interaction skill such as "follow a partner's lead." If these pieces are learned and stored separately, the robot can meet a novel situation by swapping in the required task module and the appropriate social module rather than retraining an entirely new network. Devin *et al.* (*18*) provided early evidence for this strategy: They decomposed neural policies into "task-specific" and "robot-specific" blocks that could be mixed and matched to solve previously unseen robot-task pairs without additional learning. Extending the idea, one could train distinct role modules, for instance, leader and follower behaviors, and activate the relevant component according to the robot's assignment in a team. Designing a foundation model with such composable behavior primitives would enable far more flexible transfer not only across tasks and embodiments but also across partners. Rather than learning a separate policy for every possible team scenario, the robot learns a library of behavior pieces

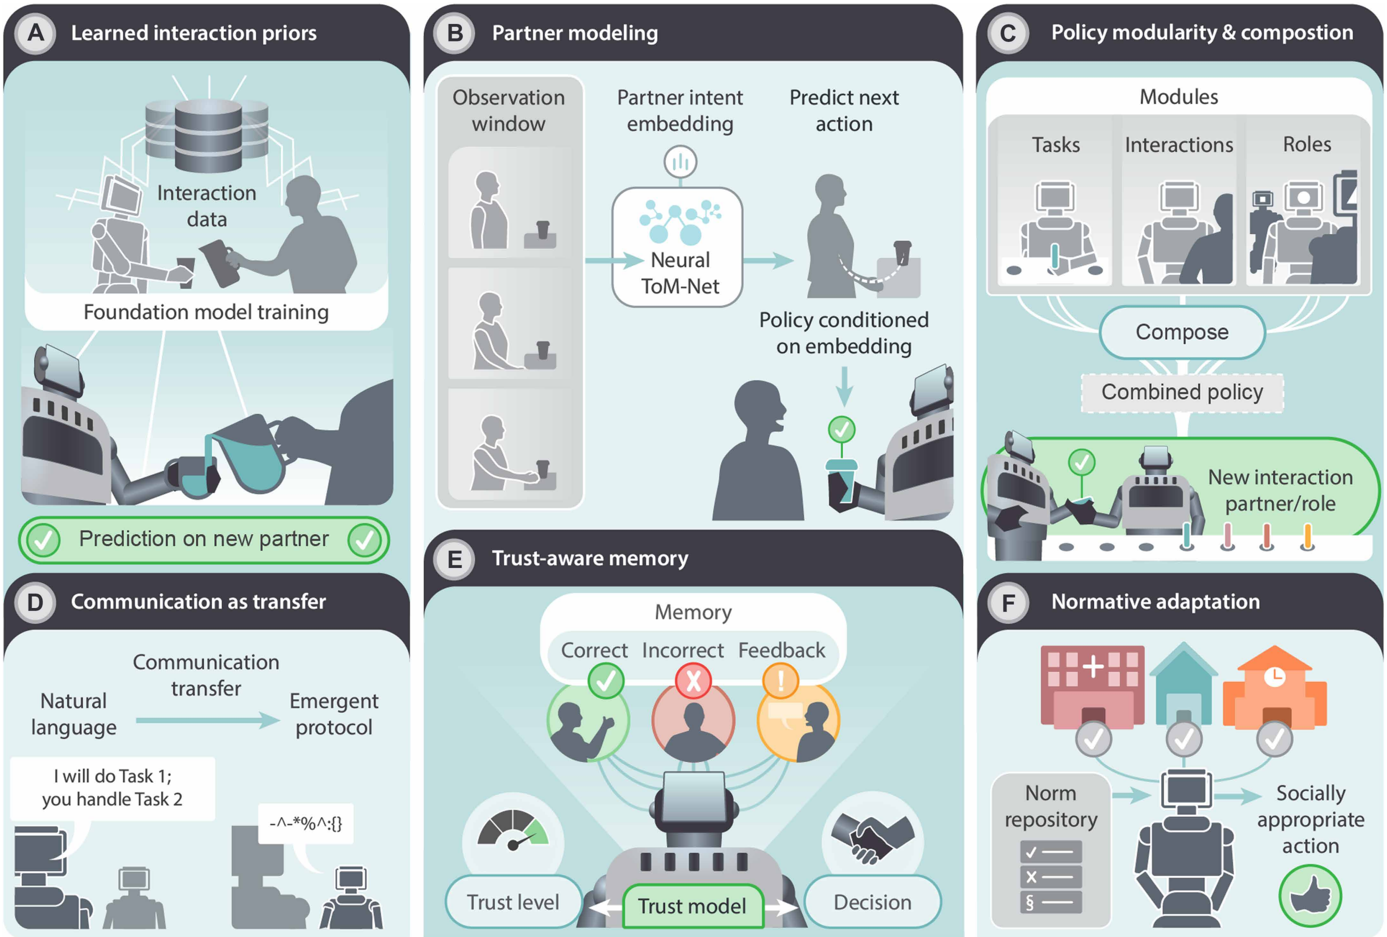

Fig. 1. Instantiations of the six design pillars for alliance-aware robotic foundation models. (**A**) Learned interaction priors: A foundation model is trained on multiagent interaction trajectories to learn patterns of coordination and support prediction about new partners. (**B**) Partner modeling: Neural ToM networks infer a partner-intent embedding from a short observation window and condition the robot's policy to adapt accordingly. (**C**) Policy modularity and composition: Task, interaction, and role modules are recomposed into a combined policy, enabling flexible transfer to new partners and social roles. (**D**) Communication as transfer: Language and other communicative channels, including legible actions and concise explanations, are used to share intent, query partners, and provide rationales, enabling more robust coordination with unfamiliar human and robotic collaborators. (**E**) Trust-aware memory: Long-term memory of interaction (correct/incorrect) outcomes, errors, and feedback updates partner-specific trust models, which in turn shapes decisions and role allocation. (**F**) Normative adaptation: Explicit or learned representations of social norms guide context-appropriate behavior across environments such as homes, clinics, and workplaces.

(*19*) that it can dynamically compose into an appropriate strategy for the current social context.

Normative adaptation and social context

