Research Statement: Physical and Perceptual Adaptive Intelligence for Fluid Human-Robot Collaborative Autonomy

Nadia Figueroa

December 10, 2024

For the last decades we have lived with the promise of one day being able to own a robot that can coexist, collaborate and cooperate with humans in our everyday lives. Recent advancements in low-cost hardware and accessible accelerated computing have driven the capabilities of robots to new levels, performing tasks outside of structured factory environments. Hence, the vision of a future with robots integrated in society, helping us in the household, in hospitals and at work, might seem attainable. Nevertheless, we are yet to see robots fluidly collaborating with humans and other robots in the human-centric dynamic spaces we inhabit. In fact, the robots currently being deployed in human close-proximity are designed to be highly conservative and risk-averse (avoiding any type of contact), slow and lightweight, stopping abruptly when humans are close, unable to communicate fluidly with humans or to handle heavy loads or exert high forces alongside or together with a human - rendering them inefficient and hard to use for non-experts. For robots to truly make a positive impact in our society and improve our quality of life, they have to be efficient, reactive, easy to communicate with and train, capable of predicting and avoiding unsafe situations in dynamic environments, as well as performing heavy-load physical assistive tasks in a safe, fluid and trustworthy manner. Achieving all of these behavioral objectives in a robot is challenging as no single optimization or data-driven approach is capable of encoding such nuanced and potentially contradicting goals and constraints in an efficient manner.

Towards Fluid Human-Robot Collaborative Autonomy Humans, on the other hand, have the remarkable capabilities to learn, adapt, and master new tasks while being efficient and safe when interacting with other humans, leading to a central question in my research. *Can we build robots that efficiently learn new tasks from humans and learn to safely interact with humans, just like humans do?* My research goal is, thus, to develop the **physical and perceptual adaptive intelligence** needed for robots to learn from and interact with humans while being able to adapt to a wide-range of capabilities and needs; from able-bodied humans that seek to teach robots to accomplish cumbersome dexterous manipulation tasks to mobility-impaired humans that require physical assistance in simple activities of daily living, like lifting and standing. In particular, **my research develops tightly coupled learning, control and robust estimation algorithms to achieve fluid human-robot collaborative autonomy with safety, efficiency and robustness guarantees.** My key insight is that fluid physical human-robot collaboration requires each of the processing units (perception, planning and control) used to govern the robot's behavior to possess two critical properties: i) be reactive and adaptive to perceptual, state changes and feedback from the connected processing layers while remaining safe and stable (Fig. 2) and ii) be aware of *perceptual, physical and cognitive* constraints of both the human and robot (Fig. 3). In order to develop these algorithmic advancements for the entire stack of the robot's processing units I leverage interdisciplinary approaches that combine insights from control theory, optimization, machine learning and cognitive science. During my PhD and postdoctoral work, I found that purely data-driven techniques lack rigorous guarantees, crucial for safety-critical tasks - in particular tasks with the human-in-the-loop. To address this, I devised creative approaches that step over traditional disciplinary boundaries, seamlessly blending concepts from control theory, robotics and machine learning [1]. Notably, I have been pioneering dynamical system (DS) based learning from demonstration (LfD), an imitation learning (IL) paradigm that assumes human and robot motion can be represented as an autonomous DS, rather than an open-ended feed-forward policy as in standard algorithmic IL strategies. Representing robot motion policies through DS enables the satisfaction of control-theoretic guarantees such as stability, convergence and

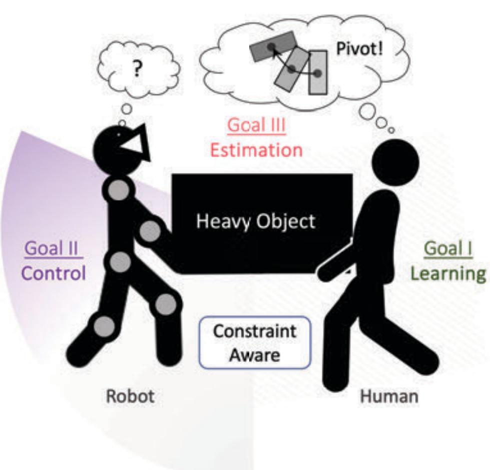

Figure 1: Towards Fluid Collaborative Human-Robot Autonomy through tightly coupled control, learning and estimation.

