

Design and analysis of complex architectures

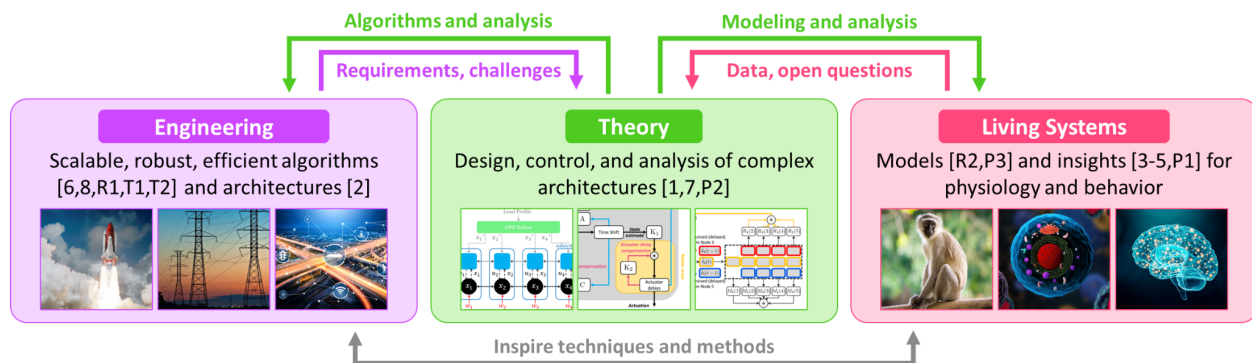
From control theory to biological and engineering systems

In engineering, control theory plays a crucial role in the design and analysis of robust and efficient systems – including robots, spacecraft, and power grids. In biology, control theory underlies sensorimotor models of organisms, and also has applications as a modeling tool for immune systems and metabolic networks. As an engineer and control theorist, I am interested in the analysis of biological systems, both for its standalone value and for the insights biology can provide for engineers and theorists.

Biology provides exceptional examples of robust, efficient architectures and behaviors. In order to stay alive, organisms must perform complex tasks in highly dynamic and unpredictable environments (e.g. evade predators, maintain homeostasis). Modern engineering systems cannot achieve comparable performance for even simple subsets of these tasks. For example, in the task of legged locomotion and balance, organisms are much more robust and efficient than their robotic counterparts. This is especially striking since biological components (e.g. neurons, cells) are in many ways less ideal than their electronic counterparts (e.g. transistors, chips) – cells in a human communicate much more noisily and slowly than chips in a robot. Clearly, there is much to be explained in biology, and much we can learn from biology as engineers. ***My research vision is to develop and apply theory to close the knowledge and performance gap between biology and engineering,*** with an emphasis on the analysis of complex, scalable systems. The three foundational directions of my research are:

- (1) Develop **theory** for the design, control, and analysis of complex architectures
- (2) Apply theory from (1) to develop models and provide insights for **living systems**
- (3) Apply theory from (1) to create scalable, robust, efficient **algorithms and architectures**

These three directions, as well as how my past work fits into them, are summarized in the figure below:



Theory: Design and control of complex architectures

Problems in **engineering** (e.g. how do we control a renewables-heavy smart grid?) and questions from **living systems** (e.g. how do neurons coordinate to provide robust sensorimotor behavior?) present new, substantial challenges to existing control theory. The aim of this research direction is to **establish unified theory on how to systematically design robust, efficient architectures and systems.**

Past work

When analyzing complex systems, it is necessary to make simplifications for tractability – however, simplified models should preserve the key features of the system that give it its distinctive behavior and architecture. I focus on features relating to distributed computation. ***These features are omnipresent in living systems and relevant for scalable engineering algorithms.*** They include:

- (a) **Local communication.** Neurons and cells communicate locally – instead of communicating to the entire set of neurons and cells in the body, they typically communicate with some set of nearby neighbors. In engineering applications, local communication can facilitate scalability and efficiency [6,R1].

- (b) **Localized behavior** is typically carried out by reflexes, which reject local disturbances without affecting the entire system. For example, touching a hot frying pan elicits a reflexive retraction of the hand, without affecting the rest of the body. This fast, reflexive action is elicited by local circuits in the arm and spinal cord, and occurs before we become conscious of pain. In engineering applications, localized behavior is desirable for large systems.

Despite their prevalence in living systems, these features are under-examined by theory. My past work provides theoretical analyses on these two features. In [P2], we show that under very mild assumptions, local communication constraints in distributed model predictive control [6,R1] induce no suboptimality compared to global communication. Thus, we can facilitate scalability and energy efficiency without compromising performance. In [1], we leverage the system level synthesis (SLS) parametrization to show how local communication and localized behavior can be incorporated as objectives or constraints in the controller optimization problem. We also show that for stable closed-loop systems, these two forms of localization are coupled; thus, these two constraints can be synergistic with one another.