Beyond understanding other agents' actions, alliance-capable robots must also grasp the social context and norms governing an environment. Every community or team has implicit rules of behavior, from workplace protocols to cultural etiquette. A robot in a new setting will falter if it violates important norms (imagine a service robot entering a library and speaking at full volume, unaware of the norm of being quiet). Thus, normative adaptation, the ability to learn and conform to the prevailing norms of whatever group the robot joins, is another pillar. Researchers have begun exploring techniques to endow robots with norm awareness. For example, the SONAR (Social Norm Adaptive Robots) architecture integrates symbolic reasoning with machine learning to help robots recognize and act upon social norms and even learn norms of different societies over time (*20*). These efforts show that adding explicit norm reasoning leads to more positive interactions and higher trust. However, many challenges remain: Norms are highly context dependent, vary by community, and can even conflict (*21*). Notably, a robot may face completely new normative settings (say, moving from a hospital ward to a private home) and must generalize its behavior appropriately without extensive retraining for each new "culture."

Trust-aware memory

Trust is the bedrock of effective teamwork. Humans develop trust through repeated interactions, remembering who proved reliable and how past collaborations fared. Likewise, robots should leverage memory of interactions to adjust future behavior and build mutual trust with their partners. This pillar emphasizes the role of long-term memory and learning: As a robot participates in ongoing teams, it should accumulate knowledge about each partner's competencies, preferences, and trustworthiness and use that to improve coordination over time. This implies maintaining a trust model for each partner, updated with experience. Indeed, machines can use an artificial trust metric to guide task and role allocation (*22*); for instance, a rehabilitation robot can increase assistance and guide the movement when its trust estimate suggests that the patient needs guidance but scale back and let the patient lead once the metric shows increased trust or a desire for more autonomous effort. Furthermore, the robot should calibrate its own behavior to earn human trust. Long-term memory of how humans responded to past actions can enable a robot to adjust its strategy to be more trustworthy or to repair trust after an error. In practice, this could be a memory module storing interaction outcomes (successes, errors, and human feedback) and using them to update the policy. Early work shows that a robot that learns a user's preferred way of being assisted (for example, a lighter grasp) can greatly improve that user's comfort and trust. Extending this idea, an alliance-oriented model would accumulate a record of team interactions, who did what and what happened, providing rich context for decisions about delegation, initiative, and coordination. In short, memory and trust go hand in hand: Memory provides data to assess reliability, whereas trust-aware policies ensure that the robot's learning "with us" prioritizes team cohesion, not just individual success.

