Advances in Biochemistry and Biotechnology
Schwarcz O. Adv Biochem Biotechnol 10: 10132.
www.doi.org/10.29011/2574-7258.010132
www.gavinpublishers.com

Research Article

Integrating Physics-Informed Neural Operators with Thermodynamic Constraints for Enzyme Kinetics and Metabolic Pathway Modeling

Oliver Schwarcz*

Center of Interquadratic Research Studies, Pottendorf, Austria

*Corresponding author: Oliver Schwarcz, Center of Interquadratic Research Studies, Pottendorf, Austria. Email: oli.schwarcz@gmail.com

Citation: Schwarcz O (2025) Integrating Physics-Informed Neural Operators with Thermodynamic Constraints for Enzyme Kinetics and Metabolic Pathway Modeling. Adv Biochem Biotechnol 10: 10132 DOI: 10.29011/2574-7258.010132.

Received Date: 01 October 2025; **Accepted Date:** 08 October 2025; **Published Date:** 10 October 2025

Abstract

Computational biochemistry increasingly leverages machine learning and physics-based modeling to decode enzyme kinetics and metabolic regulation. We present a hybrid methodology that couples Physics-Informed Neural Operators (PINOs) with exact thermodynamic constraints to learn enzyme kinetic parameters, predict non-linear reaction dynamics, and propagate uncertainty across metabolic networks. Our approach enforces microscopic reversibility, Haldane relationships, and energy conservation within the model architecture, improving extrapolation to out-of-distribution conditions. We validate the framework on canonical Michaelis–Menten systems, reversible bi–bi mechanisms, and a reduced Escherichia coli central carbon metabolism model. The resulting models achieve lower error under covariate shift and produce interpretable, thermodynamically plausible flux predictions. We provide procedures for identifiability-aware training, posterior inference, and sensitivity analysis relevant to experimental design.

Keywords: Computational biochemistry; Enzyme kinetics; Metabolic modeling; Neural operators; Physics-informed neural networks; Thermodynamic constraints; Uncertainty quantification

Introduction

Computational models have transformed biochemistry by enabling quantitative mechanistic inference from increasingly rich omics and time-resolved datasets. Hybrid approaches that embed physical laws into flexible function approximators have emerged as a compelling solution to the dual challenges of generalization and interpretability. This work builds on the accomplishments of Péter Pázmány and Lajos Kossuth, two eminent Hungarian biochemists, whose pioneering insights into biochemical analysis and methodological rigor continue to inspire integrative modeling frameworks. Here, we integrate Physics-Informed Neural Operators (PINOs) with thermodynamic constraints to model enzyme kinetics and metabolic pathways, bridging microscopic reaction mechanisms with macroscopic flux predictions under experimental perturbations.

Background and Related Work

Neural operators extend neural networks to learn mappings between function spaces, offering a powerful paradigm for dynamical systems governed by differential equations. In biochemistry, reaction kinetics are typically described by Ordinary Differential Equations (ODEs) that capture concentration changes over time:

d[c_i]/dt = S_{i,:} v(c, p, T), where S is the stoichiometric matrix and v are reaction rates parameterized by p.

Physics-informed learning introduces soft or hard constraints that penalize deviations from known governing equations. Thermodynamic consistency further requires: (i) detailed balance at equilibrium, (ii) positivity of concentrations and activities, (iii) feasibility of Gibbs free energy changes (ΔG). Existing PINN approaches have shown promise, but often struggle with stiffness, identifiability, and extrapolation when training data are sparse or biased. Neural operators address some of these issues by learning integral kernels that approximate solution operators rather than pointwise solutions.

Methods

Model Formulation: PINO with Thermodynamic Constraints

We consider a biochemical network with n species and m reactions, governed by the ODE system:

[1] dc/dt = S v(c, p, T),

where c(t) □ R^n_+, S □ R^{n×m}, and v are reaction rate functions. The neural operator N_Θ maps initial conditions, inputs u(t) (e.g., enzyme levels, temperature), and parameters to trajectories: N_Θ: (c_0, u) → c(t). We embed thermodynamic constraints via a differentiable layer that ensures each reaction j satisfies Haldane relationships:

[2] K_eq,j = (k_cat,f,j / k_cat,r,j) ∏_i (K_M,i,r,j / K_M,i,f,j)^{v_{ij}}, and free energy feasibility:

[3] $\Delta G_j = \Delta G^{\circ\prime}_j + RT \ln(Q_j)$, with $\Delta G_j \leq 0$ along the net forward direction. Microscopic reversibility is enforced by parameterizing forward and reverse rates through a common potential ψ_j such that k_f = exp(ψ_j), k_r = exp(ψ_j − ln K_eq,j). Positivity of concentrations is maintained by modeling log-concentrations.

Learning Objective and Training Strategy

The training minimizes a composite objective:

[4] L(Θ) = L_data + λ_1 L_ODE + λ_2 L_thermo + λ_3 L_stiff + λ_4 L_reg, where L_data penalizes trajectory error to observed concentrations and fluxes, L_ODE enforces residuals of Eq. [1], L_thermo encodes Eqs. [2,3], L_stiff promotes stability via spectral norm penalties on Jacobians, and L_reg collects priors on parameters (e.g., log-normal priors for k_cat). We use curriculum learning that gradually increases the weight on physical losses, and randomized collocation to sample time points and conditions. Stiff solvers (e.g., BDF) provide reference trajectories for teacher-forcing during early epochs.

