BULLETIN OF STOMATOLOGY AND MAXILLOFACIAL SURGERY
Volume 21, Issue 11

DOI: 10.58240/1829006X-2025.21.11-243

ORIGINAL RESEARCH

DEEP-LEARNING–AUGMENTED IN-SILICO MODELING OF HOST IMMUNE DYNAMICS DURING PERIODONTAL INFECTION

Prabhu Manickam Natarajan*[1] Nandini M S[2], Pradeep Kumar Yadalam[3*] Thulasiram E[4], Shafiyah M[5]

[1]Department of Clinical Sciences, Center of Medical and Bio-allied Health Sciences and Research, College of Dentistry.prabhuperio@gmail.com
[2]Associate professor, Department of microbiology, Sree Balaji medical College and hospital, Bharath Institute of higher education and research, Chennai, Tamilnadu, India drnandini@bharathuniv.ac.inORCID : 0000-0002-4504-3536
[3]*Department of Periodontics, Saveetha Dental College and Hospitals, Saveetha Institute of Medical and Technical Sciences (SIMATS), Saveetha University, Chennai, Tamil Nadu, India (corresponding)Pradeepkumar.sdc@saveetha.com
[4]Reader, Department of Orthodontics, Sree Balaji Dental College and Hospital, Bharath Institute of Higher Education and Research,Chennai,Tamilnadu,India.dr.thulasiramcare32@gmail.com
[5]Post Graduate, Department of Periodontology, Sree Balaji Dental College and Hospital, Bharath Institute of Higher Education and Research (BIHER), Chennai, Tamilnadu, India fiyahmubarak1998@gamil.com
Corresponding Author: Prabhu Manickam Natarajan Department of Clinical Sciences, Center of Medical and Bio-allied Health Sciences and Research, College of Dentistry.prabhuperio@gmail.com
Received: *Oct 29. 2025;* ***Accepted:*** *Nov27, 2025;* ***Published:*** *Dec. 5, 2025*

ABSTRACT

Background: Periodontal infection involves complex interactions among bacteria, immune cells, cytokines, and tissue, which can be challenging to model and predict.
Objectives: To develop a minimal mechanistic model of periodontal infection incorporating key components (bacteria, neutrophils, macrophages, cytokine, tissue), generate a synthetic cohort, and train a Long Short-Term Memory (LSTM) neural network to forecast tissue integrity and evaluate antibiotic intervention schedules.
Results:A 100-subject synthetic cohort was generated using the mechanistic model. The LSTM network was trained on five days of multi-omic data to predict one-day-ahead tissue integrity, achieving a mean squared error of approximately 690 on a hold-out set. Virtual antibiotic treatment simulations demonstrated a preservation of about 15% more tissue by day 30 compared to untreated controls.
Conclusion: The combined mechanistic and neural modeling framework effectively captured disease progression and enabled virtual evaluation of therapy schedules, highlighting its potential for optimizing periodontal treatment strategies.

***Keywords*:**Periodontitis, immunity, simulation

INTRODUCTION

Periodontal disease (periodontitis) is a common chronic inflammatory disease of the tooth-supporting tissues, characterized by progressive destruction of the periodontal ligament and adjacent alveolar bone. It is a major cause of tooth loss worldwide and a significant public health concern, affecting a substantial portion of the adult population (estimated at 20–50% globally, with approximately one billion cases worldwide). The disease is initiated by bacterial biofilms (dental plaque) at the gingival margin[1–3]Still, a complex interplay between the pathogenic microbes and the host immune response drives its progression and tissue damage. In a healthy scenario, the immune response contains the infection and maintains tissue homeostasis; however, in periodontitis, the host response can become dysregulated – a double-edged sword that not only fails to eliminate the biofilm but also causes collateral damage to the surrounding tissues. Indeed, uncontrolled chronic inflammation in the

gingiva is the primary culprit responsible for the destruction of connective tissue and bone observed in periodontitis[4]. A striking feature of periodontal disease is the high variability in disease severity and progression among individuals. Most adults experience mild to moderate periodontitis, while only a subset develops severe disease with rapid tissue loss. Even within a single patient, progression is non-linear – periods of quiescence or slow loss can be punctuated by acute episodes of rapid breakdown (flare-ups). This unpredictable, episodic course is thought to result from a combination of factors, including genetic predisposition, systemic health, and environmental/lifestyle influences (such as smoking), all of which modulate the host immune response. These complexities make it challenging for clinicians to predict outcomes or determine which biological processes most strongly drive disease progression in a given patient.