Communication as transfer: Learning to share and query

Last, communication is a crucial tool for real-time coordination and knowledge transfer between agents. In a multiagent setting, communicating intent can be a quick way for knowledge transfer in real time. When a robot transmits "I will clear aisle A; you handle aisle B," it conveys the minimal intent information that its teammate needs to update its plan, eliminating costly trial-and-error negotiation. Consequently, learning what to communicate, when to communicate, and how to interpret a partner's message has become a central theme in cooperative AI (*23*–*25*). The link to agent-axis generalization is straightforward: An agent endowed with a policy that develops robust communication skills can integrate with unfamiliar partners far more readily than a silent one.

Two complementary approaches can be considered. The first leverages natural language as a shared interlingua; if agents speak in grounded human language, any new teammate that understands the language can join the exchange (*26*). Large language models (LLMs) are increasingly used as an interface to parse free-form messages into structured intent and generate concise replies, thereby offloading much of the linguistic complexity to a pretrained backbone. The second approach involves agents evolving an emergent protocol tailored to the task, but the challenge is ensuring that a new agent can decipher the resulting "private language" (*26*, *27*). In both cases, communication, whether via discrete symbols, continuous vectors, or behavioral cues, acts as a real-time knowledge channel: Agents update one another's internal models and synchronize actions on the fly.

An approach to learn communication protocols is to have a period of coordination training where an agent interacts with a variety of partners (or simulated partner behaviors), forcing it to learn a robust communication method that is not overly specific to one partner's idiosyncrasies. As multirobot teams scale, a resilient, partner-agnostic communication layer, possibly mediated by LLMs or by standardized symbolic vocabularies, will be essential for ensuring that any freshly deployed agent can rapidly come up to speed with what the rest of the group already knows and intends to do. Critically, these communication protocols must remain transparent and readily interpretable by human supervisors to preserve trust and safety in mixed human-robot teams.

A complementary requirement is transparency and explainability, understood as communicative competence: Beyond exchanging intent, alliance-aware robots should furnish legible rationales for key actions, such as "why/why-not" justifications, and calibrated uncertainty so that partners can update their beliefs efficiently, recover from miscoordination, and calibrate trust (*28*). In practice, this means optimizing behaviors for legibility when appropriate, generating natural-language summaries grounded in perception-action traces during execution, and logging decision-rationale pairs to support post hoc and counterfactual queries that aid team learning and repair trust over time (*29*).

EVALUATING ALLIANCE-AWARE GENERALIZATION

Shifting our perspective toward alliance in robotics also necessitates rethinking how we evaluate success. Traditional robot learning benchmarks, whether in simulation or in controlled trials, measure performance on defined tasks with fixed agent roles. To gauge alliance-aware generalization, we must test robots under the fluid conditions of multiagent interaction. Further directions for evaluation and benchmarking should be considered.

Interaction generalization tests

Just as we test generalization to unseen environments or objects, we should test a robot's generalization to new partners. For example, a collaborative robot arm could be trained with one human partner and then tested with a different partner on the same task. Success would be measured not only by task completion but also by coordination efficiency, for example, time to reach a shared understanding or number of miscommunications. Ad hoc generalization is now recognized as a distinct generalization challenge, often missed by standard evaluations. Toolkits like ZSC-Eval generate diverse partner behaviors to benchmark an agent's adaptability to unseen collaborators (*7*). In robotics, one could envision an interaction-generalization suite where, say, a robot is trained with a particular human in the loop and then evaluated with a new person or robot; the performance drop (if any) would reveal how well the model copes with novel partners.

