Research Statement

Harish Ravichandar

Assistant Professor, School of Interactive Computing
Georgia Institute of Technology

harish.ravichandar@cc.gatech.edu
star-lab.cc.gatech.edu

April 2026

Robots are beginning to leave carefully engineered settings and enter the complex, dynamic, and resource-constrained environments in which we actually live and work. Robust and sustained success in these complex settings will require more than equipping robots with raw capabilities, either in the form of generic pretrained models or bespoke hand-designed controllers. Robots must **learn efficiently** from available data, **coordinate seamlessly** with other robots and with people, and **adapt robustly** as partners, tasks, and environments inevitably change post-deployment. I am an algorithmic roboticist who addresses these challenges by asking a foundational question: *how can we leverage the inherent physical and resource structures of our world to help make robot intelligence more efficient, collaborative, and adaptable?*

At Georgia Tech, I direct the Structured Techniques for Algorithmic Robotics (STAR) Lab. In our lab, we take the view that geometry, dynamics, embodiment, capabilities, and resource limitations are not nuisances to be abstracted away, but are key factors to be modeled, preserved, and leveraged as **generalizable structural priors**. We therefore develop computational and learning frameworks that *inject principled inductive biases into modern learning and computational methods through representations, architectures, objectives, and constraints.*

Our work has demonstrated that the structural priors that we have identified and developed can help make robots more **frugal** in their use of data and compute; **self-sufficient** in their ability to learn and improve with limited expert supervision or hand-tuning; and **translucent** by making their skills and decisions more understandable, predictable, and intuitive. My lab has advanced this agenda through a coherent body of work spanning diverse applications, from **dexterous manipulation** to **multi-robot coordination**, motivated by timely and ambitious problems (e.g., *learning dexterity from human videos*, *automating chicken deboning*, and *full-stack coordination of heterogeneous teams*).

Research

My lab's research is organized around **three complementary thrusts**: Structured motor skill learning, which studies *how geometry and dynamics can make dexterous behavior easier to learn and transfer*; Structured coordination, which develops *capability-centric frameworks for heterogeneous multi-robot teams*; and Structured adaptation, which *enables robots to adapt to uncertainty and unexpected setbacks in tasks, teammates, and environments*. Although these thrusts span different levels of robot intelligence and different application domains, they are united by a common scientific premise: **robots should not have to discover everything from scratch when inherent physical and resource *structures* can provide useful and *generalizable* guidance**. Across these thrusts, our goal is not only to improve task success, but to build methods that improve robots' *frugality*, *self-sufficiency*, and *translucency*. See Fig. 1 for a visual summary.

Thrust 1: Structured Skill Learning

Our first thrust develops structured methods for motor skill learning, with an emphasis on dexterous manipulation. Dexterous behavior is complex, high-dimensional, contact-rich, and expensive to supervise, making unstructured end-to-end mappings especially inefficient and brittle. In this thrust, we study how **structures inspired by dynamical systems theory and human skill acquisition** can serve as principled inductive biases for learning dexterous skills. The intellectual roots of this thrust are grounded in my earlier work on robot learning from demonstrations, including a 2020 *Annual Review* article [1, Annual Review'20] that helped shape the perspective underlying my current agenda.

A central theme of this thrust is leveraging dynamical systems theory to overcome the inherent limitations of purely reactive policies that map observations to actions: limited temporal coherence and the inability to consider