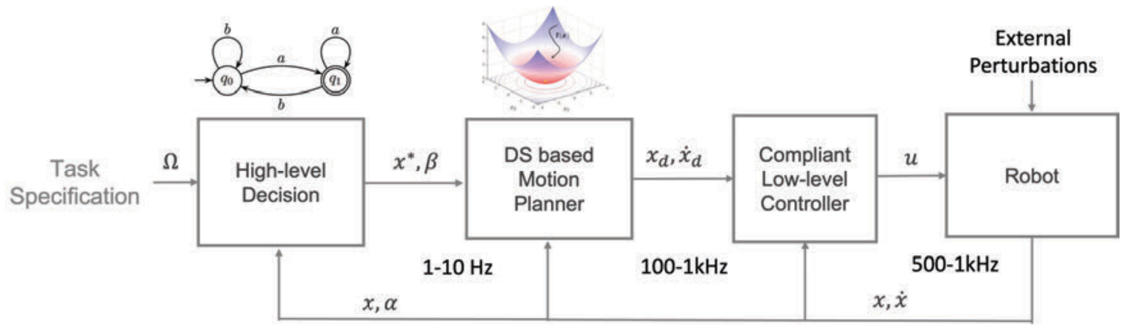

Figure 2: Proposed Reactive Framework for Robust Human-in-the-loop Robot Control

invariance [2, 3, 4, 5]. It also enjoys from sample and computational efficiency during training and inference. It can guarantee robustness to physical perturbations when coupled with impedance control laws and capable of modulating desired interaction forces, benefiting from energy-aware derivations and theoretical analyses that are fundamental to ensure safety and stability when humans and robots come in contact to achieve a shared task. The foundation of this work is detailed in my co-authored textbook, *"Learning for Adaptive and Reactive Robot Control: A Dynamical Systems Approach,"* published by MIT Press in February 2022 [6], which serves as the main textbook for my new graduate-level course.

Reactive Human-in-the-loop Control We have shown that introducing theoretical stability and convergence guarantees to robot motion policies is an elegant and efficient way to handle physical perturbations when performing tasks in the real dynamic world. Nevertheless, such theoretical guarantees are only met in practice at the motion planning layer, assuming that the robot has perfect state and environment knowledge at the required perception and control frequencies. Hence, in my current research, I propose to explicitly introduce feedback into all the processing layers of a robot's control architecture, as I depict in Fig. 2. Considering feedback in each of the robots processing units gives us a path towards achieving reactivity and adaptivity at all layers of control. However, to ensure safety, stability and convergence of the robot's behavior at all layers requires the introduction of: **(Goal I)** robot learning algorithms that take such feedback loops and interconnection between layers into account during training and execution, **(Goal II)** control algorithms that are capable of ensuring safety for both the human and the robot when they are physically interacting with the world and each other; and **(Goal III)** robust estimation algorithms that deal with faulty perception systems and noisy sensors to predict and estimate human, robot and environment states; all while being aware of both the robot and the humans perceptual, physical and cognitive **constraints**.

For the past three years, with my research group[1] and external collaborations, I have been working on addressing these challenges, with highly successful and recognized contributions summarized below.

Goal I – Efficient and Adaptive Robot Learning from and with Humans One of the most concerning limitations of the LfD approach is task generalization. While recent LfD approaches have tackled this by learning from multiple contexts, these approaches are data hungry and offer no stability and convergence guarantees. On the other hand, current DS learning techniques (used to) lack i) the expressivity to encode complex skills and ii) flexibility to generalize to new task instances as they ignore explicit task parameters that inherently change the underlying trajectories. Hence, in my group we began a quest to expand the capabilities of DS-based learning algorithms publishing our work at major robotics and robot learning venues: CORL, ICRA, IROS, ICLR and RA-L [7, 8, 9, 10, 11, 12, 13]. In terms of efficiency, we have developed parallellized statisical inference techniques [8] that achieve almost real-time learning performance of complex manipulation tasks, requiring from 1 to a handful of human demonstrations. In terms of expressivity, we have leveraged tools from differential geometry [10] to learn DS motion policies in $\mathcal{SE}(3)$ space, which is the proper mathematical manifold in which a robot's end-effector lives, showing that we can teach robots complex motion skills like opening doors, stacking plates in racks and pouring liquid into a bowl, all with robustness to perturbations (in both position and orientation). Further, we have been exploring the use of neural architectures such as neural ordinary differential equations [11] to encapsulate cyclical motions like stirring, polishing and wiping. While the control theoretic guarantees from all of these works have been derived from Lyapunov theory, we are also exploring the use of contraction theory to enforce stability in latent space representations of $\mathcal{SE}(3)$ tasks [12]. Furthermore, we have also tackled the problem of dynamically grasping objects while adapting to

