Interpretable Learning and Control of Physical Dynamical Systems

by

Alexander Ethan Cohen

Submitted to the Department of Chemical Engineering
and the Center for Computational Science and Engineering
on December 18, 2025 in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Chemical Engineering and Computation

Abstract

The laws of physics are encoded as differential equations that describe how systems evolve in time and space. These equations are traditionally derived by analyzing fundamental physical interactions. However, in many settings the underlying physics is too complex or high-dimensional to yield tractable differential equation models, resulting in partially or completely unknown governing equations. Advances in machine learning have yielded powerful predictive and generative models for complex and high-dimensional systems, but applications of these techniques to scientific problems often lack interpretability and are prone to violating physical constraints, especially when only limited data are available. Hybrid methods that integrate physical structure with data-driven models have emerged as a promising approach for accurate, interpretable modeling of complex physical systems. This thesis develops frameworks for learning interpretable dynamical models by combining first-principles structure with modern data-driven and machine learning methods. The resulting models aim to be accurate, computationally efficient, interpretable, and useful for prediction, control, and scientific understanding.

The first part of this thesis develops tools for modeling biophysical dynamics, with a focus on animal motion, behavior, and neural activity. These systems are challenging to model due to their high dimensionality, nonlinearities, and limited availability of comprehensive datasets. To address these challenges, this thesis introduces spectral mode representations as a low-dimensional and interpretable representation space for biological dynamics. While the form of the models for these biological systems is unknown, general physical symmetries and biological constraints are known. Accordingly, methods for incorporating physical constraints into dynamical models in mode space are developed and applied to animals whose dynamics can be effectively described by centerline motion. For linear dynamics, the constrained models are Hermitian, and this structure is utilized to compare dynamics across species and behavioral states. The framework is then extended to nonlinear, stochastic dynamics and neural control to build generative models of neural activity, motion, and behavior. Learning is facilitated by decomposing the vector field via a Helmholtz decomposition into gradient and divergence-free components, which are optimized separately using score matching, diffusion models, and generalized Hamiltonian dynamics. Coupling the dynamical system to neural activity enables prediction of motion from neural signals and the design of neural control strategies for steering motion. Applied to *C. elegans*, this model reproduces posture statistics across behavioral states, captures stereotyped dynamics, and links neural activity to interpretable modes of motion.

The second part of this thesis applies similar hybrid techniques to complex materials described by nonlinear partial differential equations (PDEs). These systems are challenging to model because key parameters and functions are often unknown and the equations are computationally intensive to simulate and optimize. To address these challenges, an open-source, GPU-accelerated framework is developed for optimization and control of phase-separating and pattern-forming PDEs, enabling efficient gradient-based optimization of arbitrary PDE parameters in complex domains. This tool is then leveraged to solve representative, experimentally relevant learning and control problems in energy and quantum materials, illustrating how physically structured optimization can be used to learn models for and control pattern formation. Building on this framework, this thesis develops methods for learning continuum models from molecular dynamics simulations of phase separation. Starting from particle-based simulations of a phase-separating Lennard–Jones fluid, coarse-grained continuum fields are constructed via diffusion-like smoothing operators and effective dynamical equations are learned in this continuum space. The coarse-graining procedure is linked to diffusion models from machine learning by demonstrating that the learned dynamics correspond to score functions of the coarse-grained distributions.

Finally, this thesis demonstrates the broader applicability of these ideas through several collaborative projects, including models for controlling indoor airborne disease transmission and for learning mosquito behavioral responses to multi-sensory stimuli.

Across these domains, the thesis advances a common perspective: integrating physical structure with data-driven models yields dynamical systems that are both interpretable and useful for prediction and control problems in many areas of science and engineering.

Thesis supervisor: Jörn Dunkel
Title: MathWorks Professor of Mathematics

Thesis supervisor: Martin Z. Bazant
Title: Chevron Professor of Chemical Engineering, Professor of Mathematics