

Nadia Figueroa: Research Statement

From factories to households, we envision a future where robots can break out of their cages and work *safely* and *efficiently* alongside humans in a myriad of tasks (Figure 1). Such tasks require robots to be agile, dexterous, safe, and robust while manipulating objects and closely interacting with humans. For robots to be fully adopted in such scenarios, we must i) minimize human effort while communicating and transferring tasks to robots; and ii) endow robots with the capabilities of adapting to changes in the environment, in the task objectives, and the human’s intentions while guaranteeing task-success. This is challenging for robots, as providing a single optimal solution for these objectives can be intractable and even infeasible due to problem complexity and contradicting goals. However, humans have a remarkable way of guaranteeing all of these objectives and learning, adapting, and master new tasks, leading to a central question in my research. *Can we build robots that learn tasks from humans and learn to interact with humans, just like humans do?*

To equip robots with such *human-like* capabilities, I seek to unify robot learning and control strategies that provide safe and fluid physical human-robot-interaction (pHRI) while theoretically guaranteeing task success and stability. I achieve this by devising techniques that step over traditional disciplinary boundaries, seamlessly blending concepts from control theory, robotics and machine learning. Specifically, I develop methods that leverage machine learning techniques with concepts from dynamical systems (DS) theory to tackle core robotics problems in motion planning and manipulation – central to any workplace, industrial, or household tasks.

My research focuses on *learning tasks as motion policies from data* (be it from humans, simulation, or optimization techniques). These motion policies are formulated such that they are *adaptive* to human perturbations (or intentions), changes in the environment, or even changes in the task objectives or constraints. A motion policy is a function $f(x) : \mathbb{R}^N \rightarrow \mathbb{R}^N$ that approximates a direct mapping from robot state observations to the robot actions required to achieve the motion plan and goal; with $x \in \mathbb{R}^N$ being the state of the robot. In my work [1], to generate motions robust to perturbations, the motion policy is formulated as,

$$\dot{x} = f(x, \Theta) \quad \text{with} \quad \lim_{t \rightarrow \infty} \|x - x^*\| = 0, \quad (1)$$

a first-order DS, parametrized by Θ , that converges to a target $x^* \in \mathbb{R}^N$. The mathematical form of $f(\cdot)$ can be either a probabilistic model or a deterministic mapping function, as long as the parameters Θ guarantee convergence to the target [1]. To find Θ , we follow the learning from demonstration (LfD) paradigm.

Given a dataset representing the task, we solve a constrained optimization problem that minimizes a performance objective for (1) while ensuring stability constraints. I have demonstrated that if Θ is learned or properly defined to guarantee convergence, a single motion policy as (1) can be used to solve a number of robotics problems. I have introduced several formulations of (1) that can execute *simple point-to-point* motions, such as pick-and-place and imitating motion patterns [5, 7] to more *dynamic scenarios*, such as object hand-overs and catching objects in flight [14, 15, 8]. Motion policies like (1) are a powerful tool to provide inherent physical *adaptability*. Yet, they have been limited to: a) a single target x^* , b) a fixed physical environment $e \in \mathcal{E}$, and c) a fixed set of task constraints $c \in \mathcal{C}$ observed during teaching.



Figure 1: Workplace, industrial, household tasks learned and analyzed during my doctoral research @ EPFL, Switzerland [3]

To provide *adaptation at every level* and *fluid interactive behavior* we need motion policies as,

$$\dot{x} = f(x, \mathbf{c}, \mathbf{e}, \mathbf{x}^*, \Theta) \text{ with } a) \lim_{t \rightarrow \infty} \|x - \mathbf{x}^*(x)\| = 0, b) 0 \leq \mathbf{c}(x) \leq 1, c) \|x - \mathbf{e}(x)\| > 0 \quad (2)$$

where $\mathbf{x}^*(x) \rightarrow x^*$, $\mathbf{c}(x) \rightarrow c$, $\mathbf{e}(x) \rightarrow e$ can be mapping functions that generate the desired target, constraint to guarantee, and obstacle or external agent to avoid. I posit that introducing these additional objectives to the motion policy formulation will yield the *adaptive*, *interactive* and *safe* robot behaviors necessary for real human-robot collaboration. My future research will focus on generating such motion policies that will be *aware* of the robots' environment, with parameters Θ that are *grounded* to the physical constraints of a task (e.g., follow the surface, reach an intermediate target, etc.); while avoiding collisions.

