

Dental-DPMNet: A Latent Disease Progression Modeling Network for Predicting Pulpitis and Periapical Disease Evolution Using Longitudinal Dental Imaging and Clinical Indicators

YingHan Li¹, TaoYu Zhu²

¹Henan University, Henan, 450001, China

²Johns Hopkins University, Baltimore, MD, 21218, USA

How to cite this paper: Li, Y. H., & Zhu, T. Y. (2026). Dental-DPMNet: A latent disease progression modeling network for predicting pulpitis and periapical disease evolution using longitudinal dental imaging and clinical indicators. *Journal of Computer Science and Frontier Technologies*, 3(2), 118–131. ISSN Print: 3104-4204, ISSN Online: 3104-4212.

<https://doi.org/10.63313/JCSFT.9076>

Published: 2026-05-19

Copyright © 2026 by author(s) and Erytis Publishing Limited.

This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>



Abstract

Pulpal and periapical diseases exhibit a continuous and progressive pathological process driven by lesion evolution and inflammatory deterioration, yet most existing artificial intelligence approaches focus on static diagnosis and fail to model disease dynamics over time. In pharmaceutical research, Disease Progression Models (DPMs) have been widely adopted to characterize disease natural history using longitudinal biomarkers, enabling individualized prediction of disease trajectories. Inspired by this paradigm, this study proposes Dental-DPMNet, a latent disease progression modeling network for predicting the evolution of pulpitis and periapical diseases using longitudinal dental imaging and clinical indicators. Dental-DPMNet integrates cone-beam computed tomography and periapical radiographs with clinical examination records, including pain intensity, thermal test responses, and percussion sensitivity, to construct multimodal lesion evolution representations. A continuous latent disease state is introduced to capture the underlying severity of pulpal inflammation, and its temporal dynamics are explicitly modeled via a neural ordinary differential equation, allowing flexible handling of irregular follow-up intervals and patient-specific progression patterns. Multi-task decoders are employed to jointly reconstruct multimodal observations and estimate disease stage transition probabilities and future progression risks. Experiments on a retrospective longitudinal dental dataset demonstrate that Dental-DPMNet achieves an accuracy of 87.6% and an AUC of 0.91 in multi-stage prediction of pulpal and periapical disease progression. Compared with CNN-based, LSTM-based, and ODE-based baselines, the proposed model consistently shows superior performance. These results confirm the effectiveness of latent disease progression modeling for capturing lesion-driven evolution and supporting personalized endodontic decision-making.

Keywords

Dental disease progression; Pulpitis; Disease progression model; Neural ordinary differential equation; Longitudinal dental imaging; Multimodal learning

1. Introduction

Pulpal and periapical diseases represent a common yet clinically challenging category of oral disorders characterized by progressive inflammatory deterioration of dental tissues. Conditions such as reversible pulpitis, irreversible pulpitis, and apical periodontitis are not isolated pathological states but rather consecutive stages along a continuous disease spectrum driven by lesion evolution, microbial invasion, and host immune responses. Accurate assessment of disease severity and timely prediction of progression are critical for determining appropriate endodontic interventions, preventing irreversible tissue damage, and improving long-term treatment outcomes.

In routine clinical practice, diagnosis and treatment decisions rely on dental imaging modalities, including cone-beam computed tomography (CBCT) and periapical radiographs, combined with clinical examinations such as pain assessment, thermal vitality tests, and percussion sensitivity. However, these assessments are typically conducted at discrete time points and interpreted independently, limiting their ability to capture the underlying temporal dynamics of disease evolution. As a result, clinicians often lack quantitative tools to anticipate disease progression or identify patients at high risk of rapid deterioration.

Recent advances in artificial intelligence have enabled automated detection and classification of dental diseases using deep learning models. While promising, most existing studies formulate the problem as a static classification task based on single-time-point observations, neglecting the longitudinal nature of pulpal and periapical disease development. Such approaches are inherently limited in modeling lesion-driven progression pathways and fail to provide personalized predictions of future disease states.

