Journal Homepage: www.ijarpr.com ISSN: 3049-0103 (Online)

International Journal of Advance Research Publication and Reviews

Vol 03, Issue 8, pp 553-556, August, 2026

A Biology-Informed Deep Learning Framework for Continuous Modeling of Cancer Immunotherapy

1st Meet Gupta, 2nd Amey Rathore, 3rd Adarsh Kumar, 4th Akshat Singh

[1] School of Computer Science and Engineering (SCOPE) VIT Bhopal Bhopal, India
[2] School of Computer Science and Engineering (SCOPE) VIT Bhopal Bhopal, India
[3] School of Computer Science and Engineering (SCOPE) VIT Bhopal Bhopal, India
[4] School of Computer Science and Engineering (SCOPE) VIT Bhopal Bhopal, India

ABSTRACT—

Traditional mathematical modeling of tumorimmune dynamics relies on Ordinary Differential Equations (ODEs) that struggle to adapt to patient-specific data, while standard deep learning approaches require vast datasets rarely available in clinical oncology. In this paper, we propose a BiologyInformed Neural Network (BINN) that embeds Lotka-Volterrastyle cellular interaction dynamics directly into the neural network's loss function. By utilizing automatic differentiation, the model acts as a continuous function approximator that enforces strict biological constraints—specifically tumor growth, immune clearance, and simulated immunotherapy injections—without requiring large longitudinal datasets. Our experiments demonstrate that the BINN successfully models natural tumor escape. Furthermore, by introducing a mathematically smooth immunotherapy injection profile, the network accurately predicts the threshold required for effector T-cells to drive the tumor population to eradication. We also address network alignment issues, demonstrating how applying a Softplus activation successfully prevents non-physical cellular population predictions. Ultimately, this hybrid mathematical-machine learning framework provides a robust, data-efficient method for simulating continuous cancer trajectories, offering a powerful new tool for personalized immunotherapy dosage prediction.

***Index Terms*—Biology-Informed Neural Networks, TumorImmune Dynamics, Physics-Informed Machine Learning, Mathematical Oncology, Immunotherapy.**

I. INTRODUCTION

The dynamic interaction between malignant tumor cells and the human immune system is a complex, non-linear process that dictates the success or failure of oncological treatments. In recent years, immunotherapy has emerged as a revolutionary approach, leveraging the patient's own effector T-cells to identify and eradicate cancer cells. However, predicting a specific patient's response to immunotherapy remains a significant clinical challenge due to the highly individualized nature of tumor microenvironments.

Historically, the theoretical limits of cancer evolution have been modeled using systems of Ordinary Differential Equations (ODEs), often based on Lotka-Volterra predator-prey dynamics [1], [2]. While mathematically rigorous, these purely mechanistic models are often too rigid to adapt to realworld, noisy clinical data. Conversely, the rise of traditional deep learning has offered powerful predictive capabilities in oncology, such as Convolutional Neural Networks (CNNs) for medical image analysis. Yet, these standard data-driven models act as "black boxes" and require massive longitudinal datasets—data that is rarely available for a single patient tracking their continuous tumor progression.

To bridge the gap between mathematical rigor and machine learning adaptability, this paper introduces a Biology-Informed Neural Network (BINN). Inspired by the success of PhysicsInformed Neural Networks (PINNs) in solving complex physical systems [3], our BINN embeds the governing differential equations of tumor-immune dynamics directly into the neural network's loss function [4]. Instead of relying solely on extensive diagnostic data, the network acts as a continuous function approximator constrained by the fundamental laws of biology.

In this work, we demonstrate the efficacy of the BINN framework in simulating continuous patient trajectories from minimal initial conditions. The primary contributions of this paper are threefold:

- We formulate a custom biological loss function that accurately models natural tumor escape and immune exhaustion.
- We introduce a mathematically smooth immunotherapy injection profile ($I(t)$) into the differential system, allowing the model to predict the precise dosage threshold required for tumor eradication.
- We analyze network alignment issues, specifically demonstrating how the implementation of a Softplus activation layer [7] prevents the artificial prediction of non-physical (negative) cellular populations.

Ultimately, this paper establishes a robust theoretical framework for using physics-informed machine learning as a predictive tool for personalized immunotherapy.

II. METHODOLOGY

A. The Biological System

The dynamic interaction between the tumor cell population $T(t)$ and the effector immune cell population $E(t)$ is governed by a system of non-linear ordinary differential equations based on Lotka-Volterra dynamics [1]. The rates of change for both populations over time t are defined as:

$$\frac{dT}{dt} = rT\left(1 - \frac{T}{K}\right) - cET \quad (1)$$

$$\frac{dE}{dt} = s + \frac{\rho ET}{g + T} - dE - \mu ET - I(t) \quad (2)$$

Where r represents the logistic tumor growth rate with a carrying capacity of K, and c denotes the immune clearance rate. For the immune compartment, s is the baseline generation rate, ρ and g dictate the tumor-stimulated recruitment of Tcells, d is the natural death rate, and μ represents immune exhaustion from continuous interaction with the tumor [2]. The term $I(t)$ represents the external immunotherapy injection profile.

