# Research Statement

Recent advances in hardware, compute, and learning algorithms have accelerated the physical capabilities of robots to new levels. Around the world, robots are being deployed outside of traditional factory settings, making a future in which humans and robots co-exist and collaborate seem within reach. Yet, this perception is an illusion; as they are typically deployed *in isolation from people* or, when deployed in human-centric environments, are either too lightweight and compliant to perform meaningful physical work, or so risk-averse in shared spaces that sustained physical interaction is effectively *nullified*. For robots to truly make a positive impact on society and improve quality of life, they must move, adapt, and collaborate as fluidly as humans do with each other. Even the most advanced data-driven systems struggle to achieve what I call **fluid physical interaction**: a regime in which motion, force, and intent predictions (from the human, the task, and the environment) are continuously co-regulated enabling seamless, safe exchanges of energy and information during physical interaction and collaboration. A key reason for this limitation is that many learning-based robotic systems merely automate traditional trajectory-planning and tracking pipelines built for static environments. My research closes this gap by developing structured, physically grounded abstractions of robot motion and interaction through a **unified dynamical-systems lens**.

**Towards Fluid Physical Interaction.** My ultimate goal is to understand and develop the computational tools to achieve the *physical and perceptual adaptive intelligence* needed for robots to learn from and assist humans across a wide spectrum of capabilities and interaction contexts, from dexterous manipulation taught by able-bodied users to physical assistance in activities of daily living. To achieve this, my research develops **tightly coupled constraint-aware learning, control, and robust estimation algorithms** with formal guarantees of safety, efficiency, and robustness. A key insight is that fluid physical interaction requires perception, planning, and control to remain simultaneously adaptive, reactive, and stable, while accounting for *perceptual, physical, and cognitive* constraints of both the human and the robot. Thus, my research is centered around a dynamical-systems (DS) paradigm to learning and control, in which robot motion policies are represented as autonomous, constraint-aware dynamical systems rather than open-ended feed-forward policies. This formulation enables strong control-theoretic guarantees, including stability, convergence, and invariance, while remaining sample and computation-efficient during training and inference [1, 2, 3, 4]. When combined with impedance-based interaction control, these learned dynamical policies are robust to physical perturbations and capable of modulating interaction forces in an energy-aware manner; properties that are essential for safe and stable physical collaboration. The theoretical foundation of this framework is presented in my co-authored textbook, *Learning for Adaptive and Reactive Robot Control: A Dynamical Systems Approach* (MIT Press, 2022) [5].

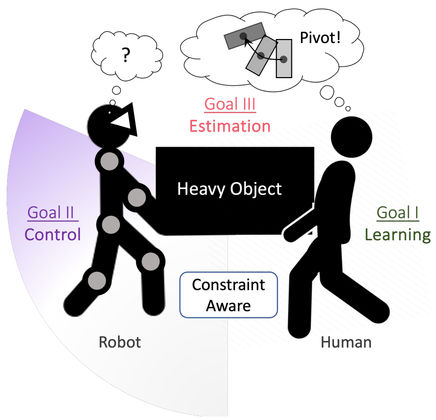


My research since founding my group at Penn, has converged into three research thrusts that advance learning, control, and estimation for physical human–robot interaction across different embodiments and task complexity. We have published 25 papers in premiere robotics venues with several receiving international recognition in these directions which I will summarize next.

**Thrust 1: Structured Learning for Generalizable and Stable Manipulation.** Teaching robots from demonstrations has become the dominant paradigm for manipulation and whole-body control, yet task generalization remains limited in low-data regimes. As a result, much of the field has shifted toward data-hungry learning approaches, which scale performance but sacrifice formal guarantees of stability and convergence. In contrast, dynamical-systems–based learning methods naturally provide such guarantees, but historically lack the expressivity and task awareness required to represent complex manipulation skills or transfer them to new task instances.

To address this gap my group has advanced DS-based learning along three axes. First, to improve **efficiency**, we developed parallelized statistical inference techniques that enable near real-time learning of complex manipulation skills [6]. Second, to increase **expressivity**, we leveraged tools from differential geometry to learn DS motion policies directly on $\mathcal{SE}(3)$, enabling robust teaching of full 6-DoF tasks such as door opening, pouring, and structured placement under perturbations [7]. We further explored neural ODEs for cyclic motion [8] and contraction-theoretic tools for stability in latent task representations [9]. Beyond motion-level learning, we have lifted stability and convergence guarantees to the **task-execution layer**. By integrating continuous DS motion policies with hybrid systems theory and temporal logic, we developed methods that detect task-level failures, learn from them, and recover online under environmental or contextual changes [10, 11]. Most recently, we introduced **Elastic Motion Policies**, a unified framework that supports real-time task transfer, continuous adaptation, and semantic modulation of learned skills, with transfer to new task instances achieved in under one second without additional demonstrations [12]. Recent extensions bridge our structured DS learning with visuomotor and generative approaches, introducing force and compliance-aware simulation pipelines that generate structured training data from minimal human demonstrations, significantly improving robustness in contact-rich manipulation.