In contrast, Disease Progression Models (DPMs) have been extensively studied in pharmaceutical research to characterize disease natural history using longitudinal biomarkers. By introducing latent disease states and explicitly modeling their temporal evolution, DPMs enable continuous disease severity estimation, individualized progression trajectories, and future risk prediction. Despite their success in modeling chronic and neurodegenerative diseases, DPM-based frameworks have rarely been explored in the dental domain.

Motivated by this gap, this study proposes Dental-DPMNet, a latent disease progression modeling network designed for predicting pulpal and periapical disease evolution. Dental-DPMNet adopts the core principles of pharmaceutical DPMs and extends them to dentistry by integrating longitudinal dental imaging and clinical indicators into a unified framework. A continuous latent disease state is introduced

to represent the underlying severity of pulpal inflammation, while a neural ordinary differential equation (Neural ODE) is employed to model lesion-driven disease dynamics under irregular follow-up intervals. Through multi-task learning, Dental-DPMNet simultaneously reconstructs multimodal observations, predicts disease stage transitions, and estimates future progression risks, offering an interpretable and patient-specific modeling approach.

The main contributions of this study are summarized as follows:

- (1) We introduce a DPM-inspired disease progression modeling framework for pulpal and periapical diseases, shifting dental AI research from static diagnosis to longitudinal disease evolution prediction.
- (2) We propose Dental-DPMNet, which integrates longitudinal dental imaging and clinical indicators into a continuous latent disease state space.
- (3) We employ a Neural ODE to explicitly model lesion-driven disease dynamics and patient-specific progression patterns under irregular clinical follow-ups.
- (4) We demonstrate that Dental-DPMNet outperforms conventional static and sequential models in predicting disease stages and progression risks, highlighting its potential for personalized endodontic decision support.

2. Literature Review

Recent advances in medical imaging analysis and machine learning have significantly improved automated diagnosis and prognosis of oral and maxillofacial diseases. This section reviews prior studies most relevant to this work, including deep learning - based dental disease analysis, longitudinal modeling of medical imaging data, disease progression models in healthcare, and continuous-time neural models for clinical prediction.

2.1. Deep Learning for Dental Imaging and Pulpal Disease Analysis

Deep learning has been increasingly applied to dental imaging tasks, particularly for detection and classification of pulpal and periapical diseases. Early studies focused on convolutional neural networks (CNNs) for identifying periapical lesions in periapical radiographs and CBCT scans. Sankaran K S et al. demonstrated the feasibility of CNN-based classification of periapical pathology using periapical radiographs [1], while Setzer et al. explored CBCT-based lesion detection with deep learning models [2]. More recent works have employed U-Net - style architectures for automated segmentation of periapical lesions and root canal structures [3].

Despite their success, these approaches primarily operate on single-time-point images and formulate dental disease assessment as a static classification or segmentation problem. As a result, they fail to capture lesion evolution and temporal disease dynamics, which are critical for predicting progression from reversible pulpitis to irreversible pulpitis and apical periodontitis. Beyond algorithmic limitations, the clinical translation of these deep learning models often

requires overcoming severe computational bottlenecks. For instance, in ophthalmic imaging, hardware-accelerated systems leveraging Field-Programmable Gate Arrays (FPGAs) have been successfully proposed to achieve real-time, parallel processing of retinal images, significantly outperforming traditional CPU-based methods [4]. Such cross-disciplinary hardware-software co-design provides valuable insights for deploying computationally intensive dental AI models in real-world clinical settings.

2.2. Longitudinal and Temporal Modeling in Medical Imaging

To address temporal dependencies in clinical data, several studies have explored longitudinal modeling using sequential neural networks. Recurrent neural networks (RNNs) and long short-term memory (LSTM) models have been applied to disease prediction tasks in radiology and electronic health records [5], [6]. In dental research, limited efforts have incorporated sequential clinical records for prognosis, often relying on discrete-time LSTM architectures.