[1]See my lab's youtube channel: `https://www.youtube.com/@figueroa-robotics-lab`

multiple potential grasps with a continuous energy-based $\mathcal{SE}(3)$ motion representation that is updated in real-time [13], achieving dynamic handovers from humans to robots while satisfying robot feasibility constraints. All of these methods are manifestations of the principle idea that being aware of the data and task constraints in our policy formulations leads to efficient learning. While introducing stability and convergence guarantees to motion policies is desirable, such guarantees no longer hold if the motion policy is used to perform a high-level task, such as transporting materials, or when the environments or context that the motion was demonstrated in change, leading to answer the following question: *Can we lift the stability and convergence guarantees to the high-level task execution layer to ensure recovery and successful task execution in the presence of failures or changes in the environment?* To tackle this, we have combined concepts of hybrid systems bisimulation with the motion-level stability guarantees and linear temporal logic specifications to i) detect failures in task execution, ii) learn from those failures and iii) recover from those failures. Such robustness to perturbations at the task and motion level had not been addressed before. Hence, warranting an oral presentation at CORL 22 with 6% acceptance rate [7]. A subsequent work explores the introduction of reactive temporal logic rules [14] to switch between learned DS motion policies triggered by perceptual changes in the task and human intention. We then addressed the generalization problem by properly embedding geometric descriptors of a task into a DS formulation, transfer to other task instances can be done in **less than 1 second** without needing more demonstrations while preserving desirable control-theoretic guarantees [9], ensuring safety and stability of both the low-level motion and the high-level task. Our most recent efforts involve real-time task transfer of a learned multi-step DS policy integrated with high-level semantic information, which we refer to as Elastic Motion Policy. I foresee this approach to have an impact on the robot learning community as we envision it to be used also for data augmentation to reduce human effort for current data-hungry learning strategies. Finally, we have also been exploring how to introduce such control theoretic guarantees into visuomotor policies with our fist attempt being an out-of-distribution recovery policy that brings the robot back to the training data manifold [15]. This work was invited for a Spotlight presentation (5/21 papers) at CoRL 2024 Workshop on Lifelong Learning for Home Robots.

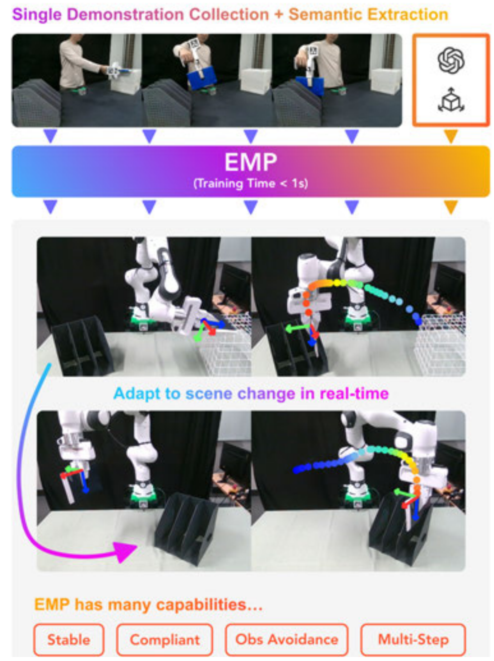

Figure 3: Elastic Motion Policies.