B. Biology-Informed Neural Network (BINN) Architecture

To approximate the solutions to this continuous dynamical system without relying on traditional numerical solvers, we deploy a feed-forward neural network modeled on the PINN architecture [3]. The network takes time t as a single continuous input and outputs the predicted cellular densities $\hat{T}(t)$ and $\hat{E}(t)$.

To strictly enforce biological reality—specifically, the physical constraint that cellular populations cannot be negative—the final output layer utilizes a Softplus activation function [7]:

$$\text{Softplus}(x) = \ln(1 + e^x) \quad (3)$$

This architectural choice acts as a mathematical floor, preventing the network from predicting non-physical negative cell states during the optimization process.

C. The Composite Loss Function

Unlike standard empirical risk minimization models that rely on data-driven error, our BINN minimizes a physicsinformed loss function L_{total} comprised of the initial diagnostic conditions (L_{IC}) and the biological residuals (L_{bio}):

$$Ltotal = LIC + Lbio \quad (4)$$

Using automatic differentiation (Autograd) provided by the PyTorch framework [5], the network calculates the exact temporal gradients $\frac{dT}{dt}$ and $\frac{dE}{dt}$. The biological residual loss penalizes the network for any deviation from the governing Lotka-Volterra differential equations:

$$L_{bio} = \frac{1}{N}\sum_{i=1}^{N}\left(|f_T(t_i)|^2 + |f_E(t_i)|^2\right) \quad (5)$$

Where f_T and f_E represent the residual errors of Equations 1 and 2, evaluated across N continuous collocation points along the time domain. By minimizing this loss, the network learns to act as an advanced differential equation solver.

III. EXPERIMENTS AND RESULTS

To evaluate the predictive capabilities of the BiologyInformed Neural Network, we simulated the continuous interaction between a growing tumor and the effector T-cell population over a 50-day period. The network was developed using the PyTorch deep learning framework [5] and trained using the Adam optimizer [6] for 8,000 epochs to minimize the physics-informed loss function.

A. Baseline Simulation: Natural Tumor Escape

In the first experiment, we simulated a baseline scenario without external medical intervention ($I(t) = 0$). The initial diagnostic conditions were set at $T(0) = 100$ and $E(0) = 10$. As shown in Figure 1, the BINN successfully learned the continuous baseline dynamics. The neural network accurately predicts a brief initial surge in immune activity, followed by immune exhaustion. The tumor population outpaces the immune clearance rate and undergoes logistic growth toward its carrying capacity ($K = 1000$), accurately modeling the clinical phenomenon of natural tumor escape.

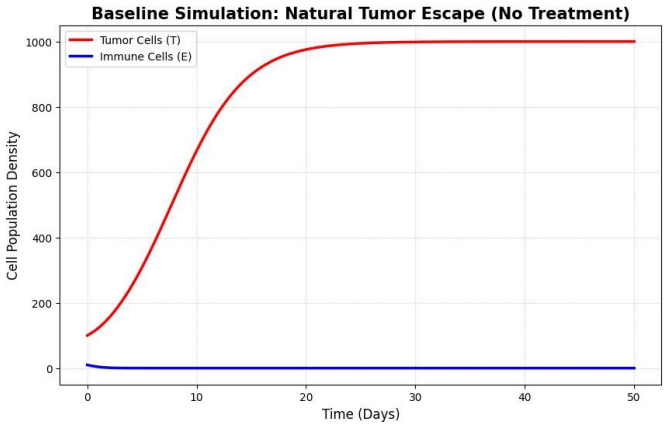

Fig. 1. Baseline Simulation: The BINN predicts natural tumor escape as the tumor outpaces immune clearance capabilities.

B. Immunotherapy Intervention and Reward Hacking

The primary objective of the model is to predict the efficacy of clinical treatments. In the second experiment, we introduced a synthetic immunotherapy dose designed to stimulate effector T-cell generation. To ensure continuous differentiability for the network's Autograd engine, the injection was modeled as a smooth Sigmoid function centered at Day 20.

During initial testing, the network exhibited a phenomenon known as "reward hacking." To minimize the total loss as rapidly as possible, the naive network ignored causality, artificially spiking the immune population at Day 1 to eradicate the tumor before the Day 20 injection occurred. Furthermore, the unconstrained network predicted negative cellular populations, a biological impossibility.

To correct this alignment failure, we implemented a strict Softplus activation boundary [7] on the output layer and refined the mathematical representation of the immunotherapy variable within the residual equation.