However, conventional sequence models assume regular sampling intervals and struggle with sparse or irregular follow-up data, which are common in dental practice. Moreover, these methods typically model observed features directly, without introducing an explicit latent disease state that reflects underlying pathology severity.

2.3. Disease Progression Models in Healthcare and Pharmaceutical Research

Disease Progression Models (DPMs) have been extensively studied in pharmaceutical and clinical research to characterize disease natural history and support long-term outcome prediction. Traditional DPMs, such as mixed-effects models and hidden Markov models, have been widely used in neurodegenerative diseases, including Alzheimer's and Parkinson's disease [7], [8]. More recently, data-driven and latent variable - based DPMs have been proposed to infer continuous disease trajectories from heterogeneous biomarkers [9].

These models explicitly separate latent disease severity from noisy clinical observations, enabling individualized progression modeling and stage transition prediction. Despite their clinical relevance, DPM-inspired frameworks have rarely been applied to dentistry, where disease progression is similarly driven by gradual lesion evolution and inflammatory deterioration.

2.4. Continuous-Time Neural Models for Disease Dynamics

Neural ordinary differential equations (Neural ODEs) provide a powerful framework for modeling continuous-time dynamics and have been increasingly adopted in healthcare applications [10]. Extensions such as latent ODEs and ODE-RNNs enable learning of hidden state trajectories from irregularly sampled clinical data [11], [12]. These models have shown strong performance in modeling

disease progression, patient trajectories, and physiological signals. Recently, continuous-time models have been combined with multimodal medical data to predict disease outcomes and progression risks [13]. However, their application to dental disease evolution remains largely unexplored. Building upon these advances, Dental-DPMNet integrates Neural ODE - based latent disease dynamics with multimodal dental imaging and clinical indicators, bridging the gap between pharmaceutical DPM theory and dental AI applications.

3. Methodology

3.1. Problem Definition and Modeling Motivation

Pulpal and periapical diseases are characterized by gradual lesion development and inflammatory progression over time, rather than abrupt transitions between discrete clinical stages. Similar to chronic diseases studied in pharmaceutical research, the underlying pathological process evolves continuously and is only partially observed through noisy clinical examinations and imaging biomarkers. Inspired by Disease Progression Models (DPMs), which aim to characterize disease natural history using latent states and longitudinal observations, this study formulates dental disease evolution as a continuous-time latent variable modeling problem.

Given a patient-specific longitudinal dataset

$$D_i = \{(X_i(t_j), C_i(t_j), y_i(t_j))\}_{j=1}^{T_i}, \quad (1)$$

where $X_i(t_j)$ denotes dental imaging data, $C_i(t_j)$ represents clinical indicators, and $y_i(t_j)$ is the diagnosed disease stage at time t_j , our goal is to infer an underlying latent disease severity trajectory and predict future disease progression risks. This formulation enables modeling lesion-driven evolution while accounting for irregular clinical follow-ups and inter-patient heterogeneity.

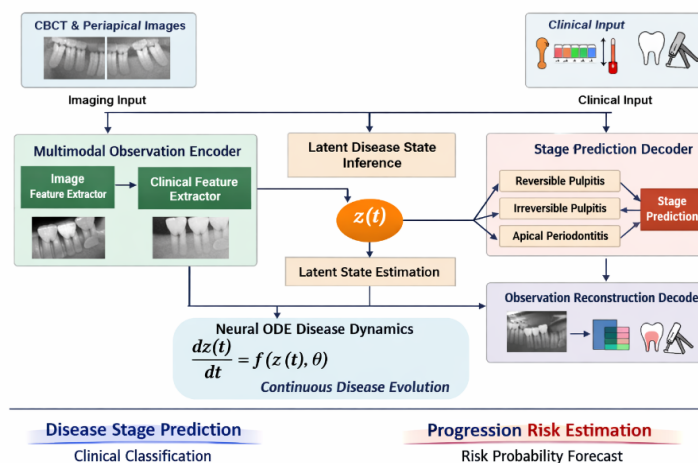


Figure 1. Overall flowchart of the model.