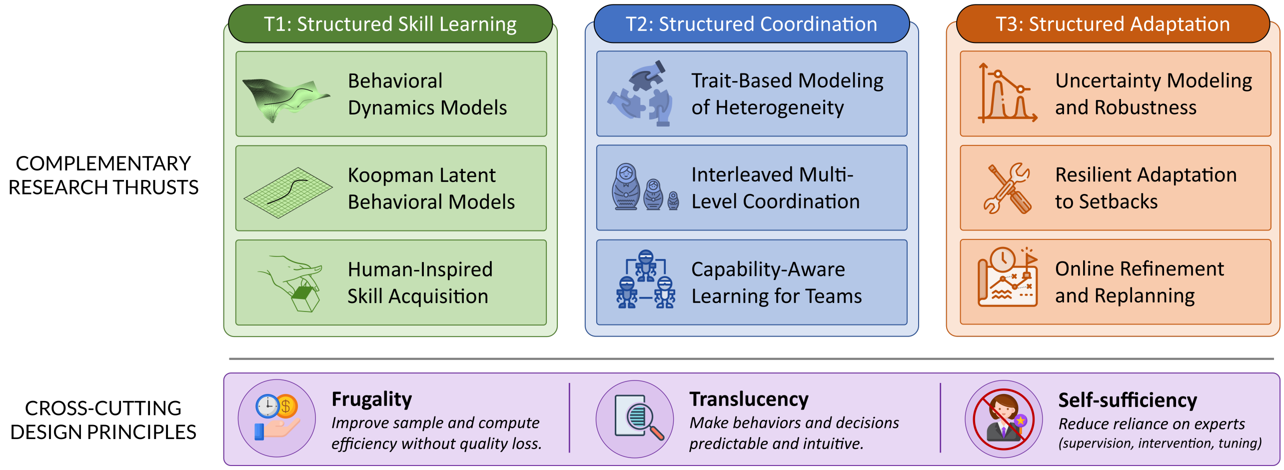

Figure 1: A Structured Approach to Robotic Intelligence. My research is organized around three complementary thrusts (structured skill learning, structured coordination, and structured adaptation) that collectively aim to make robot intelligence more frugal, translucent, and self-sufficient.

the future beyond inflexible and short-horizon "chunks". One early step in this direction was *Neural Geometric Fabrics (NGF)* [2, CoRL'22], where we and our collaborators at NVIDIA showed how **geometric and nonlinear dynamical structure** can support efficient learning of high-dimensional policies from demonstration. Building on that perspective, my group has developed a sustained line of work around **Koopman-inspired approaches for dexterous manipulation**, which has grown to become one of the most visible and distinctive parts of this thrust. In KODex [3, CoRL'23 Oral], we showed how behavioral models that capture *structured latent linear dynamics* can dramatically reduce policy-training and hyperparameter-search time compared to standard neural architectures. In subsequent collaboration with colleagues at CMU, we expanded this perspective toward richer manipulation settings with noisy vision-based perception [4, CoRL'24]. This line of work has recently matured into a broader conceptual contribution through what we call **Unified Behavioral Models (UBMs)** [5, Preprint'26], which represent dexterous skills as solutions to *coupled* dynamical systems that encode the co-evolution of nominal robot behavior and its desired impact on the environment. A UBM can serve as an implicit "planner" that predicts skill evolution over the entire task horizon, maintains temporal coherence by construction, and monitors and responds to deviations from nominal outcomes. We also developed the Koopman Behavioral Model (K-UBM), a first instantiation of our broader UBM framework that brings these conceptual benefits into practice.

In a complementary line of work, we have been developing algorithms that enable robots to **learn through active practice**, rather than only learn from human demonstrations. A key aspect of our work is the active involvement of the robot (i.e., **skill practice**) when learning from experts, rather than relying only on the passive learning paradigm of standard behavioral cloning. In CIMER, we used the idea of **emulation** (the process of learning to achieve the same *impact* on the environment as the expert, rather than merely copying the expert's action) to learn dexterous skills from state-only observations (e.g., videos) that lack action information [6, RA-L'24; 7, ICRA'25]. In AsymDex [8, CoRL'25 WS], we exploited the **asymmetric hand roles** that commonly arise in human bimanual manipulation. In ImMimic [9, CoRL'25 Oral] and our work on dynamically adaptive cutting for chicken deboning [10, ICRA'26], we further studied how robots can acquire coarse skills from limited supervision and later perform **coarse-to-fine skill refinement** through interaction and practice data. In particular, our work on automated chicken deboning is supported through a close collaboration with GTRI, funded by a **U.S. Department of Agriculture grant**.

Together, these projects broaden the thrust beyond a single modeling tool while reinforcing our three core design principles: *democratizing* dexterous learning via the frugal use of data and computation, improving *self-sufficiency* by reducing dependence on dense supervision and hand-tuning, and improving *translucency* by making learned behaviors easier to anticipate, inspect, and refine.

Thrust 2: Structured Coordination

Our second thrust develops structured approaches for coordinating multi-robot teams, with an emphasis on heterogeneous teams. A key challenge is that heterogeneous robots can differ in numerous ways: sensing, actuation, costs, roles, endurance, and resource availability. However, not all differences are meaningful and any tractable approach requires the right level of abstraction. This research thrust has helped us establish a **capability-centric view of multi-robot coordination** in which heterogeneity is modeled as a first-class attribute, and the collective impact of robots on task performance is emphasized over individual robots' identities or inconsequential idiosyncrasies.