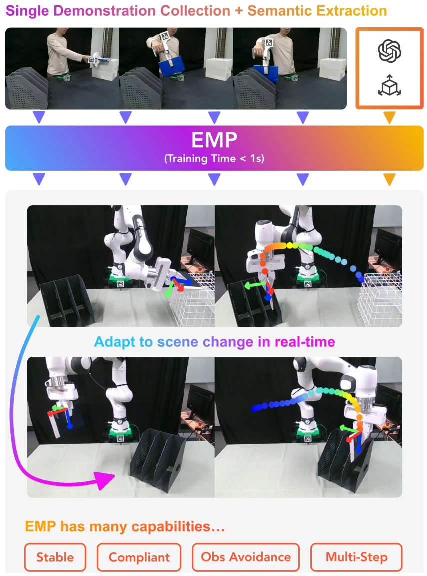

Figure 1: Elastic Motion Policies.

**Thrust 2: Safe and Adaptive Control for Physical Interaction.** Safety in robotics has traditionally been framed as collision avoidance. However, when robots must physically interact with the world and with humans, effective interaction requires robots to continuously decide *when to avoid*, *when to comply* and allow contacts, and *when to transition* between these behaviors. I showed that relaxing overly conservative safety definitions; e.g., avoiding collisions when possible while mitigating impact when unavoidable, yields safer and more efficient interaction [13]. However, guaranteeing safety in dynamic, cluttered environments without sacrificing reactivity or task progress is still a challenge. Classical global planners are too slow, while purely reactive methods suffer from local minima and freezing. To address this, I have pioneered real-time safety-aware control methods that enable robots to safely co-exist with humans while escaping local minima. At the theoretical level, we introduced **on-manifold modulation** of DS policies, which treats differentiable safety boundaries as geometric manifolds and navigates along their isosurfaces at kilohertz rates [14]. We have recently generalized this theory for non-holonomic systems with tight actuation constraints [15, 16], and combined with sampling-based model predictive control for high-dimensional joint spaces [17, 18]. Furthermore, in tasks where physical contact between robots and humans is allowed or required, safety must be achieved through *passive and compliant interaction* while simultaneously preventing self-harm and constraint violations. Classical impedance and passivity-based torque controllers are inherently unconstrained and can drive robots into unsafe or infeasible configurations under large perturbations. To address this, we developed a real-time, optimization-based torque control architecture that enforces hard safety constraints, including joint

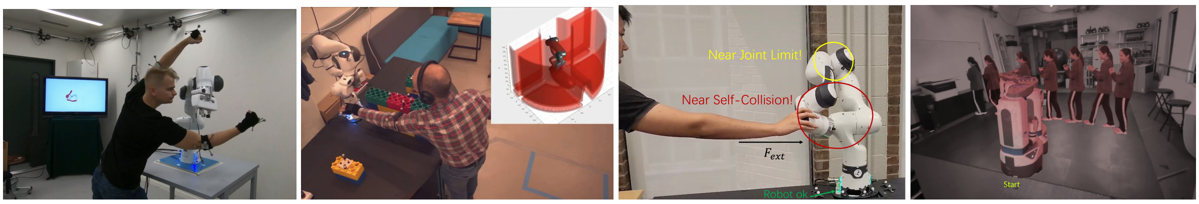

Figure 2: Joint Space Collision Avoidance [18], On-Manifold Workspace Modulation Strategy [14], Constrained Passive Interaction Control [19], Safe Real-time Human-Aware Navigation [16, 15]

limits, self-collisions, external collisions, and singularities; while guaranteeing *passivity only when feasible* [19]. Building on this idea, we introduced **viability-preserving passive torque control**, which uses viability theory to enforce state-dependent torque constraints via a quadratic program, guaranteeing infinite-horizon safety while preserving compliant interaction. Ongoing work extends this framework to whole-body control in mobile manipulators and humanoids.