To perform tasks where multiple targets, objectives, robots, and agents are involved; we can combine motion policies like (2) into sequences or networks of inter-connected *modular* motion policies. Guaranteeing stable interactions between modular motion policies and discovering the optimal granularity of these modules for different tasks is an interesting open challenge that can have applications in other domains. Finally, in order to successfully adopt motion policies like (2) in pHRI tasks, one must either have fully observable states of x, c, e or account for uncertainty in these states. Such robust state estimation is difficult to achieve with the current perception systems in pHRI scenarios where humans, robots, and objects interact and unintentionally occlude each other from perception systems.

Research to date

From High-level to Low-level Robot Learning of Complex Tasks

During my Ph.D., I focused on three research themes that allowed robots to learn tasks from high-level to a low-level perspective, from task to joint space, and from single arm to high-dimensional humanoid robots.

Learning Complex Sequential Tasks from Demonstration

A central goal of my Ph.D. thesis was learning complex sequential manipulation tasks from heterogeneous and unstructured demonstrations. Such tasks are those composed of a sequence of discrete actions. The particular challenge is learning without any prior knowledge of the number of actions (unstructuredness) or restrictions on how the human is demonstrating the task (heterogeneous). We proposed a Bayesian non-parametric learning framework that can jointly segment and discover unique discrete actions in continuous demonstrations of the task using a novel distance metric on the space of covariance matrices [6]. The learned structure of the complex tasks was then used to parametrize a sequence of motion policies (1) to execute two cooking tasks: (i) single-arm pizza dough rolling, and (ii) a dual-arm vegetable peeling task [4, 10].

Learning DS based Motion Policies from Demonstrations

I proposed a DS-based motion policy formulation and a Bayesian non-parametric learning scheme capable of automatically encoding continuous and complex motions while ensuring global asymptotic stability. The type of tasks that can be learned with this approach are unparalleled to previous work in DS-based LfD and are validated on production-line and household activities [5]. Further, I proposed a novel DS formulation and learning scheme that can encode both complex motions and varying impedance requirements along the task [7]. This approach was validated on trajectory tracking tasks where a robot arm must precisely draw letters and shapes on a surface. I also applied this formulation to other robotic domains, such as (i) navigation strategies for mobile agents and (ii) locomotion and co-manipulation tasks of bipedal humanoid robots [8].

Multi-Robot Arm Coordination and Self-Collision Avoidance

When working with multiple agents, coordination is essential for fluid and safe task execution. We endowed a set of robot arms with these *human-like* coordination capabilities by proposing i) a *task-space multi-arm coordination* motion policy [14, 15] and ii) a data-driven *joint-space self-collision avoidance* strategy [14, 15, 9]. The motion policy synchronizes the end-effectors of the robot arms to track and grasp a moving object carried by a human [14, 15]. To avoid self-collisions, I proposed to learn a self-collision boundary by sampling the workspace of the multiple robots and used it as a constraint for an optimization-based inverse kinematics solver [14, 15, 9]. We applied these control strategies in multi-arm to human-operator hand-overs of large objects, such as car parts [14, 15, 9]. With this work, we competed in the KUKA Innovation Award 2017 as one of the Top 5 finalists. We also won the award for "Best Student Paper" and were finalists for "Best Conference Paper" and "Best Systems Paper" finalist at the Robotics:Science and Systems 2016 conference for [14]. We recently extended these strategies to humanoid robots in [8, 13].