This thrust began with STRATA [11, AAMAS'21], a unified capability-based formulation, with an emphasis on task allocation. However, effective coordination is a *coupled* problem across multiple *interacting* levels of abstraction: task planning, allocation, scheduling, and motion planning. We thus invested in a broader line of work on **interleaved multi-level coordination**. In ITAGS [12, IROS'21] and GRSTAPS [13, IJRR'22], we showed that interleaving the four levels is not only possible, but often preferable to conventional sequential decompositions. A key contribution is showing that *interleaving allows decisions at one layer to remain compatible with constraints at the others*, enabling more efficient search and reducing the need for expensive sequential backtracking. D-ITAGS [14, RA-L'23] extended this to dynamic settings, introducing *targeted repair* mechanisms that preserve unaffected parts of a solution. More recently, we went beyond binary feasibility by explicitly modeling *how collective capabilities influence task performance* in Q-ITAGS [15, ISRR'24], and accounted for *exhaustible capabilities* and *battery constraints* in TRAITS [16, AAMAS'26 Oral].

A second theme of this thrust is that structured capability models can serve as useful **scaffolding for learning-based coordination**. Our work on *capability awareness and communication* [17, CoRL'23], *capability-aware shared hypernetworks (CASH)* [18, CoRL'25; 19, AAMAS'25], and *learning how collective capabilities map to team rewards* [20, RSS'23; 15, ISRR'24] helped replace brittle handcrafted heuristics with learned models that leveraged explicit representations of robots' relative capabilities and structured communication to flexibly shape team behavior. We also demonstrated that learned team behaviors can be explained by *identifying influential inter-agent communications* in GNN-based policies [21, AAMAS'25]. This thrust especially highlights *frugality* and *translucency*: exposing capabilities and shared constraints makes team behavior not only more efficient to learn and generalize, but also more legible and predictable.

A distinctive feature of this thrust is that it has grown beyond research to enable **open-source community infrastructure**. *MARBLER* [22, MRS'23] provides an open-source, zero-cost, and turnkey framework for standardized *hardware* evaluation of multi-robot reinforcement learning algorithms by leveraging Georgia Tech's Robotarium, and *JaxRobotarium* [23, CoRL'25] dramatically reduces the time required to train and deploy multi-robot policies on hardware. Together, these systems provide reproducible hardware evaluation, lower barriers to sim-to-real experimentation, and shorten the path from policy training to physical deployment.

Thrust 3: Structured Adaptation

Our third thrust asks how robotic systems can **adapt** effectively when environments, partners, and task demands are uncertain or change unexpectedly. Adaptation is a central requirement for any robotic system to operate effectively in partially observable, stochastic, and dynamic real-world environments. To ensure that adaptation does not come at the expense of reliability, we emphasize **structured adaptation**: updating only the components implicated by new information or failure while preserving valid behavior, stability, and task-relevant constraints. Across our work, this perspective has led to three complementary mechanisms: reasoning about known uncertainty before execution, repairing only the parts of a solution affected by unexpected events, and refining behavior through interaction and practice.

When relevant uncertainty can be modeled, adaptation should be anticipated rather than deferred until failure. In the multi-robot setting, we have developed methods for risk-adaptive task allocation [24, IROS'21], resource-aware adaptation of coalition formation strategies [25, AAMAS'22], risk-tolerant task allocation and scheduling [26, IROS'23], coordination under temporal uncertainty [27, RSS'23], and concurrent optimization of initially unknown rewards [20, RSS'23]. These methods expose the particular uncertainty or resource structure that matters to a decision, allowing robot teams to remain robust without discarding the efficiency and translucency of our capability-centric models.

Unexpected events often invalidate only a small part of an existing solution. D-ITAGS [14, RA-L'23] traces how disruptions such as lost robots, new tasks, or environmental changes propagate through allocation, scheduling, and motion-planning decisions, preserving unaffected parts of the solution and recomputing only what must change. The same principle motivates K-UBM [5, Preprint'26]: when predicted and observed skill evolution diverge, the model replans from the current state while preserving the learned dynamical structure. In both cases, the right representation makes it possible to localize what must change rather than recompute or relearn from scratch.

Some performance gaps become visible only after a robot acts. CIMER [6, RA-L'24; 7, ICRA'25], ImMimic [9, CoRL'25 Oral], and dynamically adaptive cutting [10, ICRA'26] use interaction and practice data to identify and refine the parts of a skill that transfer poorly from observation, another embodiment, or an initial coarse policy. In human-centered settings, related work uses task-relevant abstractions such as cognitive state [28, RO-MAN'22] and anticipated comfort [29, RO-MAN'23] to adapt team decisions and robot trajectories without requiring a complete model of the person.

Among my three design principles, this thrust most clearly ties together *self-sufficiency* and *frugality*. It supports self-sufficiency by enabling robots to improve and recover with less routine expert diagnosis and intervention. It promotes frugality by preserving valid decisions, models, and experience and adapting only the components that need to change. The central premise is that robust adaptation should change what is necessary, preserve what remains useful, and make the basis for that change intelligible.