Role-shifting stress tests

Beyond fixed-partner scenarios, alliance-capable robots should handle role changes and team reconfigurations on the fly. Evaluation scenarios can be designed to force such shifts. For example, in a collaborative delivery task, robot A might start as the leader and robot B as the assistant, but if robot A's main tool fails mid-task, robot B must take over leadership. A robust foundation model should recognize and embrace this role swap with minimal loss of efficiency. To test this, we could measure how quickly the new leader picks up the task and how smooth the handoff is. Such role-shifting tests would stress the model's compositional adaptability, essentially checking whether its modular policies truly allow recomposing behavior in real time.

Social robustness and norm adaptation

In safety research, robustness tests probe how a system handles perturbations or adversarial conditions. The social analog would probe a robot's behavior under norm-violating or inconsistent human behaviors. For instance, if a human team member unexpectedly breaks protocol (say, giving a command that contradicts a prior instruction), does the robot blindly follow the last command or pause to cross-check given the discrepancy? A socially robust robot should be resilient to the "noise" of human inconsistency or ambiguous cues. We can evaluate this by intentionally introducing inconsistent partner behaviors or varying cultural norms in test scenarios. Likewise, we can test adaptation to a new cultural context: have a robot perform a task in two virtual cultures (with different social rules) and see whether it adjusts appropriately. Success can be judged by human ratings of the robot's appropriateness or by objective checks (for example, did it avoid doing something rude, like interrupting a speaker?).

Human satisfaction and team performance

Ultimately, the litmus test for alliance-oriented robots is human acceptance and team effectiveness. Evaluations should include human-in-the-loop trials where people subjectively rate the robot's teamwork qualities (for example, "Did the robot make your job easier? Did it adapt to your preferences?") and objective outcomes (overall task success, time, and errors). Prior work suggests that when robots adapt to human norms and preferences, people report higher trust and satisfaction (*20*). For instance, a composite "alliance score" might combine trust ratings, communication efficiency, and team fluency measures. Tracking such metrics alongside task success would help quantify progress toward robots that truly learn with us in partnership.

By adopting an alliance-centered evaluation approach, researchers can better measure the competencies that matter for multiagent adaptability. Developing these benchmarks will also drive the community to build systems that excel not only in isolation but also in the rich social tapestry of real-world human-robot ecologies.

Datasets and simulations

In parallel, we need rich datasets to train foundation models to capture the diversity of multiagent interactions. Large-scale, behaviorally diverse corpora serve as fuel for training foundation models with strong interaction priors. Fortunately, several emerging datasets and simulators are targeting this need. For example, Chang *et al.* (*30*) provided a repository of human-robot joint tasks in simulated homes. Newman *et al.* (*31*) and Jiang *et al.* (*32*) provided multimodal traces (video, gaze, and joystick commands) during human-robot collaboration, useful for modeling assistive HRI. Beyond curated datasets, demonstration data collected through teleoperation or kinesthetic teaching can be mined for multiagent patterns; for instance, teleoperated or hand-guided demonstrations of a robot performing collaborative screwdriving or bolt insertion while a human coworker holds, reorients, or presents the mating part could reveal patterns of timing, compliant alignment, turn-taking, and recovery from interruptions or misalignment. Simulation is another powerful tool: Modern physics simulators like Isaac Gym and Brax can procedurally generate countless multiagent scenarios with varied partner types (*33*, *34*). Combined with techniques like generative imitation learning (*35*) or LLM-driven agents (*26*) to create synthetic partner behaviors, we can vastly expand the diversity of training scenarios. By assembling diverse, instrumented datasets that include each agent's observations, actions, and communications (when applicable), we equip foundation models with the breadth of experience needed to generalize across the manifold possibilities of who they might interact with.