Towards Task Adaptability & Efficient Learning

As a postdoctoral associate at MIT, develop techniques that will allow robots to *adapt* to small changes and physical perturbations, as well as to higher-level changes in the task objectives/targets and constraints. In order to successfully exploit the power of adaptive motion policies as (2), we need robust perception systems that can reason about object permanence, object dynamics, theory of mind, motion and intention prediction. To tackle this issue, I have been developing a multi-modal teaching architecture that leverages data from cameras, IMU sensors, natural language instructions and physical interaction with the robot, to build a probabilistic mental model representation of the scene. The goal is to account for failures in perception systems, such as loss of tracking and malfunctioning, as well as unintentional occlusions by humans or other agents in the scene. I am leading a team of undergraduate and master-level researchers that are tackling different aspects of this project including: deep-learning based object tracking, multi-modal sensor fusion for robust hand motion tracking, grounded task goal learning and building robot mental models.

Future Directions

I plan to continue developing sound mathematical formulations and learning schemes to achieve *fluid*, *safe* and *efficient* human-robot interaction. While my previous research has focused on providing *adaptability* during execution and *autonomy* during learning; in my future work I would like to additionally focus on providing *safety* and *efficiency* during interaction and learning, as well as *scalability*.

Safety in Physical Human Robot Interaction and Control

Safety in robotics can be defined from two perspectives, safety for humans or robots. A robot is safe around humans if it can either i) avoid collisions or ii) be compliant/passive; i.e., the robot absorbs and follows human physical guidance and external forces. i) is tackled with motion planning strategies, and ii) are feedback control-theoretic approaches. The approach used 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. In this research theme, I will seek to understand how humans (and/or animals) coordinate and switch between these strategies to propose biologically-inspired strategies for robots. From the *robot perspective*, avoiding collisions is the *safe* strategy, as compliance can allow for perturbations that may lead to self-damage. One could *comply while avoiding self-damage*, by leveraging optimal motion planners with control-theoretic compliant approaches.

Efficiency in Human Robot Interaction and Learning

Efficiency can also be defined from many perspectives. We can strive for data-efficient or computationally-efficient learning schemes or having efficient communication/interaction between humans and robots. *Efficiency in learning* will be achieved by branching out from the LfD paradigm and exploring other learning strategies, such as reinforcement learning or incremental learning. Regarding interaction with humans, humans are a generally lazy species, hence, the least time, effort and mental load that it takes for us to interact with a robot the better. I define this concept as *efficient interaction*. The goal of *efficient interaction* is to reduce the burden on the human to transfer tasks to robots, communicate goals, and understand the robot's intention or the expected controller behaviors. To provide fluid/efficient communication between humans and robots, we must exploit the vast amount of information available from different sensory inputs (as humans do). A research direction in my lab will focus on developing multi-modal learning/interaction strategies, guided by control-theoretic models, for efficient detection/prediction of human intentions.

Towards Modular Motion Policy Networks

The concept of *modularity* has been used in many domains. From solving complex reasoning AI problems [2] to understanding biological phenomena like the collective motion of animals [12] or the interaction of multi-scale networks of genes and proteins [11], it has been shown that the most effective approach for solving a complex task is solving a system of independent yet complementary sub-tasks. Not only is this evidenced in research, but it also has a biological basis, as we know that there are functionally specialized regions in the human brain that are tailored to solve specific cognitive processes. One could coordinate and control multiple robots to achieve an overall joint goal or even control complex high-dimensional robotic systems by decomposing them into simpler parts guided by modular interacting motion policies. This bring solutions to the challenges induced by the complex geometry of controlling high-dimensional interacting robotic systems.

Impact and Funding

My research goals are aligned with international, governmental, and industry efforts to advance the development and use of collaborative robots. The National Science Foundation (NSF) offers several funding opportunities for robotics research, such as the National Robotics Initiative (NRI)-2.0 Ubiquitous Collaborative Robots program and the FFR: Foundational Research in Robotics program. Robotics funding is also available from the USDA, NASA, and the NIOSH. Leading AI companies such as Amazon, Microsoft, and Google offer yearly research awards to develop AI-powered robots. Automation-focused companies, such as Toyota, ABB, Mitsubishi, always seek academic collaborations. Notably, my postdoctoral appointment at MIT is funded by a collaborative effort with Mitsubishi Electric Automation, Inc. Hence, there are increasing opportunities for my research to be financed by both governmental agencies and industrial collaborations.

