Graph Neural ODEs with Stability and Conservation Guarantees for Tumor Microenvironment Dynamics (Student Abstract)

Luong Doan[1,3], Tien Nguyen[2], Nhung Duong[2], Lap Nguyen[2], Tuan Do[3]
[1]Phenikaa University, Nguyen Trac, Ha Noi, Viet Nam
[2]Hanoi University of Pharmacy, Le Thanh Tong, Ha Noi, Viet Nam
[3]N2TP Technology Solutions JSC, Khuat Duy Tien, Ha Noi, Viet Nam
luongdoan000@gmail.com

Abstract

We present Graph Neural ODEs (GNODEs) for modeling tumor microenvironment dynamics with mathematically guaranteed stability and conservation properties. Unlike bulk ODEs that miss spatial heterogeneity or discrete GNNs that inadequately capture continuous biological processes, GNODEs provide continuous-time evolution with explicit adjacency-aware dynamics while maintaining provable trajectory bounds. Our framework ensures: (1) existence and uniqueness of solutions under dynamic graph topology, (2) Lyapunov stability preventing unphysical states like negative cell counts, and (3) exact conservation of biological invariants through architectural constraints. Benchmarking on synthetic tumor data demonstrates that GNODE accurately captures the dynamics of the resistant cell fraction (0.282 predicted vs 0.242 true), whereas graph-free alternatives fail completely (0.000), underscoring the importance of stability-constrained local interactions for modeling emergent resistance.

Introduction and Biological Motivation

Tumor progression and therapy resistance emerge from complex local interactions within the tumor microenvironment (TME) (Almazrouei et al. 2025). Resistant clones exploit spatial niches through paracrine signaling, cell-cell contact, and local gradients of drugs and nutrients, phenomena that occur continuously across multiple timescales. Phosphorylation cascades happen in minutes, gene expression shifts over hours, and phenotypic transitions span days. Traditional compartmental ODEs treat tumors as homogeneous, missing these critical spatial dynamics (Arredondo and Rivera 2025). While agent-based models capture heterogeneity, they lack mathematical guarantees and ML integration (Norton et al. 2019). Discrete GNNs incorporate adjacency but miss the continuous nature of biological processes like drug diffusion and metabolic exchange (Li, Hua, and Chen 2025).

The critical gap is the absence of mathematical guarantees ensuring biological plausibility. Without stability proofs, models produce negative cell counts, unbounded proliferation, or violate conservation laws, making them unsuitable for clinical prediction. GNODEs address this by defining continuous-time neural ODEs (Chen et al. 2018) where each cell i has state $\mathbf{x}_i(t) \in \mathbb{R}^d$ evolving through graph-structured message passing with provable bounds. This is biologically essential: tumors are resource-limited systems where growth must respect carrying capacity and mass conservation.

Stability-Guaranteed GNODE Framework

Core Formulation

For N cells with time-varying adjacency $\mathcal{N}(i,t)$ representing spatial proximity or functional interaction:

$$\frac{d\mathbf{x}_i}{dt} = f_\theta\big(\mathbf{x}_i, \{\mathbf{x}_j\}_{j\in\mathcal{N}(i,t)}, c_i, t\big) \tag{1}$$

where f_θ is parameterized by a GNN with messages $m_{i\leftarrow j} = \phi_\theta(\mathbf{x}_i, \mathbf{x}_j, e_{ij})$ encoding local signals (e.g., growth factors, drug concentration). Node updates aggregate neighbor influences: $\frac{d\mathbf{x}_i}{dt} = \psi_\theta\big(\mathbf{x}_i, \sum_{j\in\mathcal{N}(i)} m_{i\leftarrow j}\big)$.

Theoretical Guarantees with Biological Interpretation

Theorem 1 (Existence/Uniqueness): Under bounded degree $|\mathcal{N}(i,t)| \le D_{\max}$ (biologically: limited local contacts) and spectral normalization, the system admits unique continuous solutions even with changing topology. Between adjacency switches at times $\{t_k\}$, Lipschitz constant $L \le L^* D_{\max}$ ensures local uniqueness via Picard-Lindelöf (Coddington and Levinson 1955). *Biological meaning*: Cell trajectories are deterministic given initial conditions and local interactions.