OTHER CONSIDERATIONS

Our discussion here has centered on the learning architectures and computational frameworks for alliance-aware robots. It is important to recognize that these technical learning advances, although necessary, are not exhaustive. Achieving robust human-robot alliances in practice will also depend on parallel improvements in other domains, as we acknowledged in the Introduction. For instance, continual innovation in robot hardware (for example, developing safer, low-power actuators and durable, responsive sensors) will directly influence a robot's ability to operate safely and reliably alongside humans. Similarly, progress in high-fidelity sensing (for example, robotic hands with dense, wide-area touch sensing) may be crucial for enabling fine-grained interactions with partners, such as adjusting grip on the basis of tactile input when handing objects to a person. Furthermore, insights from embodied cognition remind us that a robot's physical design and sensorimotor capabilities fundamentally shape how it can learn and collaborate, a factor that must converge with algorithmic improvements. Beyond those areas, other complementary fronts deserve attention as well. Advancements in HRI design, such as more intuitive communication interfaces and socially grounded feedback mechanisms, will augment the performance of alliance-aware systems by supporting mutual understanding and trust, as we discussed in the "Communication as transfer: Learning to share and query" section. Safety and ethical frameworks are another related consideration, ensuring that as robots become more autonomous partners, their behavior aligns with human values and safety standards.

CONCLUSIONS

The transition from autonomy to alliance redefines what it means for robots to learn and generalize. Rather than mastering skills in isolation, future robotic foundation models must become social learners that continually adapt in dynamic human-robot ecologies. Achieving this will require reimagining algorithms, architectures, and benchmarks with interaction at the center. The reward could be worthwhile: robots that can enter new environments and immediately start learning with the people and machines around them, seamlessly joining our teams, our homes, and our communities. By reframing transfer learning as a multiagent, interactive endeavor, we can unlock a new generation of robots that are not just tools but trustworthy partners in an ever-changing world.

REFERENCES AND NOTES

1. A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu, J. Ibarz, B. Ichter, A. Irpan, T. Jackson, S. Jesmonth, N. J. Joshi, R. Julian, D. Kalashnikov, Y. Kuang, I. Leal, K.-H. Lee, S. Levine, Y. Lu, U. Malla, D. Manjunath, I. Mordatch, O. Nachum, C. Parada, J. Peralta, E. Perez, K. Pertsch, J. Quiambao, K. Rao, M. Ryoo, G. Salazar, P. Sanketi, K. Sayed, J. Singh, S. Sontakke, A. Stone, C. Tan, H. Tran, V. Vanhoucke, S. Vega, Q. Vuong, F. Xia, T. Xiao, P. Xu, S. Xu, T. Yu, B. Zitkovich, RT-1: Robotics transformer for real-world control at scale. arXiv:2212.06817 [cs.RO] (2022).
2. S. Reed, K. Zolna, E. Parisotto, S. G. Colmenarejo, A. Novikov, G. Barth-Maron, M. Gimenez, Y. Sulsky, J. Kay, J. T. Springenberg, T. Eccles, J. Bruce, A. Razavi, A. Edwards, N. Heess, Y. Chen, R. Hadsell, O. Vinyals, M. Bordbar, N. de Freitas, A generalist agent. arXiv:2205.06175 [cs.AI] (2022).
3. Octo Model Team, D. Ghosh, H. Walke, K. Pertsch, K. Black, O. Mees, S. Dasari, J. Hejna, T. Kreiman, C. Xu, J. Luo, Y. L. Tan, L. Y. Chen, P. Sanketi, Q. Vuong, T. Xiao, D. Sadigh, C. Finn, S. Levine, Octo: An open-source generalist robot policy. arXiv:2405.12213 [cs.RO] (2024).
4. E. Krimsky, S. H. Collins, Elastic energy-recycling actuators for efficient robots. *Sci. Robot.* **9**, eadj7246 (2024).
5. Z. Zhao, W. Li, Y. Li, T. Liu, B. Li, M. Wang, K. Du, H. Liu, Y. Zhu, Q. Wang, K. Althoefer, S.-C. Zhu, Embedding high-resolution touch across robotic hands enables adaptive human-like grasping. *Nat. Mach. Intell.* **7**, 889–900 (2025).
6. A. Gupta, S. Savarese, S. Ganguli, L. Fei-Fei, Embodied intelligence via learning and evolution. *Nat. Commun.* **12**, 5721 (2021).