Nonlinearity is another feature that is important in engineering and living systems – in [7], we use SLS to design nonlinear controllers that generalize and outperform the popular feedback linearization technique. This opens the door to more scalable formulations for nonlinear control.

Future work

Design of stabilizing local controllers. Theoretical understanding of this topic will (1) supplement our understanding of living systems and (2) complement data-driven techniques in engineering. Stabilizing local reflexes allow organisms to safely learn and explore. For instance, a child may touch a hot frying pan out of curiosity, but pain reflexes will cause them to retract their hand immediately, minimizing injury (i.e. promoting safety). This experience generates learning signals, which teach the child to avoid touching hot frying pans in the future. Analogously, in engineering, stabilizing local controllers offer a platform upon which data-driven strategies can operate and explore without sacrificing system safety.

Controller implementation using constrained components. Theoretical understanding of this topic will (1) allow us to better analyze biological circuits, which utilize constrained components and (2) guide engineers toward more energy-efficient implementations and components. Controller implementation typically makes use of realization theory, which translates between state-space and input-output models. However, this theory does not incorporate component-level constraints, which dominate neural and cellular implementations of control circuits – neurons and cells are highly limited in terms of bandwidth, speed, and signal-to-noise ratio. Component-level constraints in living systems are a result of extreme energy efficiency requirements; thus, this theory can also guide engineers toward more energy-efficient components and controller implementations.

Living systems: Novel models and insights for physiology and behavior

I aim to use **control and architectural theory** to derive models, principles, and insights for physiology and behavior. Where applicable, I use **scalable algorithms** to model large living systems in more detail. My research at the intersection of distributed control and sensorimotor modeling earned me the **competitive PGSD funding award from the Natural Sciences and Engineering Research Council of Canada (NSERC)**. An eventual aim of this research direction is to **create models that have diagnostic and therapeutic impact** – for instance, sensorimotor models that inform therapeutic directions for degenerative diseases. Here, collaborations with domain experts play an essential role in shaping problems that have meaningful scientific impact.

Past work

Internal feedback in the sensorimotor system (collaboration with Dr. Terry Sejnowski, Salk Institute). This work offers the first explanation for the large quantities of internal feedback observed in the cortex. The standard model of sensorimotor processing involves signal flows from sensory inputs to sensory areas, then from sensory areas to motor areas. However, massive amounts of signal flow in the opposite

direction (i.e. from motor areas back toward sensory areas) are observed in the cortex; we refer to these signals as *internal feedback*. What function does internal feedback serve, and why does the brain contain so much of it? We answer these questions using control theory in [3-5]; in particular, [5] argues that the incorporation of local communication and localized behavior necessitates large amounts of internal feedback. This paper was a **Best Student Paper Finalist at the 2022 American Control Conference**. In [P1], we build on these works and further integrate theory with physiology (e.g. cell types and experimental observations) for neuroscience audiences.

Layered model for robust multi-legged locomotion (collaboration with Dr. Bing Brunton, University of Washington) [P3]. This work produces the first model that integrates realistic kinematics with leg dynamics for *Drosophila melanogaster* (fruit fly). We use control theory in conjunction with an artificial neural network to create a layered model of locomotion. Using a novel controller formulation from previous work [4], we are also able to model and study muscle delays in the leg. This end-to-end model provides a unifying, organism-agnostic framework which brings together multiple aspects of legged locomotion: realistic gait patterns, multi-leg coordination, and physiology-based dynamics.

Future work

Distributed control models for cancer immunotherapy (collaboration with Dr. Peter Lee, City of Hope). This project aims to provide models and therapeutic directions for cancer and the immune system. Immune cells use local and distributed mechanisms to activate and suppress cytotoxic T cells, which play a central role in adaptive immunity. These distributed mechanisms allow the body to mount rapid and effective responses, but can also lead to dysfunctional states, ranging from autoimmunity to cancer. Principles from distributed control can be applied to model and better understand the types of perturbations (e.g. wounding) that lead to dysfunctional states, as well as how we may perturb the system (e.g. via medical interventions) to escape these states. We are particularly interested in the role of immune checkpoints and how they can be leveraged to provide effective therapies without inducing autoimmunity.