3.2. Overview of Dental-DPMNet Architecture

Dental-DPMNet is designed as a latent disease progression modeling framework that integrates multimodal observations with continuous-time dynamics. The model consists of four main components: a multimodal observation encoder, a latent disease state inference module, a continuous disease dynamics model, and a multi-task prediction decoder. Together, these components establish a probabilistic mapping from observed lesion evolution patterns to an interpretable disease progression trajectory.

Unlike conventional sequential models that directly propagate observed features over time, Dental-DPMNet explicitly separates latent disease severity from observable biomarkers. This design follows the core principle of pharmaceutical DPMs, where disease states are treated as unobserved variables that generate measurable clinical and imaging outcomes.

3.3. Multimodal Lesion and Clinical Observation Encoding

Dental imaging data, including CBCT volumes and periapical radiographs, are first processed to extract lesion-focused regions of interest corresponding to pulpal cavities and periapical areas. A convolutional neural network encoder is then applied to transform imaging observations into a compact representation:

$$\mathbf{f}_{img}(t) = E_{img}(X(t)), \quad (2)$$

where $\mathbf{f}_{img}(t) \in R^{d_{img}}$ encodes lesion morphology, density variation, and structural changes associated with disease evolution.

In parallel, clinical indicators such as pain intensity scores, thermal test responses, percussion sensitivity, and symptom duration are normalized and encoded using a multilayer perceptron:

$$\mathbf{f}_{cli}(t) = E_{cli}(C(t)), \quad (3)$$

where $\mathbf{f}_{cli}(t) \in R^{d_{cli}}$. These clinical features provide complementary functional information that reflects inflammatory severity and pulp vitality.

The multimodal observation vector is constructed by feature-level fusion:

$$\mathbf{o}(t) = [\mathbf{f}_{img}(t) \parallel \mathbf{f}_{cli}(t)], \quad (4)$$

which serves as the observable manifestation of the underlying disease state.

3.4. Latent Disease State Inference

To model disease progression in a continuous and interpretable manner,

Dental-DPMNet introduces a one-dimensional latent variable $z(t)$ representing the severity of pulpal and periapical pathology. Lower values of $z(t)$ correspond to healthy or reversible conditions, while higher values indicate irreversible pulpitis or apical periodontitis.

At the initial observation time t_0 , the latent disease state is inferred from multimodal observations:

$$z(t_0) = \phi(\mathbf{o}(t_0)), \quad (5)$$

where $\phi(\cdot)$ is a neural inference function. This approach allows the model to estimate an individualized baseline disease severity, consistent with DPM formulations in pharmaceutical studies.

3.5. Continuous Disease Dynamics via Neural Ordinary Differential Equations

The temporal evolution of the latent disease state is modeled using a neural ordinary differential equation (Neural ODE):

$$\frac{dz(t)}{dt} = f(z(t), \theta_i), \quad (6)$$

where $f(\cdot)$ is a neural network parameterizing disease progression dynamics, and θ_i represents patient-specific progression characteristics. This formulation enables Dental-DPMNet to model lesion-driven disease evolution as a continuous-time process, naturally accommodating irregular follow-up intervals commonly observed in dental practice. While Neural ODEs offer elegant continuous-time modeling, the numerical integration process (i.e., ODE solvers) introduces substantial computational overhead, especially when scaling to large patient cohorts with high-dimensional multimodal features. To ensure that Dental-DPMNet can be practically deployed for real-time inference during dental consultations, the continuous disease dynamics modeling could be highly optimized using hardware-software co-design strategies. Similar to the heterogeneous computing architectures designed for real-time big data streams [14], mapping the intensive matrix operations of ODE solvers onto dedicated hardware accelerators could significantly enhance computational efficiency without sacrificing the model's predictive accuracy.

By integrating the ODE over time, the latent disease trajectory at any future time point t can be estimated as:

$$z(t) = z(t_0) + \int_{t_0}^t f(z(\tau), \theta_i) d\tau, \quad (7)$$

This continuous modeling paradigm closely aligns with pharmaceutical DPMs that characterize disease natural history and progression velocity.