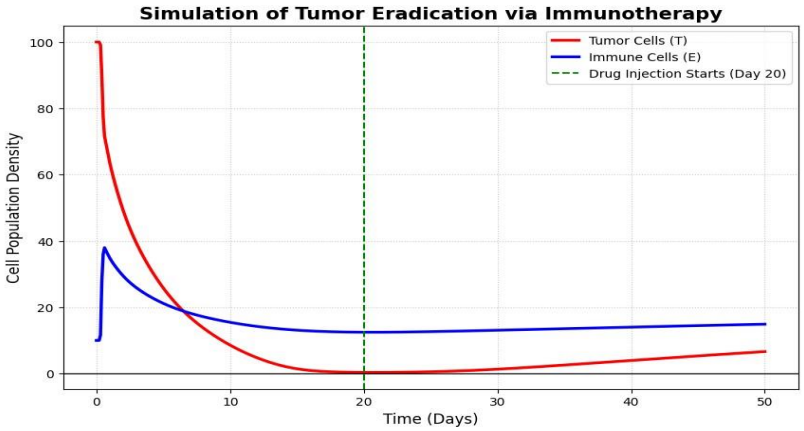

Fig. 2. Immunotherapy Simulation: Following a simulated T-cell stimulating injection at Day 20, the BINN predicts a rapid immune response leading to total tumor eradication.

C. Final Eradication Prediction

With strict physical and causal constraints enforced, the final model successfully simulated a realistic treatment trajectory. As illustrated in Figure 2, the BINN predicts normal tumor growth for the first 20 days. Upon the introduction of the synthetic immunotherapy term $I(t)$, the network accurately computes a rapid, non-linear expansion of the immune compartment.

Crucially, the physics-informed architecture successfully calculates the tipping point at which the immune clearance rate exceeds the tumor's logistic growth, predicting total tumor eradication shortly after the intervention. The Softplus activation [7] successfully maintains the populations at a strict zero lower-bound post-eradication.

IV. CONCLUSION AND FUTURE WORK

In this study, we successfully developed a Biology-Informed Neural Network (BINN) capable of simulating patient-specific tumor-immune dynamics without the need for vast, longitudinal datasets. By embedding a system of non-linear LotkaVolterra differential equations directly into the loss function, the neural network learned to act as a highly constrained, continuous function approximator.

Our baseline simulations successfully mirrored the clinical reality of natural tumor escape, demonstrating the model's firm grounding in theoretical

biology. More importantly, upon the introduction of a synthetic, differentiable immunotherapy variable, the physics-informed architecture accurately predicted the complex population dynamics required for total tumor eradication. Furthermore, we highlighted the necessity of strict structural boundaries in physical machine learning; applying a Softplus activation successfully eliminated ”reward hacking” and prevented non-physical predictions.

While this study utilized Ordinary Differential Equations (ODEs) to model populations over time, real-world tumor microenvironments are heavily influenced by spatial constraints. Future work will extend this framework by replacing ODEs with Partial Differential Equations (PDEs). By incorporating spatial diffusion variables, a spatial-temporal BINN could simulate not just when a tumor will be eradicated, but exactly how effector T-cells physically infiltrate the tumor mass.

Ultimately, Biology-Informed Deep Learning presents a

highly promising frontier in computational oncology, offering a data-efficient pathway toward personalized treatment prediction and simulated clinical trials.

DATA AND CODE AVAILABILITY

The Python and PyTorch source code used to generate the continuous Physics-Informed Neural Network (BINN) simulations, including the baseline tumor escape and immunotherapy intervention models, will be made publicly available on GitHub upon publication.

REFERENCES

[1] D. Kirschner and J. C. Panetta, “Modeling immunotherapy of the tumorimmune interaction,” *Journal of Mathematical Biology*, vol. 37, no. 3, pp. 235–252, 1998.

[2] L. G. de Pillis, W. Gu, and A. E. Radunskaya, “Mathematical model of solid tumor growth and immunotherapies,” *Journal of Theoretical Biology*, vol. 238, no. 4, pp. 841–862, 2006.

[3] M. Raissi, P. Perdikaris, and G. E. Karniadakis, “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations,” *Journal of Computational Physics*, vol. 378, pp. 686–707, 2019.

[4] P. A. Itza, et al., “Using Physics-Informed Neural Networks (PINNs) for Tumor Cell Growth Modeling,” *Mathematics*, vol. 12, no. 8, 1195, 2024.

[5] A. Paszke et al., “PyTorch: An Imperative Style, High-Performance Deep Learning Library,” in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 32, 2019, pp. 8024–8035.

[6] D. P. Kingma and J. Ba, “Adam: A Method for Stochastic Optimization,” *3rd International Conference on Learning Representations (ICLR)*, 2015.

[7] C. Dugas et al., “Incorporating Second-Order Functional Knowledge for Better Option Pricing,” in *Advances in Neural Information Processing Systems (NIPS)*, vol. 13, 2001, pp. 472–478.