Large-scale models of metabolic dynamics (collaboration with Dr. Fangzhou Xiao, UCSD). This project aims to provide large-scale models of whole cells, as well as microbial communities and their interactions with human bodies. Analyses of large-scale chemical reaction networks, including metabolism, often make steady-state assumptions – the analysis of large-scale temporal dynamics is an open question. In recent work led by Dr. Xiao [R2], the metabolic network is partitioned into a dynamical system and a controller. The dynamical system is determined by the stoichiometry matrix with mass-action dynamics; the controller governs the sensitivity of reaction fluxes to metabolite concentrations. We can apply optimal, predictive, or distributed control to provide dynamical trajectories for metabolites of interest. Early work using model predictive control captures oscillations in glycolysis and growth arrest of cells under stress, in small networks. Future work aims to utilize scalable model predictive control from [6,R1] to provide larger models for whole cells and microbial communities.

Overall, my work in this area is heavily driven by collaboration with neuroscientists and biologists. In addition to the topics described, I am excited to work with local experts on diverse questions in sensorimotor control and broader topics of their choosing.

Engineering: Scalable algorithms and architectures for large-scale networks

Many modern engineering problems involve complex, large-scale dynamical networks. Leveraging advancements in *control and architectural theory* and findings from *living systems*, I aim to **establish a comprehensive set of tools that allow for systematic design of scalable, robust, and efficient architectures and algorithms**.

Past work

My existing work focuses on distributed, scalable control algorithms. These algorithms enjoy **complexity that scales independently of network size**, under mild assumptions. This is important for large-scale systems, for which centralized controllers can be intractable. I often use the system level synthesis (SLS)

framework as a basis for building scalable algorithms. I have also written MATLAB and Python toolboxes for SLS [T1,T2], which include plug-and-play code for the algorithms described below.

- (1) **Scalable robust control** [8]. Robust control provides stability and performance guarantees in the presence of model uncertainties; this is applicable in conjunction with data-based modeling methods, and was historically vital to the development of aircraft. We propose a scalable algorithm for robust control synthesis subject to structured uncertainty.
- (2) **Scalable model predictive control (MPC)**. MPC is widely used across engineering applications, and also has applications in biological modeling, as described in the previous section. In [6], we propose a scalable robust MPC algorithm which can be both synthesized and implemented in a distributed manner. We derive feasibility and stability guarantees for this algorithm in [R1] by synthesizing a terminal set and its associated Lyapunov function. This work is the first to provide an exact, distributed computation of these quantities, which allows us to make theoretical guarantees without introducing conservatism and performance degradation.
- (3) **Layered controllers for efficiency and performance** [2]. This work illustrates an important principle; by appropriately combining algorithms in layered architectures, we can mitigate the weaknesses of individual algorithms and produce scalable, efficient, and high-performing behavior. We consider a power grid-inspired system with changing optimal flows, and apply a distributed two-layer controller consisting of online and offline control. While offline control alone is insufficient and online control alone is computationally heavy, the layered controller produces near-optimal performance with a 20-fold reduction in computational load.

Future work

Systematic design of layered controllers. This topic has high potential for advancing performance, efficiency, and safety of control algorithms. Multiple algorithms can be integrated into one layered algorithm, which takes on the best aspects of its constituent algorithms [2,P3]. Rather than creating a brand new algorithm or tuning existing algorithms for small improvements, we can cleverly combine existing algorithms for large improvements. I am interested in developing techniques that allow for systematic integration of various control algorithms in this manner. I am particularly interested in the use of stabilizing local controllers (described on pg.1) and how they can be combined with adaptive algorithms to create safe and efficient behavior.

Sparsity for efficient perception (collaboration with Ivan D. Jimenez Rodriguez, Caltech). Natural scenes, particularly those perceived by autonomous vehicles (e.g. cars, drones), are dominated by predictable, slow-changing elements (e.g. trees, buildings, sky). For camera-based perception, new information introduced between consecutive frames is typically highly sparse. Visual systems of organisms take advantage of this sparsity; predictive feedback is used to filter out predictable portions of visual signals. This considerably lowers the processing burden required for perception. Inspired by this, we seek ways to leverage sparsity for efficient perception. We are especially interested in using neural ordinary differential equations (Neural ODEs), which can approximate underlying dynamics of predictable portions of the scene. Neural ODEs can be combined with techniques from distributed control to produce sparse representations which are amenable to efficient, parallelizable computation.

In addition to the listed topics, I am eager to work with local engineering collaborators to apply scalable algorithms to real-life engineering problems, and also to develop algorithms tailored for specific applications and systems.

Funding

In addition to typical early-career funding (NSF-CAREER/PECASE, DoD Young Investigator), I plan to apply for a variety of grants, including DoD Multidisciplinary University Research Initiative (MURI), NSF Emerging Frontiers in Research and Innovation (EFRI), and NSF Integrative Research in Biology (IntBIO); NSF Dynamics, Control and Cognition (DCC), NSF Understanding the Rules of Life: Emergent Networks (URoL:EN), NSF Collaborative Research in Computational Neuroscience (CRCNS), etc.