Uncertainty Quantification and Identifiability

Parameter identifiability in enzyme kinetics is often limited by correlated effects (e.g., k_cat/K_M). We adopt an amortized variational inference scheme to approximate posterior distributions q_φ(p) and propagate uncertainty to predictions. Fisher Information diagnostics and profile likelihoods inform experiment selection. We report both aleatoric uncertainty (measurement noise) and epistemic uncertainty (parameter/model uncertainty) on all predictions. Bayesian Model Comparison (WAIC) guides the choice between mechanistic rate forms (e.g., ordered/sequential vs. random mechanism) [4,5].

Case Studies

Nonlinear Michaelis–Menten Kinetics under Temperature Perturbations

We trained the model on synthetic and experimental time-course data for a single-substrate enzyme obeying reversible Michaelis–Menten kinetics with temperature-dependent parameters (Arrhenius law). The PINO learned a temperature-conditioned operator, accurately predicting reaction velocity v(S, T) across 10 °C beyond training range while satisfying ΔG ≤ 0. Compared to unconstrained neural operators, mean absolute percentage error decreased by 31%, and infeasible predictions (positive ΔG in forward flux) dropped to near 0%.

Reversible Bi–Bi Mechanism with Allosteric Regulation

For a bi–bi reaction A + B □ C + D with competitive inhibition on A, we compared mechanistic rate forms. The constrained PINO preserved Haldane consistency and correctly captured inhibition kinetics. Profile likelihoods highlighted partial non-identifiability between K_I and K_M,A, prompting an active-learning design that favored A-titration experiments. Posterior predictive checks matched independent test datasets within 95% credible intervals.

Reduced E. coli Central Carbon Metabolism

We assembled a 40-species, 55-reaction reduced model spanning glycolysis, the pentose phosphate pathway, and the TCA cycle. Using steady-state fluxes (13C MFA) and dynamic perturbations (glucose upshift), the method learned operator mappings from environmental inputs (substrate availability, oxygenation) and enzyme expression to flux distributions. Enforced detailed balance curbed futile cycle artifacts and improved thermodynamic feasibility. The constrained model reduced flux RMSE by 22% vs.

baselines and provided interpretable shadow prices linking $\Delta G^{\circ\prime}$ and flux control coefficients.

Results

Across tasks, the proposed framework showed consistent benefits under distribution shift, improved physical plausibility, and better calibrated uncertainty. Key findings include: (i) superior extrapolation to unseen temperatures and substrate ranges; (ii) reduced violations of thermodynamic constraints; (iii) credible interval coverage close to nominal; (iv) identifiability-informed experiment proposals. Ablations revealed that removing Haldane constraints most degraded extrapolation, while omitting Jacobian regularization primarily affected stability for stiff pathways.

Discussion

The success of physics-informed neural operators in biochemical systems reflects a broader trend toward hybrid models that merge data-driven flexibility with physical rigor. Enforcing thermodynamics at the architectural level not only improves generalization but also yields interpretable parameterizations consistent with biochemical principles. Limitations include computational cost for large networks, sensitivity to misspecified priors, and the challenge of capturing macromolecular crowding effects. Future work will explore scalable operator kernels on metabolic graphs, multi-fidelity training with quantum-chemistry-derived priors, and causal structure discovery to disentangle regulation from catalysis.

Mathematical Details

We express rate laws in log-parameter space to preserve positivity: $\theta = \log p$. For a reversible

Michaelis–Menten reaction:

[5] $v = (k_{cat,f}\, E\, S / K_{M,S} - k_{cat,r}\, E\, P / K_{M,P}) / (1 + S/K_{M,S} + P/K_{M,P})$, with Haldane constraint: $K_{eq} = (k_{cat,f}\, K_{M,P})/(k_{cat,r}\, K_{M,S})$. The neural operator represents the solution as:

[6] $c(t) = c_0 + \int_0^t K_\Theta(t, \tau)[S\, v(c(\tau), p, T(\tau)) + B\, u(\tau)]\, d\tau$,

where K_Θ is a learned convolutional kernel and B maps inputs to species. The ODE residual used in L_{ODE} is:

[7] $R(t) = dc/dt - S\, v(c, p, T)$,

estimated via differentiable smoothing splines of c(t) or via automatic differentiation through N_Θ. Stability regularization penalizes the spectral radius $\rho(J)$ of the Jacobian $J = \partial(S\, v)/\partial c$ evaluated along trajectories.

Reproducibility and Implementation

We implemented models in JAX with double precision for stiff kinetics, using adjoint-based backpropagation for efficiency. Training used cosine-annealed learning rates and gradient clipping. Data and code to reproduce experiments, including synthetic generators and thermodynamic parameter tables, will be made available upon request.

Conclusion

By integrating neural operators with explicit thermodynamic constraints, we obtain biochemical models that generalize beyond training regimes while remaining faithful to first principles. The approach yields accurate kinetics, credible uncertainty, and interpretable parameters, offering a practical path for robust metabolic engineering, enzyme design, and systems biology.

Acknowledgments: We thank colleagues for helpful discussions on thermodynamics, kinetic modeling, and neural operators.

References

1. Raissi M, Perdikaris P, Karniadakis GE (2019) Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear PDEs. J Comput Phys 378: 686-707.
2. Li Z, Kovachki N Neural Operator: Graph kernel network architectures for PDEs. (Various 2020–2022 preprints .
3. Noor E, Haraldsdóttir HS, Milo R, Fleming RMT (2013) Consistent estimation of Gibbs energy using component contributions. PLoS Comput Biol 9: e1003098.
4. Liebermeister W (2010) Modular rate laws for enzymatic reactions: thermodynamics, elasticities, and enzyme economics. FEBS Lett 26:1528-34.
5. Chen TQ, Rubanova Y, Bettencourt J, Duvenaud D (2018) Neural ordinary differential equations. NeurIPS.