Looking Ahead

My research agenda in structured robot learning and decision making pursues a principled middle ground between two unsatisfying extremes in robotics: brittle hand-engineering on one hand, and brute-force scaling on the other. My central claim is that *generalizable structural priors improve robot intelligence*. By identifying and exploiting the right generalizable priors, whether from geometry, dynamics, embodiment, heterogeneity, resources, or the human context, we can build robotic systems that are more efficient, more independent, and more reliable in practice.

Looking ahead, my lab aims to deepen and extend this agenda along three timely directions. First, we will investigate how structural priors can serve as *effective scaffolding to operationalize the open-world "common sense" of large pretrained models*. Large vision, language, and action models can provide broad semantic knowledge, but their predictions alone do not guarantee physically feasible, temporally coherent, or reliable behavior. We plan to use these models to identify goals, objects, constraints, and task-relevant features, while geometric, dynamical, and behavioral structures govern skill evolution, monitor execution, and support replanning.

Second, we aim to build structures for coordination that *model and integrate the impact of low-level execution on higher-level planning*. Most high-level multi-robot planners treat execution time, resource consumption, success likelihood, and task requirements as direct inputs, even though these quantities emerge from embodiment, interaction, and team composition. We will develop hierarchical frameworks in which experience at the skill-execution layer updates coordination models, while task allocation and scheduling decisions shape skill selection, learning, and resource use.

Our third goal is to develop adaptive methods that *substantially reduce the need for routine expert intervention after deployment*. Robot learning still relies heavily on experts to diagnose failures, collect targeted data, tune objectives, and decide what to retrain. I plan to develop methods that detect meaningful deviations, localize likely causes, select informative experience, and update only the affected components while preserving validated behavior and constraints.

Together, these directions will advance a lasting research agenda in robotics: showing how the right structure can operationalize broad learned knowledge, connect decisions across levels of abstraction, and enable robots to become more frugal, self-sufficient, and predictable in practice.

References

[1] **Harish Ravichandar**, Athanasios S. Polydoros, Sonia Chernova, and Aude Billard. Recent advances in robot learning from demonstration. *Annual Review of Control, Robotics, and Autonomous Systems*, 3(1), 2020.

[2] Mandy Xie, Karl Van Wyk, Ankur Handa, Stephen Tyree, Dieter Fox, **Harish Ravichandar**, and Nathan Ratliff. Neural geometric fabrics: Efficiently learning high-dimensional policies from demonstration. In *Proceedings of the Conference on Robot Learning (CoRL)*, 2022.

[3] Yunhai Han, Mandy Xie, Ye Zhao, and **Harish Ravichandar**. On the utility of koopman operator theory in learning dexterous manipulation skills. In *Proceedings of the Conference on Robot Learning (CoRL)*, 2023. Oral Spotlight.

[4] H. Chen, A. Abuduweili, A. Agrawal, Yunhai Han, **Harish Ravichandar**, C. Liu, and J. Ichnowski. KOROL: Learning visualizable object feature with koopman operator rollout for manipulation. In *Proceedings of the Conference on Robot Learning (CoRL)*, 2024.

[5] Yunhai Han, Linhao Bai, Z. Xiao, Zhaodong Yang, Y. Choudhary, K. Jha, C. Kong, Shishir Kousik, and **Harish Ravichandar**. Going with the flow: Koopman behavioral models as implicit planners for visuo-motor dexterity, 2026. Manuscript under review.

[6] Yunhai Han, Zhenyang Chen, K. Williams, and **Harish Ravichandar**. Learning prehensile dexterity by imitating and emulating state-only observations. *IEEE Robotics and Automation Letters*, 2024.

[7] Yunhai Han, Zhenyang Chen, K. Williams, and **Harish Ravichandar**. CIMER: Learning prehensile dexterity by imitating and emulating state-only observations. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2025.

[8] Zhaodong Yang, Yunhai Han, and **Harish Ravichandar**. AsymDex: Leveraging asymmetry and relative motion in learning bimanual dexterity. In *Workshop on Whole-body Control and Bimanual Manipulation: Applications in Humanoids and Beyond, Conference on Robot Learning (CoRL)*, 2025. Oral Spotlight.

[9] Yifan Liu, W. C. Shin, Yunhai Han, Zhenyang Chen, **Harish Ravichandar**, and Danfei Xu. ImMimic: Cross-domain imitation from human videos via mapping and interpolation. In *Proceedings of the Conference on Robot Learning (CoRL)*, 2025. Oral Spotlight.