7. X. Wang, S. Zhang, W. Zhang, W. Dong, J. Chen, Y. Wen, W. Zhang, "ZSC-Eval: An evaluation toolkit and benchmark for multi-agent zero-shot coordination" in *Advances in Neural Information Processing Systems*, A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, C. Zhang, Eds. (Curran Associates, 2024), vol. 37, pp. 47344–47377.
8. S. Dey, T. Yoshida, R. H. Foerster, M. Ernst, T. Schmalz, R. M. Carnier, A. F. Schilling, A hybrid approach for dynamically training a torque prediction model for devising a human-machine interface control strategy. arXiv:2110.03085 [cs.CE] (2021).
9. S. Dey, S. Ravindran Nair, "ReMAP: Neural model reprogramming with network inversion and retrieval-augmented mapping for adaptive motion forecasting" in *Advances in Neural Information Processing Systems*, A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, C. Zhang, Eds. (Curran Associates, 2024), vol. 37, pp. 25195–25227.
10. S. Dey, B. Paassen, S. R. Nair, S. Boughorbel, A. F. Schilling, Continual learning from simulated interactions via multitask prospective rehearsal for bionic limb behavior modeling. arXiv:2405.01114 [cs.LG] (2024).
11. S. V. Albrecht, P. Stone, Autonomous agents modelling other agents: A comprehensive survey and open problems. *Artif. Intell.* **258**, 66–95 (2018).
12. S. Baron-Cohen, H. Ring, J. Moriarty, B. Schmitz, D. Costa, P. Ell, Recognition of mental state terms: Clinical findings in children with autism and a functional neuroimaging study of normal adults. *Br. J. Psychiatry* **165**, 640–649 (1994).
13. N. C. Rabinowitz, F. Perbet, H. F. Song, C. Zhang, S. M. A. Eslami, M. Botvinick, Machine theory of mind. arXiv:1802.07740 [cs.AI] (2018).
14. P. Stone, G. Kaminka, S. Kraus, J. Rosenschein, Ad hoc autonomous agent teams: Collaboration without pre-coordination. *Proc. AAAI Conf. Artif. Intell.* **24**, 1504–1509 (2010).
15. C. Wang, A. Rahman, I. Durugkar, E. Liebman, P. Stone, N-agent ad hoc teamwork. arXiv:2404.10740 [cs.AI] (2024).
16. S. Nikolaidis, M. Kwon, J. Forlizzi, S. Srinivasa, Planning with verbal communication for human-robot collaboration. *ACM Trans. Hum. Robot Interact.* **7**, 1–21 (2018).
17. M. Chen, S. Nikolaidis, H. Soh, D. Hsu, S. Srinivasa, Trust-aware decision making for human-robot collaboration: Model learning and planning. *ACM Trans. Hum. Robot Interact.* **9**, 1–23 (2020).
18. C. Devin, A. Gupta, T. Darrell, P. Abbeel, S. Levine, "Learning modular neural network policies for multi-task and multi-robot transfer" in *IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2017), pp. 2169–2176.
19. S. Barrett, A. Rosenfeld, S. Kraus, P. Stone, Making friends on the fly: Cooperating with new teammates. *Artif. Intell.* **242**, 132–171 (2017).
20. D. Dell'Anna, A. Jamshidnejad, SONAR: An adaptive control architecture for social norm aware robots. *Int. J. Soc. Robot.* **16**, 1969–2000 (2024).
21. B. F. Malle, E. Rosen, V. B. Chi, M. Berg, P. Haas, "A general methodology for teaching norms to social robots" in *29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)* (IEEE, 2020), pp. 1395–1402.
22. C. Centeio Jorge, E. J. De Visser, M. L. Tielman, C. M. Jonker, L. P. Robert, Artificial trust in mutually adaptive human-machine teams. *Proc. AAAI Symp. Ser.* **4**, 18–23 (2024).
23. J. N. Foerster, Y. M. Assael, N. de Freitas, S. Whiteson, "Learning to communicate with deep multi-agent reinforcement learning" in *Advances in Neural Information Processing Systems*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, R. Garnett, Eds. (Curran Associates, 2016), vol. 29; https://proceedings.neurips.cc/paper_files/paper/2016/file/c7635bfd99248a2cdef8249ef7bfbef4-Paper.pdf.