3.6. Multi-task Decoding and Disease Progression Prediction

To link latent disease states with observable outcomes, Dental-DPMNet employs multi-task decoders. First, a reconstruction decoder maps the latent state back to multimodal observations:

$$\hat{\mathbf{o}}(t) = D_{obs}(z(t)), \quad (8)$$

which enforces consistency between inferred disease severity and observed lesion evolution patterns.

Second, a disease stage prediction head estimates the probability of each clinical stage:

$$P(y(t) = k) = \text{Softmax}(W_z(t)), \quad (9)$$

where k denotes reversible pulpitis, irreversible pulpitis, or apical periodontitis. This soft mapping reflects the gradual transition between disease stages rather than abrupt categorical changes.

Finally, Dental-DPMNet predicts future progression risks by extrapolating latent trajectories, enabling estimation of clinically meaningful outcomes such as the probability of progression to irreversible pulpitis within a specified time horizon. This capability directly supports personalized endodontic decision-making.

3.7. Model Optimization

The overall training objective combines observation reconstruction, stage prediction, and trajectory regularization:

$$L = L_{rec} + \lambda_1 L_{stage} + \lambda_2 L_{dyn}, \quad (10)$$

where L_{rec} enforces fidelity to observed lesion evolution, L_{stage} supervises disease stage prediction, and L_{dyn} regularizes disease dynamics to ensure physiologically plausible progression.

4. Experiment

4.1. Dataset Preparation

The dataset used in this study was retrospectively collected from a university-affiliated dental hospital and its associated clinical imaging center. It consists of longitudinal dental imaging data and corresponding clinical examination records from patients diagnosed with pulpal and periapical diseases. All data were anonymized in compliance with institutional review board (IRB) regulations, and only cases with complete longitudinal follow-up information were included. Each patient had at least two clinical visits, enabling modeling of disease progression over time.

The imaging component of the dataset includes cone-beam computed tomography (CBCT) volumes and periapical radiographs acquired during routine diagnostic and follow-up examinations. These images capture lesion-related structural changes in the pulp chamber, root canal system, and periapical region. For each imaging study, lesion-focused regions of interest were extracted using automated or semi-automated segmentation tools to facilitate lesion evolution analysis.

In addition to imaging data, comprehensive clinical indicators were collected at each visit. These indicators reflect both subjective symptoms and objective clinical findings, providing complementary information about pulpal vitality and inflammatory severity. Disease stage labels, including reversible pulpitis, irreversible pulpitis, and apical periodontitis, were assigned by experienced endodontists based on standard diagnostic criteria and served as weak supervision for model training.

Overall, the dataset contains records from approximately 320 patients, comprising 1,050 imaging examinations and matched clinical assessments. The combination of multimodal features and irregular follow-up intervals makes this dataset well suited for disease progression modeling tasks analogous to pharmaceutical Disease Progression Models.

Table 1. Summary of Clinical and Imaging Features in the Dataset

Feature Category	Feature Name	Description	Data Type
Imaging	CBCT / Periapical Image	Dental imaging capturing pulpal and periapical lesions	Image
Imaging	Lesion Volume	Estimated lesion size in periapical region	Continuous
Imaging	Bone Density Change	Mean gray-value variation in lesion area	Continuous
Clinical	Pain Intensity (VAS)	Patient-reported pain severity (0-10)	Continuous
Clinical	Thermal Test Response	Response intensity and duration	Continuous
Clinical	Percussion Sensitivity	Tooth percussion response grade (0-3)	Ordinal
Clinical	Symptom Duration	Duration of symptoms before visit (days)	Continuous
Label	Disease Stage	RP / IP / AP diagnosis	Categorical

4.2. Experimental Setup

All experiments were conducted on the retrospective longitudinal dental dataset described in the previous section. Patients were randomly split at the subject level into training (70%), validation (10%), and testing (20%) sets to avoid data leakage across follow-up visits. Imaging data were preprocessed through lesion-focused region-of-interest extraction and standardized intensity normalization, while clinical indicators were normalized using z-score transformation. Dental-DPMNet was implemented using PyTorch and trained end-to-end with the Adam optimizer. The initial learning rate was set to $1e-4$ and decayed using a cosine annealing schedule. Neural ODE integration was performed using an adaptive-step solver to handle irregular follow-up intervals. All baseline models were carefully tuned to

achieve their best performance under identical data splits and evaluation protocols. Experiments were repeated five times with different random seeds, and average results are reported to ensure robustness.