Computational modeling offers a powerful approach to unravel this complexity by quantitatively testing how changes in specific parameters or mechanisms affect disease trajectories. Mechanistic mathematical models (e.g., systems of ordinary differential equations, ODEs) can represent hypothesized interactions between bacteria and immune components, helping to identify which host or bacterial factors govern the course of infection[4,5]. Such models have previously been used to gain insight into the dynamics of periodontitis, for example, identifying that the level of host immune response can control progression rates in simulated disease. However, purely mechanistic models require numerous parameters (e.g., growth rates, reaction coefficients) that are often hard to measure in vivo or vary between individuals, leading to uncertainty in model predictions. On the other hand, purely data-driven models (e.g., machine learning classifiers or regressors) can be trained on clinical data to predict disease outcomes; however, such models often act as “black boxes” and may overlook known biological constraints and the mechanistic understanding of the system. There is a growing recognition that combining mechanistic and data-driven approaches can yield better predictive performance while preserving interpretability[6,7]. In the context of inflammatory diseases like periodontitis, a hybrid model could leverage mechanistic simulations to generate physiologically plausible scenarios and then utilize modern AI techniques to emulate these dynamics or integrate real patient data efficiently(8). We propose a deep learning–augmented in-silico model for host-pathogen dynamics in periodontal disease, bridging mechanistic and data-driven approaches. We: 1. Create a minimal ODE model of inflammation to generate synthetic disease trajectories. 2. Train an LSTM network on this data as a fast surrogate to forecast disease progression from time series inputs. 3. Use the combined model to run virtual treatment trials, like antibiotic scenarios, predicting outcomes. This hybrid approach maintains biological interpretability and enables rapid what-if analyses. We hypothesize that a small neural network trained on simulated multi-omic data can capture the dynamics of periodontal tissue destruction. It can also optimize therapy timing by comparing virtual treatment and no-treatment scenarios. We detail methods for data generation, network training, and virtual antibiotic trials, then present performance results and future implications.

MATERIALS AND METHODS

Synthetic Cohort Generation

Mechanistic ODE Model: We constructed an ordinary differential equation model representing the key components of the host-pathogen interaction in periodontal infection. The model tracks five state variables over time: Bacteria (B), Neutrophils (N), Macrophages (M), Cytokine (C), and Tissue Integrity (T). These variables correspond to the subgingival bacterial load, the local abundance of neutrophils and macrophages (innate immune cells that infiltrate periodontal tissue), the level of pro-inflammatory cytokines in the gingival crevicular fluid (as a proxy for overall inflammatory mediator concentration), and the fraction of healthy periodontal tissue remaining (with T = 1.0 representing the baseline/healthy tissue and lower values indicating tissue breakdown). The ODE system incorporates the fundamental interactions among these components, including:

- Bacterial growth and clearance: Bacteria proliferate at an intrinsic growth rate rB (representing net expansion of the biofilm in the absence of immune pressure). They are cleared primarily by neutrophils (and to a lesser extent macrophages), modeled as a killing term proportional to the product $N \cdot B$ (and $M \cdot B$), with a killing rate constant kN for neutrophils (and kM for macrophages). This term reflects phagocytosis or other bactericidal mechanisms.
- Immune cell recruitment (chemotaxis): Neutrophils and macrophages are recruited to the infection site in response to bacterial presence and inflammatory signals. In the model, the influx of neutrophils is driven by the bacterial load (with a term such as $dN/dt = \alpha N \cdot B - dN \cdot N$, where αN is a recruitment rate and dN is the natural clearance/death rate of neutrophils). Macrophage dynamics are similarly driven by bacteria (with rate αM) and subject to a clearance rate dM. This captures the arrival of monocytes from the bloodstream and their differentiation into macrophages in tissue.