References

- [1] A. Billard, S. S. Mirrazavi Salehian, and N. Figueroa. *Learning for Adaptive and Reactive Robot Control: A Dynamical Systems Approach*. MIT Press, Cambridge, USA, 2022. Book to be published within the Intelligent Robotics and Autonomous Agents Series in Fall 2021.
- [2] K. Chen. *Deep and Modular Neural Networks*, pages 473–494. Springer Berlin Heidelberg, Berlin, Heidelberg, 2015.
- [3] N. Figueroa. *From High-Level to Low-Level Robot Learning of Complex Tasks: Leveraging Priors, Metrics and Dynamical Systems*. PhD thesis, EPFL, Lausanne, Switzerland, 2019.
- [4] N. Figueroa and A. Billard. Learning complex manipulation tasks from heterogeneous and unstructured demonstrations. In Proceedings of Workshop on Synergies between Learning and Interaction. IEEE/RSJ International Conference on Intelligent Robots and Systems, 2017.
- [5] N. Figueroa and A. Billard. A physically-consistent bayesian non-parametric mixture model for dynamical system learning. volume 87 of *Proceedings of Machine Learning Research*, pages 927–946. PMLR, 29–31 Oct 2018.
- [6] N. Figueroa and A. Billard. Geometric invariance of covariance matrices for unsupervised clustering, segmentation and action discovery in robotic applications. In preparation for re-submission to the Journal of Machine Learning Research (JMLR), 2019.
- [7] N. Figueroa and A. Billard. Locally active globally stable dynamical systems: Theory, learning and experiments. Under review for publication in The International Journal of Robotics Research, 2019.
- [8] N. Figueroa, S. Faraji, M. Koptev, and A. 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, 2020.
- [9] N. Figueroa, S. S. Mirrazavi Salehian, and A. Billard. Multi-arm self-collision avoidance: A sparse solution for a big data problem. *Proceedings of the ICRA 2018 Third Machine Learning in Planning and Control of Robot Motion (MLPC) Workshop*, page 6, 2018.
- [10] N. Figueroa, A. L. P. Ureche, and A. Billard. Learning complex sequential tasks from demonstration: A pizza dough rolling case study. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 611–612, 2016.
- [11] A. Groß, B. Kracher, J. M. Kraus, S. D. Kühlwein, A. S. Pfister, S. Wiese, K. Luckert, O. Pötz, T. Joos, D. Van Daele, L. De Raedt, M. Kühl, and H. A. Kestler. Representing dynamic biological networks with multi-scale probabilistic models. *Communications Biology*, 2(1):21, 2019.
- [12] F. J. H. Heras, F. Romero-Ferrero, R. C. Hinz, and G. G. de Polavieja. Deep attention networks reveal the rules of collective motion in zebrafish. *PLOS Computational Biology*, 15(9):1–23, 09 2019.
- [13] M. Koptev, N. Figueroa, and A. Billard. Real-time self-collision avoidance in joint space for humanoid robots. Under review for publication in the IEEE Robotics and Automation Letters (RA-L) with ICRA 2021 Option., 2020.
- [14] S. S. Mirrazavi Salehian, N. Figueroa, and A. Billard. Coordinated multi-arm motion planning: Reaching for moving objects in the face of uncertainty. In *Proceedings of Robotics: Science and Systems*, Ann Arbor, Michigan, June 2016.
- [15] S. S. Mirrazavi Salehian, N. Figueroa, and A. Billard. A unified framework for coordinated multi-arm motion planning. *The International Journal of Robotics Research*, 37(10):1205–1232, 2018.