**Thrust III: Constraint-Aware State Estimation for Interactive & Assistive Tasks** Adaptive learning and safe control require accurate estimation of state, intent, and dynamics under uncertainty. In human–robot interaction, this challenge is compounded by unknown human biomechanics and inherently imperfect perception due to occlusions, noise, and limited sensing. Therefore, I develop *constraint-aware estimation* methods that explicitly reason about physical and perceptual limitations and provide uncertainty-aware outputs for downstream control. To address faulty perception, my group introduced **permanence-aware filters** for multi-object environments and human skeleton tracking that enforce object permanence, biomechanical feasibility, and occlusion reasoning. These methods maintain physically consistent state estimates through long-term occlusions in challenging human-centric settings [20, 21]. Beyond state estimation, we recently developed **Rapid Mismatch Estimation (RME)**, a controller-agnostic probabilistic framework that adapts end-effector dynamics online using proprioception alone [22]. By combining neural priors with variational inference, RME enables safe, compliant adaptation to unknown objects and changing loads within hundreds of milliseconds, without external force sensing. Finally, for physical assistance tasks, we developed a coupled estimation–control framework that jointly reason about human intent, confidence, and shared physical constraints, enabling adaptive assistive forces during collaborative manipulation of large and heavy objects [20].

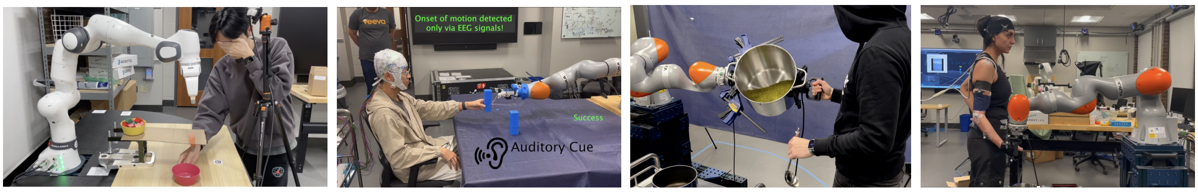


Figure 3: OPF tracking occluded objects with no disruptions [23], Real-time intent estimation from brain signals [24], Constraint-aware co-manipulation [25] and muscle-in-the-loop control.

**Emerging Directions & Outlook** Building on the capabilities developed so far, my research agenda is expanding toward emerging directions that extend **fluid physical interaction** to richer and longer-horizon forms of human–robot collaboration. I am actively pursuing federal and medical funding (e.g., NSF, NIH, VA, AI-health programs) to support these efforts. A first direction focuses on **assistive and human neuromotor robotics**, where physical interaction is the primary mode of operation. With clinicians and biomedical researchers, I am extending my constraint-aware learning and control framework through what I term **muscle-in-the-loop control**, integrating bio-signal–driven estimation and control (e.g., EEG, EMG and ultrasound) to adapt robot behavior directly to human neuromotor dynamics. The goal is to develop robots that support, amplify, and adapt to human capability during rehabilitation and physical assistance. Second, I am exploring **strategic and rhythmic physical interaction**, in which robots engage with humans beyond short-horizon tasks. Activities such as competitive physical games and cooperative rhythmic motion (e.g., partner dance) require entrainment, multi-timescale intent prediction, and mutual adaptation. This would provide a principled testbed for understanding physical coordination, timing, and co-regulation under uncertainty; which translate to assistive robotics, gait rehabilitation, collaborative manipulation, and socially intelligent interaction. Finally, I am interested in developing **adaptive continual learning** for physical interaction, enabling robots to refine internal models and behaviors over long-term deployment without catastrophic forgetting or unsafe exploration.

Overall, my **long-term vision** is to establish a unified foundation in which learning, estimation, and control are co-designed around physical interaction, enabling robotic systems that are adaptive, trustworthy, and deeply human-centered.

# References

[1] Nadia Figueroa and Aude Billard. A physically-consistent bayesian non-parametric mixture model for dynamical system learning. In Aude Billard, Anca Dragan, Jan Peters, and Jun Morimoto, editors, *Proceedings of The 2nd Conference on Robot Learning*, volume 87 of *Proceedings of Machine Learning Research*, pages 927–946. PMLR, 29–31 Oct 2018.

[2] Y. Shavit, N. Figueroa, Seyed Sina Mirrazavi Salehian, and A. Billard. Learning augmented joint-space task-oriented dynamical systems: A linear parameter varying and synergetic control approach. *IEEE Robotics and Automation Letters*, 3(3):2718–2725, 2018.

[3] Nadia Figueroa, Salman Faraji, Mikhail Koptev, and Aude Billard. A dynamical system approach for adaptive grasping, navigation and co-manipulation with humanoid robots. In *2020 IEEE International conference on robotics and automation (ICRA)*, pages 7676–7682. IEEE, 2020.

[4] N. Figueroa and A. Billard. Locally active globally stable dynamical systems: Theory, learning, and experiments. *The International Journal of Robotics Research*, 41(3):312–347, 2022.