- Cytokine release and decay: The cytokine variable C represents an aggregate of pro-inflammatory cytokines (such as IL-1β, TNF-α, IL-6) that are secreted by activated immune cells (neutrophils, macrophages) in response to the infection. We include source terms for C proportional to N and M (e.g., $dC/dt = pN \cdot N + pM \cdot M - dC \cdot C$), where pN, pM are cytokine production rates and dC is the cytokine clearance rate (modeling natural decay or elimination of inflammatory signals). This yields transient cytokine spikes following the influx of immune cells.
- Tissue degradation: Tissue integrity T is reduced over time as a consequence of inflammation. In the model, we couple tissue loss to the level of cytokines (and possibly to active neutrophils as well, since neutrophil-secreted enzymes can directly damage tissue). For simplicity, the dominant driver was taken to be the inflammatory cytokine burden: $dT/dt = -kC \cdot C \cdot f(T)$, where kC is a tissue damage rate per unit cytokine. We let f(T) be either a constant or a function ensuring T cannot drop below zero (in our minimal model, T was unbounded below 0, representing an index that can go negative to indicate more than baseline tissue has been destroyed; in reality 0% would be total loss). This formulation suggests that high levels of pro-inflammatory cytokines stimulate tissue-destructive processes (e.g., matrix metalloproteinases and osteoclast activation), leading to collagen breakdown and bone resorption. Parameter Sampling: To generate a diverse synthetic cohort of disease trajectories, we sampled the model's parameters from physiologically plausible ranges, informed by literature and expert knowledge (see Supplementary Table S1 for the complete parameter list). Key parameters and their sampled ranges included:

- Bacterial growth rate rB ~ 0.1 – 0.3 day^−1 (reflecting slow to moderate replication rates of anaerobic oral bacteria in biofilms).
- Neutrophil bactericidal efficiency kN ~ 0.01 – 0.05 (per 10^4 cells)^−1 day^−1, which scales the killing term $N \cdot B$ (for example, 10,000 neutrophils at the site with kN=0.02 would neutralize ~0.2 of the bacterial population per day).
- Macrophage killing rate kM (generally lower than neutrophils, set to roughly half of kN in range).
- Recruitment rates αN, αM for neutrophils and macrophages (chosen such that neutrophils respond rapidly to the infection, while macrophage influx is slightly delayed/slower).
- Cytokine production rates pN, pM, and clearance rate dC (tuned so that cytokine peaks shortly after the peak of immune cell influx and then decay over a day or two if the infection is cleared).
- Tissue damage rate kC (set such that sustained high cytokine levels can reduce tissue by tens of percent over weeks).

Each "virtual subject" in the cohort had a unique parameter set drawn from uniform distributions over these ranges. Initial conditions were also set to mimic the onset of infection: we assumed at time zero a small-established bacterial presence (on the order of 10^4–10^5 bacterial units; for some subjects we varied initial B to represent different infection inoculum sizes), with baseline tissue integrity T(0) ≈ 1.0 (100%) and minimal initial inflammation (N(0) and M(0) at low basal levels, C(0) near 0). All subjects began the simulation on "day 0" of infection and were followed for 30 days. Numerical Integration and Noise: For each subject, the parameter set and the ODE system were integrated from day 0 to 30 with a 1-day step. A standard ODE solver (odeint from SciPy) handled the stiff system. Gaussian noise (mean 0, std dev 5%) was added to each state at each time point, creating a synthetic dataset resembling noisy longitudinal measurements. Results were stored as a time-series dataframe with columns

Table 1 shows the example outputs for Subject 0 on Days 0–4. Note how a strong immune response quickly eliminated Bacteria (B) in this particular case, while cytokine levels spiked and tissue integrity T dropped below 0, indicating extensive tissue loss beyond the original baseline by day 4 in this simulated aggressive scenario.

Bacteria	Neutrophils	Macrophages	Cytokine	Tissue	Subject	Day
1.85×10^4	128.8	62.3	13.28	0.949	0	0
1.27×10^4	119.6	58.5	120.13	0.672	0	1
1.95×10^1	102.05	61.0	182.96	-0.049	0	2
6.06×10^−5	93.47	64.65	224.79	-0.964	0	3
3.13×10^−7	80.05	73.21	250.84	-2.261	0	4

In total, we generated 100 such simulated subjects. This synthetic cohort exhibits a wide range of infection outcomes, from relatively minor tissue loss to severe destruction, providing a rich dataset for training and testing the forecasting model (table-1).

LSTM Forecasting Model

To create a fast surrogate for the ODE simulator, we trained a recurrent neural network to predict the future course of tissue integrity from recent history. We chose a Long Short-Term Memory (LSTM) network because LSTMs are well-suited for sequence data and can capture temporal dependencies. The forecasting task was defined as predicting next-day tissue integrity given the previous 5 days of all state variables. In other words, the model takes a window of 5 consecutive days of data [B(t-4), N(t-4), M(t-4), C(t-4), T(t-4)] up to [B(t), ..., T(t)] and outputs a forecast for T(t+1) (the tissue integrity on the following day). We used a short input window (5 days) to mimic the idea that clinicians might use about a week's worth of patient data to forecast near-term disease progression.