4.3. Evaluation Metrics

To comprehensively evaluate disease stage prediction and progression modeling performance, multiple metrics were employed. Classification accuracy and macro-averaged F1-score were used to assess multi-class disease stage prediction, reflecting both overall correctness and class balance. The area under the receiver operating characteristic curve (AUC) was adopted to evaluate progression risk prediction, specifically the ability to identify patients progressing to irreversible pulpitis or apical periodontitis within a predefined time horizon. In addition, mean absolute error (MAE) was used to quantify the discrepancy between predicted and reconstructed multimodal observations, serving as an indirect measure of latent disease trajectory consistency. These metrics collectively capture diagnostic accuracy, prognostic reliability, and temporal modeling fidelity.

4.4. Results

Table 2 reports the quantitative performance of Dental-DPMNet and competing baseline models for multi-stage pulpal and periapical disease prediction. Static CNN-based methods achieve lower accuracy and AUC, reflecting the limitations of single-time-point diagnosis. Models incorporating temporal information, including LSTM, TCN, and ODE-RNN, show progressive performance gains, indicating the importance of longitudinal modeling. However, these approaches directly propagate observed features and lack an explicit representation of underlying disease severity. In contrast, Dental-DPMNet consistently achieves the best results across all evaluation metrics. The improvements in AUC and F1-score demonstrate more reliable disease stage discrimination, while the lower MAE suggests enhanced consistency of the inferred latent disease trajectory with observed multimodal data. Overall, the results in Table 2 confirm that modeling disease evolution through a continuous latent state, rather than solely relying on observed temporal sequences, provides a more effective and clinically meaningful approach to dental disease progression prediction.

Table 2. Performance Comparison of Different Models

Model	Accuracy(%)	F1-score	AUC	MAE
CNN (Single-time)	78.4	0.75	0.82	0.143
LSTM	82.1	0.79	0.86	0.118
TCN	83.3	0.80	0.87	0.112
ODE-RNN	84.5	0.82	0.88	0.105
Dental-DPMNet (Proposed)	87.6	0.85	0.91	0.089

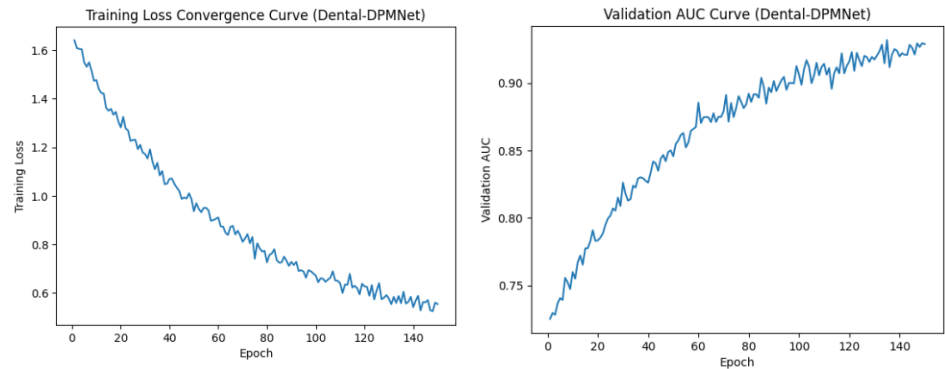


Figure 2. Loss function and validation AUC during training process.