Theorem 2 (Lyapunov Stability): We construct energy function $V(\mathbf{X}) = \sum_i V_i(\mathbf{x}_i) + \sum_{(i,j)} V_{ij}(\mathbf{x}_i, \mathbf{x}_j)$ with $V \ge 0$ and $\dot{V} \le 0$ along trajectories. Enforced through weight normalization $\|W\| \le 1$ and stability penalty $\mathcal{L}_{\text{stab}} = \lambda \max(0, \rho(J_\theta) - 1)$ where $\rho(J_\theta)$ is the Jacobian's spectral radius (Khalil 2002). *Biological meaning*: Tumor growth remains bounded by resource constraints; no explosive proliferation or negative populations.

Theorem 3 (Conservation Laws): Zero-sum message passing architecturally guarantees $\sum_i f_\theta^{(k)} = 0$ for conserved quantities. *Biological meaning*: Total cell mass, energy, or drug amount cannot be created/destroyed, only redistributed, essential for realistic pharmacokinetics.

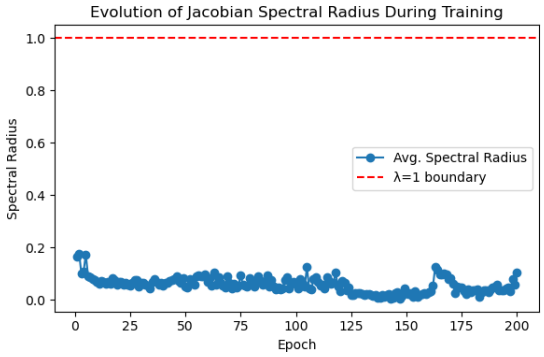

Figure 1: Jacobian spectral radius remains below stability boundary ($\rho < 1$) throughout training, guaranteeing bounded dynamics and preventing unphysical states.

Scalability for Tissue-Level Modeling

For large tumors (10^5 cells): degree capping preserves Lipschitz bounds (cells interact only locally), normalized aggregation prevents message explosion, and domain decomposition enables GPU parallelization while respecting stability at boundaries.

Validation: Multi-Scale ABM-PDE Simulator

Given scarce longitudinal single-cell TME data, we developed a ground-truth simulator coupling agent-based cells with a PDE-governed microenvironment. The system evolves 200-300 cells with 75-dimensional states (gene expression, metabolism, signaling) undergoing stochastic transitions: Sensitive $\rightarrow$ Tolerant $\rightarrow$ Resistant, influenced by local drug/nutrient concentrations. Drug and nutrient fields diffuse on a 25^3 grid via finite difference methods with proven stability: CFL condition (Abe et al. 2014) $\lambda = D\Delta t/h^2 \leq 1/6$ ensures non-negativity, scheme is L^∞-stable, and mass balance is exactly preserved. Output includes 70% sampling with noise ($\sigma = 0.05$), mimicking experimental limitations.

Results: Stability Enables Biological Accuracy

Table 1 reveals that GNODE uniquely captures resistant subpopulation dynamics (0.282 vs true 0.242), the critical emergent phenomenon driving treatment failure. The Graph-Free ODE catastrophically fails (0.000 resistant cells) despite lower MSE, demonstrating that aggregate error metrics miss biological relevance. This failure occurs because resistance emerges from local niches: tolerant cells shield resistant neighbors from drug exposure, a spatial phenomenon invisible to graph-free methods. The discrete GNN underperforms (0.110) due to its inability to model continuous diffusion and signaling timescales. Crucially, only GNODE and the baseline GNN operate without needing the final or "end-stage" topology in advance, which allows them to capture local interactions as they evolve.