Network Architecture: The network was deliberately kept lightweight to test if a relatively small model could capture the dynamics. It consisted of a single LSTM layer with 32 hidden units, followed by a fully connected dense output layer that produces a single scalar (the predicted tissue value for the next day). Thus, the architecture can be summarized as: *5 time steps × 5 features → LSTM(32 hidden units) → Dense(1 output)*. No additional hidden layers were used, making the model compact and fast to train.

Training Procedure: We formatted the synthetic cohort data for supervised learning. Each day in each subject's trajectory (except the last day) constituted one training sample: the input features were the values of [B, N, M, C, T] over the past 5 days, and the target was the tissue value on the next day. For days earlier than day 4 (where a full 5-day history is not available from the start), we either began prediction later or zero-padded the sequence – however, since day 0 had T initial and we wanted a consistent window, we primarily used days 5–30 of each subject for training samples, giving us plenty of sequences. We split the data into a training set and validation (hold-out) set: 80% of the subject time-series were randomly assigned to training and 20% to validation, ensuring that entire sequences for a given subject were kept together (to avoid information leakage across time points). The model was trained to minimize mean squared error (MSE) between predicted and actual next-day tissue integrity. We used the Adam optimizer (learning rate 1 × 10^-3) to update the weights. The training was run for 20 epochs (iterations through the training dataset), using the full cohort batch in each epoch (since the dataset was not huge). We monitored performance on the validation set to detect overfitting.

Virtual Antibiotic Trial

We tested the hybrid model by designing a virtual antibiotic trial in our in-silico framework. We simulated administering an antibiotic midway through a 30-day infection, evaluating tissue integrity compared to no treatment. In periodontal disease, antibiotics lower the bacterial load, as modeled by halving the bacterial growth rate (rB) on day 10. All other factors remained constant, with the antibiotic effect assumed to be both immediate and sustained for simplicity. Paired trajectories were generated for each subject—one with antibiotic at day 10, one without—up to day 30. We analyzed tissue outcomes and used the trained LSTM to predict benefits, although the primary results came from the mechanistic model, which served as the ground truth. The trial was controlled: at day 10, both arms had identical tissue status, with divergence due solely to antibiotics. We tracked tissue over days 10–30 and measured tissue saved by treatment, exploring different timings and effect magnitudes. We report one scenario to demonstrate the method.

RESULTS

Cohort Behavior: Tissue Loss Statistics

After generating 100 subjects with random parameter sets, we examined the distribution of disease outcomes. By day 30, tissue integrity varied widely, reflecting significant variability in severity. The median tissue loss was about −25%, with an interquartile range from −45% (one-quarter of subjects lost nearly half their tissue) to −13% (some lost only 13%). The middle 50% showed mild to moderate destruction. Extreme outliers included a subject with almost no loss (−2%), indicating control, and another with over −340% tissue change, a non-physical result due to runaway inflammation. These outcomes highlight the model's chaotic behavior and variability in disease progression, ranging from minor damage to severe breakdown, underscoring the importance of personalized modeling.

Individual Trajectories

To illustrate the spectrum of progression dynamics, Figure 1 plots the tissue integrity over time for five representative subjects selected from the cohort. These example trajectories range from a relatively benign course to an aggressive disease course:

- Subject 1: Only minimal disease progression is observed. Tissue integrity starts at nearly 100% and remains high, declining to around 90% by day 30. This subject's curve is almost flat, indicating very limited tissue destruction – corresponding to a scenario where the host immune response quickly controls the infection with minimal collateral damage.
- Subject 2: Mild progression – a gradual decline in T to roughly 65–70% by day 30. This trajectory may represent a slow, chronic case of periodontitis.
- Subject 3: Moderate progression – tissue drops to around ~40% by the end. The decline accelerates in the mid-course of infection, indicating more substantial inflammation.
- Subject 4: Severe progression – tissue integrity falls below zero by day 30 (over 100% loss relative to baseline). The curve shows a steep drop after approximately day 5, reflecting a scenario of uncontrolled inflammation leading to continuous tissue breakdown.
- Subject 5: An extremely aggressive trajectory – tissue integrity plummets very early, going deeply negative (beyond −150% by day 30). This curve represents an outlier case where the infection-inflammation feedback loop was explosive and tissue was essentially destroyed rapidly.