Goal II – Safe and Adaptive Control for Physical Interaction Safety in robotics can be defined from two perspectives, safety for humans or robots. A robot is safe for humans if it can either i) avoid unwanted collisions or ii) be compliant and passive to desired contacts; i.e., the robot absorbs and follows human physical guidance and external forces. So, which one is more important? During my postdoctoral work, I explored this question. I showed that allowing a relaxed definition of safety: avoid collision, but if collision is unavoidable then minimize impact [16], yields not only success in a complex task, like a robot dressing a human, but also the task is performed more efficiently than with purely collision avoidance requirements. Nevertheless, the correct approach depends on the task requirements, the robot capabilities, and expected behavior. Humans (and most animals) have an inherent intuition that triggers the decision of *when to comply?*, *when to avoid?* and *when to switch?*. At the core of this problem is *coordination*, we either coordinate with another agent *to avoid* it or *to comply* with it. Our recent efforts in this regard have been published in major robotics venues RA-L, IJRR, TRO and ICRA [17, 18, 19, 20] and summarize next.

The primordial task of guaranteeing safety for humans in close-proximity to robots is to *avoid unwanted collisions*. The problem of collision avoidance is a long-standing problem in robotics which becomes even more challenging to guarantee in dynamic environments with complex shaped obstacles (e.g., the human body or cluttered human environments). State-of-the-art approaches that ensure no collision or no local minima tend to be slow and not reactive, whereas purely reactive methods are prone to local minima and potential task failures (robot freezing problem). To this end, building on my PhD work on real-time robot collision avoidance using learning-based boundary functions, we have developed safe real-time algorithms that are capable of escaping local minima through a combination of sampling-based MPC and DS motion policies in high-dimensional robot's joint space, allowing safe real-time human-robot co-existence [17, 18]. Furthermore, in task space we have developed a theoretical framework to ensure no local minima when applying collision-aware modulation to DS motion policies. In particular, we introduce the notion of on-manifold modulation that leverages the isosurfaces created by differentiable boundary functions as manifolds to navigate on [19]. This method achieved real-time, 1kHZ, reactive safe-collision avoidance of highly cluttered dynamic workspaces in 3D space. We have recently extended this idea to the more general safety filter approach, control-barrier functions, to guarantee no local minima and ensure task success for real-time navigation of non-holonomic systems with tight input constraints, allowing humans to walk freely alongside robots.

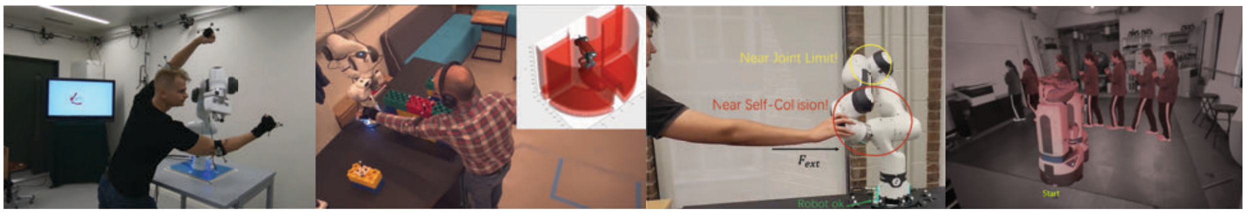

Figure 4: Real-time Joint Space Collision Avoidance [18], On-Manifold Workspace Modulation Strategy [19], Constrained Passive Interaction Control [20], Current work on Safe Real-time Human-Aware Navigation

In tasks where robots *are allowed* to make contact with humans (and viceversa) we must ensure that the robots are *passive and compliant* while also avoiding self-harm. Due to the unconstrained nature of passivity-based impedance control laws, the robot is vulnerable to infeasible and unsafe configurations upon unruly physical perturbations. Hence, we proposed a novel control architecture that allows a torque-controlled robot to guarantee safety constraints (such as kinematic limits, self-collisions, external collisions and singularities) and is passive only when feasible [20] through a real-time quadratic programming based approach. While successful, one of the main limitations of optimization-based torque control is the need for accurate dynamics models. This is particularly challenging to achieve when the robot is tasked to manipulate objects of unknown mass autonomously or together with a human. Adapting to unknown dynamics model mismatch errors through multi-modal sensing and providing feasible solutions to highly constrained optimization problems for real-time torque control are the current research directions of my group in this topic.

