A187 · Augmenting protein stability predictions from generative models with non-equilibrium thermodynamics and physics-based potentials

Presenter: Kevin Borisiak — Yale University, Department of Physics

Authors: Kevin Borisiak, Jakub Otwinowski, Armita Nourmohammad

Session: Day 2 · Fri Oct 2, 2026 · 2:15–4:15 PM

Predicting protein stability is a central task in protein design. Recently, deep learning based methods such as folding-model confidence scores and inverse-folding likelihoods have largely displaced physics-based energy functions for this task, trading physical motivation and interpretability for speed. Can we recover that interpretability without sacrificing performance? Generative models trained on molecular simulations are a natural bridge: they provide efficient, steerable samplers of protein ensembles from which physical observables can be computed. But the learned distribution can be biased by its training data, and the generative trajectories that produce it are statistical constructs rather than physical paths. Here we present a framework grounded in non-equilibrium statistical physics that computes protein folding stability by importance sampling from BioEmu, a machine-learning emulator of molecular dynamics (MD). Treating BioEmu as a deterministic flow model, we integrate the likelihood ODE along generation trajectories to obtain the non-equilibrium work performed during sampling. Further evaluating the generated samples with physics-based potentials yields accurate estimators of the free energy, enthalpy, and entropy of folding. The likelihoods further decompose into per-residue contributions, quantifying each residue's share of the protein's total conformational entropy. We evaluate our pipeline against Protherm thermodynamic data and structural data from the PDB. More broadly, non-equilibrium reweighting offers a general route to physically grounded, interpretable observables from generative models of biomolecules.

Augmenting protein stability predictions from generative models with non-equilibrium thermodynamics and physics-based potentials

Kevin Borisiak[1,4*] Jakub Otwinowski[3,4] Armita Nourmohammad[1,2,3,4]

[1]Department of Physics, Yale University
[2]Department of Bioemedical Engineering, Yale University
[3]Department of Immunobiology, Yale School of Medicine
[4]Center for Systems Engineering and Immunology, Yale School of Medicine

Abstract

Predicting protein stability is a central task in protein design. Recently, deep learning based methods such as folding-model confidence scores and inverse-folding likelihoods have largely displaced physics-based energy functions for this task, trading physical motivation and interpretability for speed. Can we recover that interpretability without sacrificing performance? Generative models trained on molecular simulations are a natural bridge: they provide efficient, steerable samplers of protein ensembles from which physical observables can be computed. But the learned distribution can be biased by its training data, and the generative trajectories that produce it are statistical constructs rather than physical paths. Here we present a framework grounded in non-equilibrium statistical physics that computes protein folding stability by importance sampling from BioEmu, a machine-learning emulator of molecular dynamics (MD). Treating BioEmu as a deterministic flow model, we integrate the likelihood ODE along generation trajectories to obtain the non-equilibrium work performed during sampling. Further evaluating the generated samples with physics-based potentials yields accurate estimators of the free energy, enthalpy, and entropy of folding. The likelihoods further decompose into per-residue contributions, quantifying each residue's share of the protein's total conformational entropy. We evaluate our pipeline against Protherm thermodynamic data and structural data from the PDB. More broadly, non-equilibrium reweighting offers a general route to physically grounded, interpretable observables from generative models of biomolecules.

*Presenting author