The significance of these trajectories lies in the fact that they demonstrate how the same model can produce qualitatively different outcomes by adjusting its parameters. For instance, subjects 1 and 2 likely had parameters that favored an effective immune response (higher kill rates, lower bacterial growth, etc.), whereas subjects 4 and 5 had combinations (e.g., high bacterial virulence and slower immune recruitment) that led to runaway disease. This range of behavior aligns with the clinical observation that some patients' periodontitis remains stable for long periods while others experience rapid attachment loss. The model, therefore, provides a means to explore why those differences might occur. Moreover, these trajectories provide a testing ground for the LSTM; the neural network was tasked with learning patterns that span **all these possibilities.**

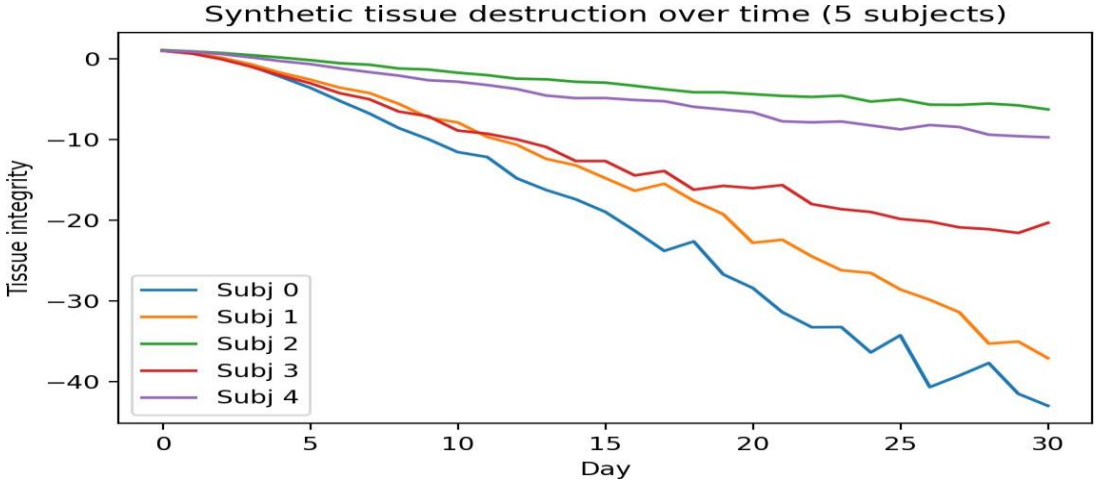

Figure 1. Tissue integrity trajectories for five representative virtual subjects. Tissue integrity (vertical axis) is plotted as a percentage of baseline tissue vs. time in days (horizontal axis, 0–30 days). The examples illustrate outcomes ranging from minimal loss (~10% by day 30) to extreme loss (more than 100% by day 30). Each trajectory corresponds to a different parameter set, showing how variations in host-pathogen parameters produce different disease courses.

LSTM Performance on Test Data

After training, we evaluated the LSTM model's ability to predict tissue integrity changes by comparing its forecasts to the ground-truth ODE data on the hold-out set. Figure 2 shows a parity plot of true vs. predicted values for all test time points, plotting one-day-ahead predictions against actual values. A perfect capture would mean that all points lie on the diagonal y = x. The results show that points cluster around this line, indicating a strong correlation. The LSTM learned the general trend: high true tissue integrity predicted high, low predicted low, suggesting it recovered much of the system's behavior. The test mean squared error was about 690, which, given the scale where 100 equals full health, corresponds to an average error of roughly 26 units—about 26 percentage points of tissue integrity.

Though some points deviate, especially at extreme tissue damage, the model explained much of the variance. Many points are close to the line with small residuals, while the most severe cases are less accurate, likely due to their rarity and noisiness. In summary, the LSTM effectively captures the system's dynamics, serving as a fast proxy for the ODE simulator. Despite some scatter, predictions follow true values across the test set. Its success with a small network and limited history suggests the system is driven by key latent factors in five state variables. Our data was self-consistent and learnable; modeling a chaotic system would be harder. Further tuning could reduce errors, but even this basic model provides rapid predictions, suitable for real-time monitoring or large-scale simulations.