[5] A. Billard, Seyed Sina Mirrazavi Salehian, and N. Figueroa. *Learning for Adaptive and Reactive Robot Control: A Dynamical Systems Approach*. MIT Press, Cambridge, USA, 2022.

[6] Sunan Sun, Haihui Gao, Tianyu Li, and Nadia Figueroa. Directionality-aware mixture model parallel sampling for efficient linear parameter varying dynamical system learning. *IEEE Robotics and Automation Letters (RA-L)*, pages 1–8, 2024.

[7] Tianyu Li and Nadia Figueroa. Task transfer with stability guarantees via elastic dynamical system motion policies. In *7th Annual Conference on Robot Learning*, 2023.

[8] Farhad Nawaz, Tianyu Li, Nikolai Matni, and Nadia Figueroa. Learning complex motion plans using neural odes with safety and stability guarantees. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2024.

[9] Hadi Beik-Mohammadi, Søren Hauberg, Georgios Arvanitidis, Nadia Figueroa, Gerhard Neumann, and Leonel Rozo. Neural contractive dynamical systems. In *The Twelfth International Conference on Learning Representations (ICLR)*, 2024.

[10] F. Nawaz, S. Peng, L. Lindemann, N. Figueroa, and N. Matni. Reactive temporal logic-based planning and control for interactive robotic tasks. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2024.

[11] Yanwei Wang, Nadia Figueroa, Shen Li, Ankit Shah, and Julie Shah. Temporal logic imitation: Learning plan-satisficing motion policies from demonstrations. In *6th Annual Conference on Robot Learning*, 2022.

[12] Tianyu Li, Sunan Sun, Shubhodeep Shiv Aditya, and Nadia Figueroa. Elastic motion policy: An adaptive dynamical system for robust and efficient one-shot imitation learning. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2025.

[13] Shen Li, Nadia Figueroa, Ankit Shah, and Julie A. Shah. Provably Safe and Efficient Motion Planning with Uncertain Human Dynamics. In *Proceedings of Robotics: Science and Systems*, Virtual, July 2021.

[14] Christopher K. Fourie, Nadia Figueroa, and Julie A. Shah. On-manifold strategies for reactive dynamical system modulation with non-convex obstacles. *IEEE Transactions on Robotics (TRO)*, pages 1–20, 2024.

[15] Anusha Srikanthan, Yifan Xue, Vijay Kumar, Nikolai Matni, and Nadia Figueroa. Admm-mcbf-lca: A layered control architecture for safe real-time navigation. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2025.

[16] Yifan Xue and Nadia Figueroa. No minima, no collisions: Combining modulation and control barrier function strategies for feasible dynamical collision avoidance, 2025.

[17] Mikhail Koptev, Nadia Figueroa, and Aude Billard. Neural joint space implicit signed distance functions for reactive robot manipulator control. *IEEE Robotics and Automation Letters*, 8(2):480–487, 2023.

[18] Mikhail Koptev, Nadia Figueroa, and Aude Billard. Reactive collision-free motion generation in joint space via dynamical systems and sampling-based mpc. *The International Journal of Robotics Research (IJRR)*, 0(0):0, 2024.

[19] Zhiquan Zhang, Tianyu Li, and Nadia Figueroa. Constrained passive interaction control: Leveraging passivity and safety for robot manipulators. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2024.

[20] Yifei Shao, Tianyu Li, Shafagh Keyvanian, Pratik Chadhuari, Vijay Kumar, and Nadia Figueroa. Constraint-aware intent estimation for dynamic human-robot object co-manipulation. In *Robotics: Science and Systems*, 2024.

[21] Enrico Martini, Harshil Parekh, Shaoting Peng, Nicola Bombieri, and Nadia Figueroa. A robust filter for marker-less multi-person tracking in human-robot interaction scenarios. In *IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, 2024.

[22] Mateusz Jaszczuk and Nadia Figueroa. Rapid mismatch estimation via neural network informed variational inference. In *9th Annual Conference on Robot Learning (CORL)*, 2025.

[23] Shaoting Peng, Margaret X. Wang, Julie Shah, and Nadia Figueroa. Object permanence filters for robust tracking with interactive robots. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2024.

[24] Ho Jin Choi, Satyajeet Das, Shaoting Peng, Ruzena Bajcsy, and Nadia Figueroa. On the feasibility of eeg-based motor intention detection for real-time robot assistive control. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2024.

[25] Yifei Simon Shao, Tianyu Li, Shafagh Keyvanian, Pratik Chaudhari, Vijay Kumar, and Nadia Figueroa. Constraint-Aware Intent Estimation for Dynamic Human-Robot Object Co-Manipulation. In *Proceedings of Robotics: Science and Systems*, Delft, Netherlands, July 2024.