Goal III – Constraint-Aware Robust Estimation for Interactive & Assistive Tasks In order to successfully exploit the power of the adaptive motion policies and control algorithms developed in **(Goal I and II)** for interactive and assistive tasks, we need robust estimation algorithms that can reason about object permanence, object dynamics, theory of mind, motion, and human intention prediction. In fact, I argue that a redesign of state and intention estimation algorithms is needed to address the often ignored *perceptual and physical constraints* that arise in human-robot interactive scenarios. While *physical constraints*, like kinematic and dynamical limits are known beforehand for a robot they are particularly difficult to estimate for a human partner. Furthermore, *perceptual constraints* can come in the form of faulty or limited perception systems which may arise due to sensors with limited field of view, occlusion/self-occlusions or sensory noise hindering the proper estimation of the state/intent of the human. To tackle these issues, and with the assumption that perception will always be faulty, we have developed state estimation algorithms, which I refer to as permanence filters, for rigid multi-object [21] and human skeleton tracking [22]. These filters provide an informed estimate of predicted states aware of perceptual (occlusions) and physical (skeleton biomechanics), even in the presence of long-term occlusions, along with a measure of confidence to enhance estimation and control. They have been evaluated on challenging scenarios of severe occlusions and multi-person self-occlusions, and showcase superior performance over data-driven tracking algorithms.

Towards Physical Assistance When a robot and a human come in contact with each other through an object to accomplish a shared physical task, they form a tightly coupled (dyadic) dynamical system with each active sub-system of the dyad *differing in its perceptual, physical and cognitive capabilities.* For tasks such as co-manipulating a heavy or large object, intent estimation (be it at the action, motion or task level) is useless for a controller if it is not feasible for both the human and the robot to perform. Hence, to ensure that the robot controller is receiving a feasible command in such scenarios, we developed a tightly coupled estimation and control algorithm that is aware of and can guarantee both human and robot's *internal and*

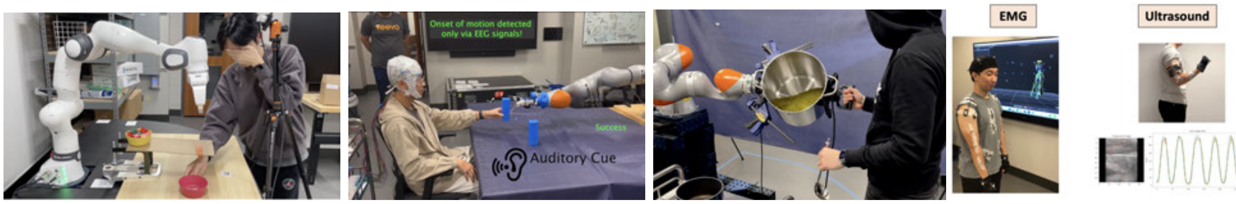

Figure 5: Object permanence filter tracking occluded objects with no disruptions [21], Real-time intent estimation from surface EEG (electroencelography) signals [24], Constraint-aware intent estimation and control [23] and most recent unpublished work on estimation of human joint angle, torque and exercise completion rates from surface EMG (electromyography) and wearable ultrasound sensors.