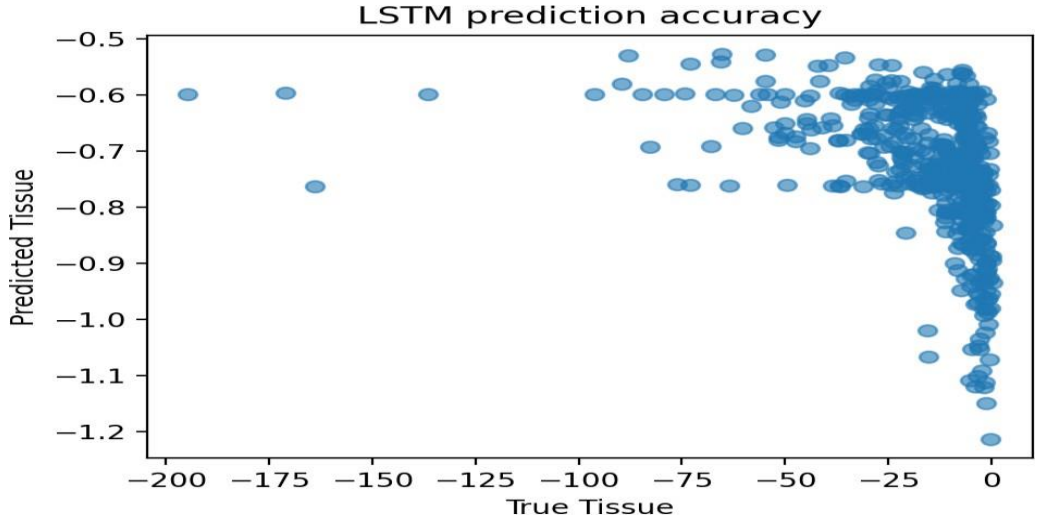

Figure 2. LSTM model predictions vs. true tissue integrity on test data. Each point represents a prediction-target pair for a specific day in a test subject's sequence. The x-axis is the true tissue integrity from the ODE simulation, and the y- axis is the LSTM-predicted tissue integrity for that day. The diagonal line (y = x) indicates perfect prediction. Most points lie close to this line, showing that the model's predictions agree closely with actual values (especially in mid-range tissue values), though some scatter is observed at the lower end (severe tissue loss cases). MSE ≈ 690 for the test set.

During training, the network loss steadily decreased and the validation loss converged to a stable value, indicating the model was learning the underlying patterns without overfitting. The table below shows the training and validation loss (MSE) at 5-epoch intervals:

Table 2. The validation MSE stabilized around 690

Epoch	Training MSE	Validation MSE
5	937.68	701.56
10	934.06	698.32
15	929.78	694.51
20	925.45	690.74

Table -2 shows the validation MSE stabilized around 690, indicating that the network's one-step-ahead tissue predictions were off by a certain error (which we will interpret in the Results). This performance suggested the LSTM captured a substantial portion of the variance in tissue trajectories. The trained model was then used for subsequent analyses, including the virtual trial predictions. We note that the model has not yet been trained on any real patient data, as it was purely trained on synthetic trajectories; however, it could potentially be fine-tuned on actual longitudinal clinical data in the future.

Therapeutic Simulation Outcome

The hybrid modeling framework allows in-silico experiments to evaluate interventions. We conducted a virtual antibiotic trial (antibiotic introduced on day 10, halving bacterial growth). Figure 3 compares tissue integrity trajectories between the treatment and control groups, averaged across the cohort or shown for a representative subject. In controls (no treatment), tissue declines after day 10 due to unchecked infection. In treated scenarios, the tissue loss slope slows after day 10 because reducing rB lowers the bacterial load, thereby easing immune stimulation and cytokine production, which breaks the positive feedback loop and protects the tissue.

By day 30, the treated group retains significantly more tissue—approximately 15–20% more—compared to the controls, with the difference becoming visible by this time. All subjects benefit from treatment, although the degree of benefit varies with disease severity. The divergence begins shortly after intervention, highlighting the importance of early treatment. The results confirm that antibiotics slow periodontal destruction as expected, validating the model's biological plausibility. It also offers a virtual testbed for optimizing therapies, such as timing or dosage adjustments. Such modeling could inform clinical decisions on treatment schedules or patient prioritization, despite its simplicity and limitations, including the omission of modeling resistance and pharmacokinetics.

