Personal Statement

Optimization-Based Generative Systems through Objective Landscapes

Research Overview

I completed my PhD in AI at the University of Zurich and ETH AI Center, defending my dissertation, *Optimization-Based Composable Generation with Learnable and Scientific Priors*, in July 2026. My background in physics lets me use intuitions from the physical sciences to find new angles on AI problems, while mathematical training helps me turn these ideas into explicit, testable formulations. I develop generative systems whose outputs are produced by optimization, sampling, or search over objective landscapes combining scientific and learned priors with physical laws, simulations, measurements, and task-specific constraints.

My dissertation develops this idea along two complementary directions. First, I built differentiable scientific priors for inverse problems, combining physical equations, biomechanics, and patient data through transparent optimization. Second, I developed efficient-to-sample learnable priors, where the prior is an energy landscape that can be sampled, optimized, and composed with new constraints. Across four first-author papers, I led formulation, implementation, evaluation, and open-source dissemination.

Differentiable Scientific Priors

In GliODIL (*Nature Communications*, 2025), I developed a physics-informed method for individualized glioma radiotherapy planning and recurrence evaluation. It embeds a Fisher–Kolmogorov tumor-growth model in a differentiable objective combining patient imaging, PDE constraints, and expected growth patterns. A related physics-regularized multimodal image-assimilation framework (NeurIPS, 2024) jointly infers tumor concentration, tissue composition, deformation, and patient-specific parameters under growth and biomechanical constraints. These projects shaped my view that scientific AI should make its priors, constraints, and uncertainties visible.

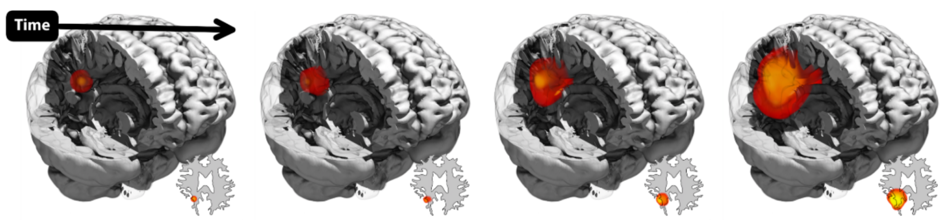

Physics-informed optimization assimilates patient data into a tumor-growth model for individualized prediction.

Efficient-to-Sample Learnable Priors

Many domains lack complete differentiable simulators, so the prior itself must be learned while remaining usable at inference. In Energy Matching (NeurIPS, 2025), I developed a framework that gives flow-based approaches the flexibility of energy-based models. A single time-independent scalar field connects transport-aligned generation with equilibrium modeling: far from the data manifold it guides efficient paths from noise to data, while near the manifold an entropic term supports Boltzmann exploration. The method advanced the state of the art in energy-based modeling across image generation and controlled protein optimization while retaining a prior that composes with observations, constraints, and property objectives. I designed the objective, implemented the training and sampling pipelines, led the evaluation, and released the code.

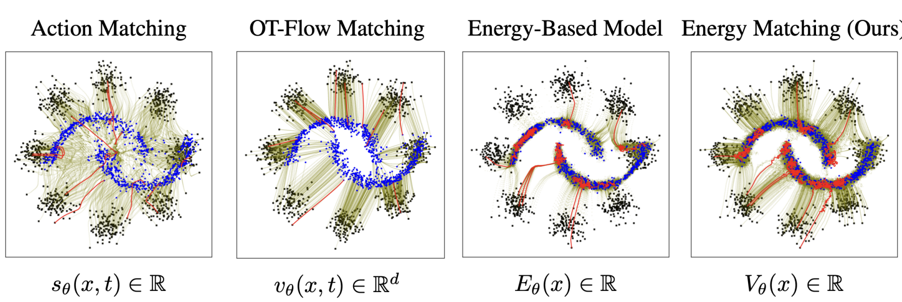

Transport from black noise samples to blue two-moons data. Energy Matching retains a time-independent scalar landscape while moving efficiently toward and along the data manifold.

Across both directions, my long-term aim is to learn and compose objective landscapes that integrate data-driven priors, simulators, tools, and external evidence within mathematically explicit inference loops. These landscapes should make generation and reasoning controllable, interpretable, and adaptable to new evidence without retraining.

Discrete Objectives and Molecular Navigation

Graph Energy Matching (GEM; 2026 preprint) extends this perspective to discrete scientific objects and molecular design. I developed its graph energy, proposal mechanism, corrected mixing phase, and main evaluations. Gradient-ranked edits preserve valid categorical graphs. Its explicit energy turns generation into molecular optimization: a learned molecular prior can be composed at inference with property objectives and structural constraints, supporting targeted navigation among candidate molecules for drug discovery. GEM is competitive with strong discrete generators. Together, Energy Matching and GEM provide the technical basis for my next question: how can explicit objective landscapes support general-purpose reasoning and planning over proofs, programs, plans, and interaction trajectories?

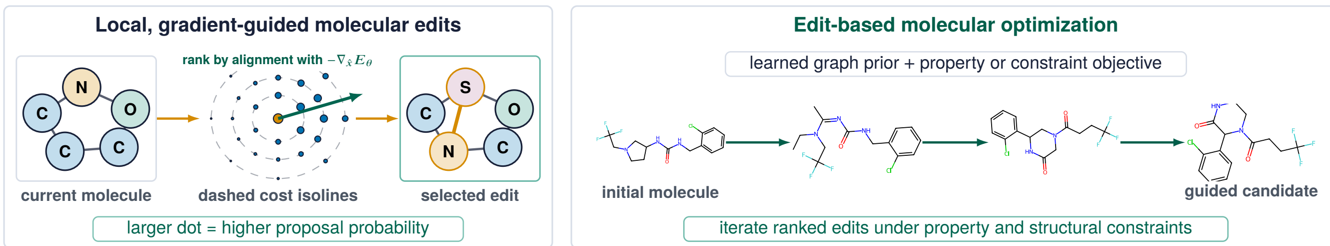

Graph Energy Matching turns learned objectives into efficient discrete search. Gradient-ranked categorical edits preserve valid graph states and guide molecular optimization under new property and structural constraints.