24. S. Sukhbaatar, A. Szlam, R. Fergus, "Learning multiagent communication with backpropagation" in *Advances in Neural Information Processing Systems*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, R. Garnett, Eds. (Curran Associates, 2016), vol. 29; https://proceedings.neurips.cc/paper_files/paper/2016/file/55b1927fdafef39c48e5b73b5d61ea60-Paper.pdf.
25. A. Fachantidis, M. Taylor, I. Vlahavas, Learning to teach reinforcement learning agents. *Mach. Learn. Knowl. Extr.* **1**, 21–42 (2019).
26. P. Zhou, Y. Feng, H. Julaiti, Z. Yang, Why do AI agents communicate in human language? arXiv:2506.02739 [cs.AI] (2025).
27. A. Lazaridou, M. Baroni, Emergent multi-agent communication in the deep learning era. arXiv:2006.02419 [cs.CL] (2020).
28. A. D. Dragan, K. C. T. Lee, S. S. Srinivasa, "Legibility and predictability of robot motion" in *8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (IEEE, 2013), pp. 301–308.
29. U. Ehsan, P. Tambwekar, L. Chan, B. Harrison, M. O. Riedl, "Automated rationale generation: A technique for explainable AI and its effects on human perceptions" in *IUI '19: Proceedings of the 24th ACM International Conference on Intelligent User Interfaces* (Association for Computing Machinery, 2019), pp. 263–274.
30. M. Chang, G. Chhablani, A. Clegg, M. D. Cote, R. Desai, M. Hlavac, V. Karashchuk, J. Krantz, R. Mottaghi, P. Parashar, S. Patki, I. Prasad, X. Puig, A. Rai, R. Ramrakhya, D. Tran, J. Truong, J. M. Turner, E. Undersander, T.-Y. Yang, PARTNR: A benchmark for planning and reasoning in embodied multi-agent tasks. arXiv:2411.00081 [cs.RO] (2024).
31. B. A. Newman, R. M. Aronson, S. S. Srinivasa, K. Kitani, H. Admoni, HARMONIC: A multimodal dataset of assistive human–robot collaboration. *Int. J. Robot. Res.* **41**, 3–11 (2022).
32. S. Jiang, H. Li, R. Ren, Y. Zhou, Z. Wang, B. He, Kaiwu: A multimodal manipulation dataset and framework for robot learning and human-robot interaction. arXiv:2503.05231 [cs.RO] (2025).
33. V. Makoviychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire, A. Handa, G. State, Isaac Gym: High performance GPU-based physics simulation for robot learning. arXiv:2108.10470 [cs.RO] (2021).
34. C. D. Freeman, E. Frey, A. Raichuk, S. Girgin, I. Mordatch, O. Bachem, Brax—A differentiable physics engine for large scale rigid body simulation. arXiv:2106.13281 [cs.RO] (2021).
35. J. Ho, S. Ermon, "Generative adversarial imitation learning" in *Advances in Neural Information Processing Systems*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, R. Garnett, Eds. (Curran Associates, 2016), vol. 29; https://proceedings.neurips.cc/paper_files/paper/2016/file/cc7e2b878868cbae992d1fb743995d8f-Paper.pdf.

Acknowledgments

Funding: This research was not supported by any external funding. **Author contributions:** Conceptualization: S. Dey. Methodology: S. Dey. Visualization: S. Dey, S.V.A., S. Dosen, and R.R. Investigation: S. Dey, S.V.A., S. Dosen, and R.R. Writing—original draft: S. Dey. Writing—review and editing: S. Dey, S.V.A., S. Dosen, and R.R. Supervision: S. Dey, S.V.A., S. Dosen, and R.R. Project administration: S. Dey. **Competing interests:** The authors declare that they have no competing interests.

10.1126/scirobotics.aea1822

Science Robotics

From autonomy to alliance: Robotic foundation models must learn with us, not just for us

Sharmita Dey, Robert Riener, Strahinja Dosen, and Stefano V. Albrecht

Sci. Robot. **11** (113), eaea1822. DOI: 10.1126/scirobotics.aea1822

Science Robotics (ISSN 2470-9476) is published by the American Association for the Advancement of Science. 1200 New York Avenue NW, Washington, DC 20005. The title *Science Robotics* is a registered trademark of AAAS.