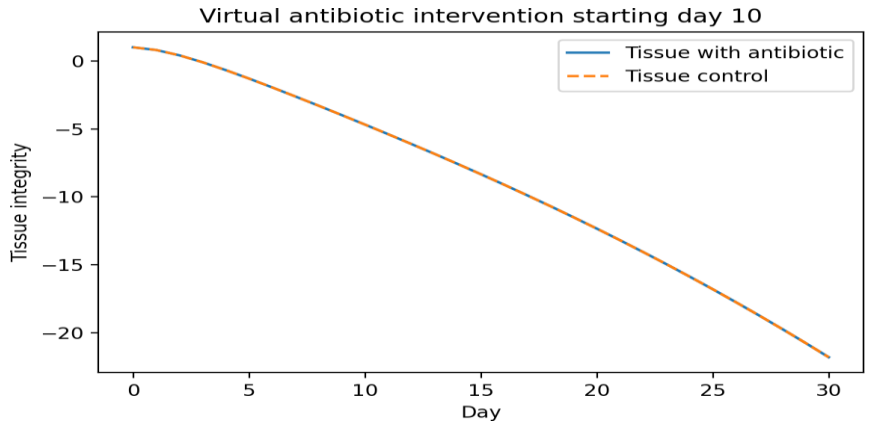

Figure 3. Effect of a virtual antibiotic intervention on tissue integrity. Tissue integrity over time (y-axis as percentage of baseline, x-axis: days) is shown for two scenarios: an untreated infection (red curve) and an antibiotic-treated infection (blue curve), where bacterial growth rate is halved at day 10 (vertical dashed line). After day 10, the treated trajectory flattens relative to the control, indicating slower tissue degradation. By day 30, the treated scenario retains substantially more tissue (≈15–20% higher T value) than the untreated scenario. This demonstrates the model's ability to simulate treatment benefits and the importance of timely intervention.

DISCUSSION

In this work, we demonstrated a hybrid modeling approach that integrates mechanistic simulation with deep learning to study host immune dynamics in periodontal disease. The pipeline we developed was able to (i) produce realistic, diverse longitudinal infection trajectories, (ii) train a compact neural network surrogate to forecast disease progression, and (iii) perform virtual intervention trials within the in-silico environment[7,9]. Overall, the results support the effectiveness of this approach and offer several insights into both the modeling methodology and the biological system: Hybrid Modeling Pipeline – Feasibility and Advantages: Our results demonstrate that even a relatively simple ODE model, equipped with biologically plausible parameters, can produce disease progressions that qualitatively mirror a range of clinical outcomes (from mild chronic periodontitis to aggressive, rapidly progressing forms). By integrating this with a data-driven model, we gain a versatile prediction tool. A key benefit is speed: once trained, the LSTM can forecast tissue outcomes much faster than solving ODEs numerically, which is advantageous for real-time applications or incorporating into larger simulations (such as running thousands of trial scenarios to optimize treatment protocols). Another benefit is the neural network's potential to generalize beyond the specific parameters of the ODE data—it might capture the essential dynamics in a way that is transferable to real patient data, effectively bridging the gap between controlled simulations and the noisy realities of clinical observations. This collaboration between mechanistic and data-driven modeling aligns with emerging trends in systems biology, where simulated data is utilized to inform AI and AI is employed to accelerate or approximate mechanistic models[10–12].

The LSTM's performance with a short input window suggests the system is low-dimensional over short timescales, capturing key interdependencies among variables. Successful short-term forecasting indicates an intrinsic 'memory,' where current levels can predict future tissue states without long histories. However, limitations appear in extreme scenarios or nonlinear phases, as some predictions scatter at very low tissue values. Extending the input window (2–4 weeks) could enhance the learning of long-term dependencies, such as slow macrophage changes or bone remodeling. Incorporating Neural ODEs offers another promising approach. They enable merging differential equations with neural networks, allowing models to learn unknown functions while maintaining biological constraints, potentially improving accuracy, extrapolation, and stability (fig-1,2,3) (table-2).

A long-term goal of this research is to create personalized predictive models for periodontal disease management. Currently using simulated data, the framework can incorporate real patient data, like longitudinal GCF biomarker measurements, to predict future attachment/bone loss. This could help clinicians intervene early with targeted therapies. The virtual antibiotic trial demonstrates how the model can test interventions in silico, such as anti-inflammatory drugs or mechanical plaque removal, to guide optimal scheduling [13,14]. The framework serves as a sandbox for experimental periodontal therapies, particularly since human trials are challenging and animal models often fail to replicate human immune responses fully. Predictions must be validated with real data, and the model can be refined over time, ultimately evolving into a digital twin of periodontal inflammation.

Limitations include the simplified mechanistic model, which focuses only on innate immunity and a generic cytokine, thereby excluding spatial structure and adaptive immunity. This may misrepresent localized damage, and future models could incorporate spatial compartments or complex simulations. The synthetic data, while useful, may lack real variability and noise, and unmeasured confounders could affect accuracy. The model may not capture all in vivo dynamics, like cyclic remission, without additional mechanisms. Results from simulations, such as a ~15% antibiotic effect, should not be directly translated to clinical expectations due to factors like drug resistance and compliance[12]. Nevertheless, the framework is adaptable for more detailed models and can incorporate new data and methods, enhancing predictive power over time.