Figure 1 shows the spectral radius consistently below unity during training, confirming stable learned dynamics. Conservation violations approach machine precision ($<$

Model	MSE	Resistant Fraction	KL Div
GNODE	107.9	**0.282** (True: 0.242)	0.47
Discrete GNN	312.0	0.110	1.92
Neural ODE (Flat)	72.9	0.188	0.35
Graph-Free ODE	71.4	**0.000**	6.97
LSTM	78.6	0.162	0.20

Table 1: Benchmark Performance on Resistant Fraction Prediction

10^{-14}), ensuring mass balance across arbitrary time horizons. These aren't just mathematical properties; they guarantee the model respects fundamental biological constraints like finite resources and physical crowding.

Biological Impact

The provable stability of GNODEs transforms them from black-box predictors to mechanistically interpretable models suitable for precision oncology. Bounded trajectories respect tumor carrying capacity determined by oxygen and nutrient availability. Conservation laws maintain drug mass balance, critical for accurate pharmacokinetic/pharmacodynamic modeling. The stable Jacobian reveals which perturbations (drug pulses, immune infiltration, stromal remodeling) push the system toward resistance-emergence thresholds.

GNODE analysis indicates optimal temporal sampling: rapid early points capture signaling dynamics, while later sparse sampling tracks phenotypes. This enables gradient-based therapy optimization, designing adaptive protocols to maintain tumor control and minimize resistance.

Current limitations center on data availability, though spatial transcriptomics and multiplexed imaging are rapidly advancing. The identifiability challenge, multiple parameter sets satisfying constraints, requires biological priors from mechanistic knowledge. Future work will integrate immune dynamics, angiogenesis, and tissue-level PDEs for comprehensive tumor-immune-drug-stroma modeling.

Conclusion

GNODEs provide the first continuous-time framework for TME dynamics with mathematical guarantees matching biological reality. Through bounded degrees (local interactions), spectral normalization (growth limits), and architectural conservation (mass balance), we ensure trajectories remain biologically plausible while capturing the adjacency-driven resistance that determines treatment outcomes. The framework's proven stability makes it suitable for clinical applications where uncontrolled predictions could harm patients. Our results demonstrate that these guarantees enable accurate emergence prediction; only GNODE correctly captures resistant dynamics through local spatial mechanisms. As spatial single-cell technologies mature, GNODEs offer a principled foundation for learning interpretable tumor evolution from complex tissue data, bridging mathematical rigor with biological mechanisms for improved cancer therapy.

References

Abe, K.; Higashimori, N.; Kubo, M.; Fujiwara, H.; and Iso, Y. 2014. A Remark on the Courant-Friedrichs-Lewy Condition in Finite Difference Approach to PDE's. *Advances in Applied Mathematics and Mechanics*, 6(5): 693–698.

Almazrouei, K.; Mishra, V.; Pandya, H.; Sambhav, K.; and Bhavsar, S. 2025. Tumor Microenvironment and Its Role in Cancer Progression: An Integrative Review. *Cureus*, 17(9): e92707.

Arredondo, J. A.; and Rivera, A. 2025. Recent advances in ODEs modeling of tumor-immune responses: a focus on delay effects. *Mathematical Biosciences and Engineering*, 22(12): 3060–3087.

Chen, R. T. Q.; Rubanova, Y.; Bettencourt, J.; and Duvenaud, D. K. 2018. Neural Ordinary Differential Equations. In Bengio, S.; Wallach, H.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.

Coddington, E. A.; and Levinson, N. 1955. *Theory of Ordinary Differential Equations*. McGraw-Hill.

Khalil, H. K. 2002. *Nonlinear Systems*. Upper Saddle River, NJ: Prentice Hall, 3rd edition. ISBN 978-0-13-067389-3.

Li, S.; Hua, H.; and Chen, S. 2025. Graph neural networks for single-cell omics data: a review of approaches and applications. *Briefings in Bioinformatics*, 26(2).

Norton, K.-A.; Gong, C.; Jamalian, S.; and Popel, A. S. 2019. Multiscale agent-based and hybrid modeling of the tumor immune microenvironment. *Processes*, 7(1): 37.