[10] Zhaodong Yang, A. P. Hu, and **Harish Ravichandar**. Towards automated chicken deboning via learning-based dynamically-adaptive 6-DoF multi-material cutting. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2026.

[11] **Harish Ravichandar**, Kenneth Shaw, and Sonia Chernova. STRATA: Unified framework for task assignments in large teams of heterogeneous agents. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS), Journal Track*, 2021.

[12] Glen Neville, Andrew Messing, **Harish Ravichandar**, Seth Hutchinson, and Sonia Chernova. An interleaved approach to trait-based task allocation and scheduling. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021.

[13] Glen Neville, Andrew Messing, Sonia Chernova, Seth Hutchinson, and **Harish Ravichandar**. GRSTAPS: Graphically recursive simultaneous task allocation, planning, and scheduling. *The International Journal of Robotics Research*, 41(2), 2022.

[14] Glen Neville, Sonia Chernova, and **Harish Ravichandar**. D-ITAGS: A dynamic interleaved approach to resilient task allocation, scheduling, and motion planning. *IEEE Robotics and Automation Letters*, 2023.

[15] Jiazhen Liu, Glen Neville, Sonia Chernova, and **Harish Ravichandar**. Quality-optimized interleaved task allocation and scheduling. In *Proceedings of the International Symposium on Robotics Research (ISRR)*, 2024.

[16] Jinwoo Park, **Harish Ravichandar**, and Seth Hutchinson. Modeling and optimizing the provisioning of exhaustible capabilities for simultaneous task allocation and scheduling. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2026. Oral Presentation.

[17] Pierce Howell, Max Rudolph, Kevin Fu, Reza Torbati, and **Harish Ravichandar**. Towards generalizable heterogeneous multi-robot coordination via trait awareness and communication. In *Proceedings of the Conference on Robot Learning (CoRL)*, 2023.

[18] Kevin Fu, Shalin Jain, Pierce Howell, and **Harish Ravichandar**. Capability-aware shared hypernetworks for flexible heterogeneous multi-robot coordination. In *Proceedings of the Conference on Robot Learning (CoRL)*, 2025.

[19] Kevin Fu, Pierce Howell, Shalin Jain, and **Harish Ravichandar**. Learning flexible heterogeneous coordination with capability-aware shared hypernetworks. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2025.

[20] Sukriti Singh, Anusha Srikanthan, Vivek Mallampati, and **Harish Ravichandar**. Concurrent constrained optimization of unknown rewards for multi-robot task allocation. In *Proceedings of Robotics: Science and Systems (RSS)*, 2023.

[21] Siva Kailas, Shalin Jain, Jiazhen Liu, and **Harish Ravichandar**. Evaluating and improving graph-based explanation methods for multi-agent coordination. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2025.

[22] Reza Torbati, S. Lohiya, S. Singh, M. S. Nigam, and **Harish Ravichandar**. MARBLER: An open platform for standardized evaluation of multi-robot reinforcement learning algorithms. In *Proceedings of the IEEE International Symposium on Multi-Robot Systems (MRS)*, 2023. Best Paper Award.

[23] Shalin Jain, Jiazhen Liu, Siva Kailas, and **Harish Ravichandar**. JaxRobotarium: Training and deploying multi-robot policies in 10 minutes. In *Proceedings of the Conference on Robot Learning (CoRL)*, 2025.

[24] Max Rudolph, Sonia Chernova, and **Harish Ravichandar**. Desperate times call for desperate measures: Towards risk-adaptive task allocation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021.

[25] Anusha Srikanthan and **Harish Ravichandar**. Resource-aware adaptation of heterogeneous strategies for coalition formation. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS), Extended Abstract*, 2022.

[26] Jinwoo Park, Andrew Messing, **Harish Ravichandar**, and Seth Hutchinson. Risk-tolerant task allocation and scheduling in heterogeneous multi-robot teams. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023.

[27] Andrew Messing, Jacopo Banfi, Martina Stadler, Ethan Stump, **Harish Ravichandar**, Nicholas Roy, and Seth Hutchinson. A sampling-based approach for heterogeneous coalition scheduling with temporal uncertainty. In *Proceedings of Robotics: Science and Systems (RSS)*, 2023.

[28] Jack Kolb, **Harish Ravichandar**, and Sonia Chernova. Leveraging cognitive states in human-robot teaming. In *Proceedings of the IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, 2022. Best Student Paper Award Finalist.

[29] Oscar Jed Chuy, Hritik Sapra, Xiang Zhi Tan, **Harish Ravichandar**, and Sonia Chernova. Benefits of multi-objective trajectory adaptation in close-proximity human-robot interaction. In *Proceedings of the IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, 2023.