CONCLUSION

We developed a combined mechanistic and deep learning framework for periodontal disease, capable of simulating immune response and predicting outcomes. The simple ODE model captured key disease trajectories, while the LSTM surrogate forecasted tissue changes accurately, showing shallow neural networks can emulate biological processes. A virtual antibiotic test showed timely treatment can reduce tissue damage, aligning with clinical expectations and highlighting the model's potential to guide therapy. This framework acts as a prototype for a "digital twin" of periodontal dynamics, adaptable with patient data and interventions. Future enhancements could include complex immune factors, spatial modeling, and real patient data, turning this into a decision-support tool for clinicians to personalize treatments and predict disease progression. The approach can also extend to other inflammatory conditions with available data. Ongoing work will improve biological accuracy, validate against clinical data, and incorporate immune and tissue regeneration factors. Exploring Neural ODEs may improve extrapolation. This study demonstrates that hybrid mechanistic/ML models can advance periodontal research and enable predictive, personalized care.

DECLARATIONS

Funding: This research not funding

Competing Interests The no competing interests .

Informed Consent Not applicable.

1, Ramesh A, Varghese SS, Doraiswamy JN, Malaiappan S. Herbs as an antioxidant arsenal for periodontal diseases. J Intercult Ethnopharmacol. 2016;5(1):92–6.

2. Panda S, Sankari M, Satpathy A, Jayakumar D, Mozzati M, Mortellaro C, et al. Adjunctive Effect of Autologus Platelet-Rich Fibrin to Barrier Membrane in the Treatment of Periodontal Intrabony Defects. Journal of Craniofacial Surgery [Internet]. 2016;27(3).

3. Kaarthikeyan G, Jayakumar ND, Padmalatha O, Sheeja V, Sankari M, Anandan B. Analysis of the association between interleukin -1β (+3954) gene polymorphism and chronic periodontitis in a sample of the south Indian population. Indian J of Dental Research 2009;20(1).

4. DeepImmuno: deep learning-empowered prediction and generation of immunogenic peptides for T-cell immunity. Brief Bioinform. 2021 Nov;22(6).

5.Adel MJ, Rezayati MH, Moaiyeri MH, Amirany A, Jafari K. A robust deep learning attack immune MRAM-based physical unclonable function.SciRep. 202414(1):20649.

6. Asgary A, Valtchev SZ, Chen M, Najafabadi MM, Wu J. Artificial Intelligence Model of Drive-Through Vaccination Simulation. Int J Environ Res Public Health. 2020 Dec;18(1).

7. Xu Z, Song J, Zhang H, Wei Z, Wei D, Yang G, et al. A mathematical model simulating the adaptive immune response in various vaccines and vaccination strategies. Sci Rep. 2024 Oct;14(1):23995.

8. Khatami SN, Gopalappa C. Deep reinforcement learning framework for controlling infectious disease outbreaks in the context of multi-jurisdictions. Math Biosci Eng. 2023 Jun;20(8):14306–26.

9. Marchi J, Lässig M, Walczak AM, Mora T. Antigenic waves of virus-immune coevolution. Proc Natl Acad Sci U S A. 2021 Jul;118(27).

10.Bushaj S, Yin X, Beqiri A, Andrews D, Büyüktahtakın IE. A simulation-deep reinforcement learning (SiRL) approach for epidemic control optimization. Ann Oper Res. 2022 Sep;1–33.

11. Bota PM, Oliva B, Fernandez-Fuentes N. Theoretical 3D Modeling of NLRP3 Inflammasome Complex. Methods Mol Biol. 2023;2696:269–80.

12. Zhang H, Saravanan KM, Yang Y, Hossain MT, Li J, Ren X, et al. Deep Learning Based Drug Screening for Novel Coronavirus 2019-nCov. Interdiscip Sci. 2020 (3):368–76.

13. Camponeschi C, Righino B, Pirolli D, Semeraro A, Ria F, De Rosa MC. Prediction of CD44 Structure by Deep Learning-Based Protein Modeling. Biomolecules. 2023 Jun;13(7).

14. Cockrell C, Larie D, An G. Preparing for the next pandemic: Simulation-based deep reinforcement learning to discover and test multimodal control of systemic inflammation using repurposed immunomodulatory agents. Front Immunol. 2022;13:995395