external physical constraints. Constraints that include kinematic limits and collision avoidance constraints; while seamlessly adapting the assistive force based on the confidence of the intent estimation [23]. When a robot and human are physically interacting, there is an exchange of energy and power, this produces not only vulnerabilities to kinematic limits but also to dynamic limits. Thus, besides robust estimation capabilities, a fluid and safe regulation of interaction forces and power exchange of the coupled dynamics is necessary in order to ensure stable and fluid physical interaction, in particular for high-load safety-critical tasks. Furthermore, when interacting with a human that has mobility issues, tracking purely kinematics is insufficient for proper state or intent estimation. We need to be able to disambiguate a person's intent with their capability to actually achieve that task; so that the robot can adapt its physically assistive forces accordingly. My current pursuit to achieve this is by tracking human biosignals, in real-time, through wearable sensing technologies. Via active collaborations with colleagues from the Penn Engineering and the Penn Medical school we have developed real-time state and intent estimation processing pipelines from surface EEG (electroencelography) signals [24], surface EMG (electromyography) and US (ultrasound) sensors. For our first foray in this research direction we evaluated the feasibiliy of using non-invasive EEG signals to detect a human's intent to move left or right arm [24].; achieving real-time performance coupled with a robot assistant providing an object to the closest predicted arm. Most recently we have been developing sparse neural network based processing pipelines to use Ultrasound images and EMG sensors to predict in real-time joint angle, torques and exercise completion rates for isolated joints when doing heavy-lifting exercises like bicep curls and shoulder lifts (Fig. 5). Through these developments we are currently waiting for the IRB approval of a protocol for a user study in which our US processing pipeline is coupled to an exercise assistive robotic arm, as depicted in Fig. 6. The goal of this study is to validate our novel muscle-in-the-loop robotic exercise assistant, which we envision can be used in a range of applications from augmented high-performance training to physical therapy and muscle rehabilitation. These efforts have led me to build strong collaborations with faculty from the Department of Physical Medicine and Rehabilitation, as showcased by one of our recent published works on learning realistic joint space range of motion boundary functions for diagnosis, monitoring and robotic rehabilitation therapy interventions of stroke survivors [25].

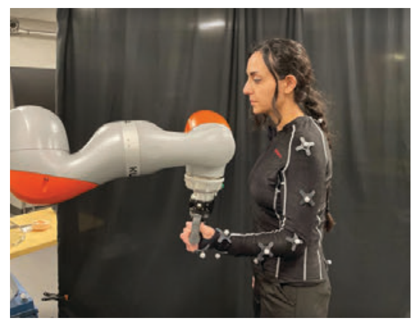

Figure 6: Muscle-in-the-loop adaptive robotic physical therapy.

Summary and Future Outlook Over the past three years at Penn, I have focused on developing foundational algorithms for learning **(Goal I)**, control **(Goal II)**, and estimation **(Goal III)** to enable physical and perceptual adaptive intelligence for fluid collaborative human-robot autonomy. Looking ahead, my goals remain centered on addressing the challenges of physical human-robot interaction. I aim to leverage control-theoretic and geometry-informed approaches, inspired by principles from physics, biology and psychology, to realize truly collaborative robots that are both safe and efficient for humans. As robots increasingly engage with our human-centric world, understanding *how to control robots to maintain stable physical interactions with dynamic human partners* will have broader applications, such as enabling robots to collaborate with other robots or perform contact-rich manipulation tasks in safety-critical scenarios. The outcomes of my research are expected to drive significant advancements in both the theoretical foundations and practical applications of robotics. These advancements have the potential to transform industries such as manufacturing, healthcare, and service robotics, while setting new benchmarks for future research in developing physically adaptive robots designed to interact seamlessly with humans and their environments.

References

[1] Nadia Figueroa. *From High-Level to Low-Level Robot Learning of Complex Tasks: Leveraging Priors, Metrics and Dynamical Systems*. PhD thesis, EPFL, Lausanne, Switzerland, 2019.

[2] 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.

[3] 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.

[4] 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.

[5] 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.

[6] 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.

[7] 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.

[8] 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.

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

[10] S. Sun and N. Figueroa. Se(3) linear parameter varying dynamical systems for globally asymptotically stable end-effector control. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2024.

[11] 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.

[12] 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.

[13] Ho Jin Choi and Nadia Figueroa. Towards feasible dynamic grasping: Leveraging gaussian process distance fields, se(3) equivariance and riemannian mixture models. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2024.

[14] 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.

[15] George Jiayuan Gao, Tianyu Li, and Nadia Figueroa. Out-of-distribution recovery with object-centric keypoint inverse policy for visuomotor imitation learning, 2024.

[16] 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.

[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] 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.

[20] 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.

[21] 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.

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

[23] 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.

[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] Shafagh Keyvanian, Michelle J. Johnson, and Nadia Figueroa. Learning realistic joint space boundaries for range of motion analysis of healthy and impaired human arms. In *2023 IEEE-RAS 22nd International Conference on Humanoid Robots (Humanoids)*, pages 1–8, 2023.