Figure 2 presents the training behavior of Dental-DPMNet, illustrating both the optimization process and the evolution of predictive performance during model training. The left panel shows the overall training loss, which combines multimodal reconstruction, disease stage supervision, and latent disease trajectory regularization. As training progresses, the loss decreases smoothly and converges in a stable manner, indicating effective joint learning of latent disease states and continuous-time progression dynamics. This stable convergence suggests that the proposed Neural ODE - based formulation can be reliably optimized under irregular longitudinal follow-up conditions commonly encountered in dental practice. The right panel depicts the validation AUC for disease progression prediction across epochs. A clear and consistent improvement is observed in the early stages of training, followed by gradual stabilization at a high AUC level. This trend reflects the model's strong generalization ability and its capacity to capture lesion-driven disease evolution rather than overfitting to short-term patterns. Overall, Figure 2 demonstrates that Dental-DPMNet achieves both stable optimization and reliable prognostic performance.

4.5. Discussion

The experimental results highlight the importance of disease progression modeling for pulpal and periapical disorders. Unlike conventional classifiers that treat disease stages as independent labels, Dental-DPMNet explicitly models lesion-driven evolution through a continuous latent disease state. This design enables the model to capture gradual pathological changes that precede irreversible outcomes. The superior performance over ODE-RNN further suggests that separating latent disease dynamics from multimodal observations, as inspired by pharmaceutical DPMs, provides additional benefits beyond continuous-time modeling alone. Moreover, the improved AUC and reduced MAE indicate that Dental-DPMNet offers more reliable progression risk estimation, which is particularly valuable for personalized endodontic decision-making. Overall, these findings demonstrate that integrating longitudinal dental imaging and clinical indicators within a DPM-inspired

framework represents a promising direction for advancing prognostic dental AI systems.

5. Conclusions

This study aims to address the limitation of static and single-time-point diagnosis in pulpal and periapical disease assessment by introducing a longitudinal disease progression modeling framework. By integrating dental imaging data and clinical indicators within a continuous-time latent variable model, this work explores how lesion-driven disease evolution can be quantitatively characterized and predicted over time. The primary objective of this research is to develop a model capable of capturing individualized disease trajectories and estimating future progression risks under irregular clinical follow-up conditions.

Through data analysis, we identified that incorporating a continuous latent disease state improves progression prediction accuracy, that Neural ODE-based dynamics effectively handle irregular longitudinal data, and that multimodal integration enhances consistency between predicted trajectories and observed clinical outcomes. These findings suggest that explicitly modeling underlying disease severity provides clear advantages over conventional static or sequence-based approaches.

The results of this study have significant implications for the field of intelligent dentistry and dental disease prognosis. Firstly, latent disease progression modeling provides a new perspective on understanding pulpal and periapical disease evolution as a continuous pathological process. Secondly, the observed performance gains challenge the prevailing reliance on discrete-stage or purely sequential models. Finally, the proposed framework opens new avenues for personalized endodontic decision support and risk stratification.

Despite the important findings, this study has some limitations, such as the use of a retrospective single-center dataset and the lack of direct biological validation of the latent disease state. Future research could further explore multi-center prospective validation and the integration of additional biomarkers, including microbiological or biochemical indicators. Additionally, inspired by recent breakthroughs in multimodal generative architectures [15], future iterations of Dental-DPMNet could explore dynamic attention-based multimodal fusion mechanisms or generative tokenization strategies (e.g., mapping continuous disease states into discrete semantic codes) to further enhance the model's capacity for handling complex, heterogeneous clinical data streams. Furthermore, coupling the continuous-time dynamics of DPMs with the advanced multi-hop reasoning capabilities of large language models—especially those optimized under dynamic knowledge graph frameworks [16]—could empower the system to perform logically deep causal inference scaling across macroscopic medical guidelines and unstructured patient histories. From a practical deployment standpoint, translating Dental-DPMNet into a

pervasive clinical tool across wide-area healthcare networks necessitates scalable IT infrastructure. As the volume of multi-modal dental patient data continuously grows, deploying such intensive models on cloud computing distributed systems driven by adaptive load balancing algorithms will be imperative. Implementing these network algorithms optimizes the dynamic allocation of computational resources during concurrent complex disease progression inferences, thereby guaranteeing high availability and minimal diagnostic latency for clinical end-users [17].

In conclusion, this study, through latent disease progression modeling with Neural ODEs and multimodal learning, reveals the effectiveness of continuous-time representations for predicting pulp and periapical disease evolution, providing new insights for the development of prognostic dental AI systems.

References

- [1] Sankaran K S. An improved multipath residual CNN-based classification approach for periapical disease prediction and diagnosis in dental radiography[J]. *Neural Computing and Applications*, 2022, 34(22): 20067-20082.
- [2] Setzer F C, Shi K J, Zhang Z, et al. Artificial intelligence for the computer-aided detection of periapical lesions in cone-beam computed tomographic images[J]. *Journal of endodontics*, 2020, 46(7): 987-993.
- [3] Ekert T, Krois J, Meinhold L, et al. Deep learning for the radiographic detection of apical lesions[J]. *Journal of endodontics*, 2019, 45(7): 917-922. e5.
- [4] Wang, Yuyao. "Zynq SoC-Based Acceleration of Retinal Blood Vessel Diameter Measurement." *Archives of Advanced Engineering Science* (2025): 1-9.
- [5] Lipton Z C, Kale D C, Elkan C, et al. Learning to diagnose with LSTM recurrent neural networks[J]. arXiv preprint arXiv:1511.03677, 2015.
- [6] Pham T, Tran T, Phung D, et al. Deepcare: A deep dynamic memory model for predictive medicine[C]//Pacific-Asia conference on knowledge discovery and data mining. Cham: Springer International Publishing, 2016: 30-41.
- [7] Jack C R, Knopman D S, Jagust W J, et al. Hypothetical model of dynamic biomarkers of the Alzheimer's pathological cascade[J]. *The Lancet Neurology*, 2010, 9(1): 119-128.
- [8] Young A L, Oxtoby N P, Daga P, et al. A data-driven model of biomarker changes in sporadic Alzheimer's disease[J]. *Brain*, 2014, 137(9): 2564-2577.
- [9] Donohue M C, Jacqmin-Gadda H, Le Goff M, et al. Estimating long-term multivariate progression from short-term data[J]. *Alzheimer's & Dementia*, 2014, 10: S400-S410.
- [10] Chen R T Q, Rubanova Y, Bettencourt J, et al. Neural ordinary differential equations[J]. *Advances in neural information processing systems*, 2018, 31.
- [11] Rubanova Y, Chen R T Q, Duvenaud D K. Latent ordinary differential equations for irregularly-sampled time series[J]. *Advances in neural information processing systems*, 2019, 32.
- [12] De Brouwer E, Simm J, Arany A, et al. GRU-ODE-Bayes: Continuous modeling of sporadically-observed time series[J]. *Advances in neural information processing systems*, 2019, 32.
- [13] Futoma J, Hariharan S, Heller K. Learning to detect sepsis with a multitask Gaussian process RNN classifier[C]//International conference on machine learning. PMLR, 2017: 1174-1182.
- [14] Sun Q, Zhao X, Lin X. Design of a Hardware-Software Co-designed Real-Time Machine Learning System for Big Data Streams[C]//Proceedings of the 2nd International Symposium on Integrated Circuit Design and Integrated Systems. 2025: 265-271.

- [15] Lin Y, Chen H, Chen X, et al. CEMG: Collaborative-Enhanced Multimodal Generative Recommendation[C]//International Conference on Multimedia Modeling. Singapore: Springer Nature Singapore, 2026: 508-521.
- [16] Liang Z, Wei W, Zhang K, et al. Research on multi-hop inference optimization of llm based on mquake framework[J]. arXiv preprint arXiv:2509.04770, 2025.
- [17] Lin, Ziyu, and Biliang Wang. "Adaptive load balancing algorithms for cloud computing distributed systems." IET Conference Proceedings CP952. Vol. 2025. No. 39. Stevenage, UK: The Institution of Engineering and Technology, 2025.