

FEDERATED GENERATIVE AI FRAMEWORK FOR PERSONALIZED DENTAL CANCER THERAPY A MULTI- MODEL DECISION SUPPORT APPROACH

Mr. Mahajan Sachin Pandurang^{1*}, Dr. Wani Sanjay T.²

¹Dept. of computer science, P. V. P. College Loni, Tal- Rahata, Dist-AhilyaNagar, Maharashtra, Email: sachinmahajan679@gmail.com

²Women College of Home science and BCA, Loni, Tal- Rahata, Dist-AhilyaNagar, Maharashtra Email: stwani14@gmail.com

*Corresponding Author: Mr. Mahajan Sachin Pandurang,

Abstract: Dental cancer poses a major global challenge due to late diagnosis, biological heterogeneity, and variable response. Traditional planning depends on clinician experience and generic protocols. This study proposes a novel Generative AI-enabled precision oncology framework for multi-modal clinical decision support in dental cancer. The proposed framework integrates heterogeneous patient information including radiological imaging, genomic sequencing, histopathological slides, electronic health records, and therapeutic response patterns. The framework combines Vision Transformers (ViTs), Graph Neural Networks (GNNs), Generative Adversarial Networks (GANs), Denoising Diffusion Probabilistic Models (DDPMs), Reinforcement Learning (RL), and Federated Learning (FL) to generate individualized treatment recommendations. Explainable Artificial Intelligence (XAI) mechanisms using SHAP provide transparent reasoning behind AI-driven clinical decisions. Experimental evaluation on 1,847 dental cancer cases demonstrated a clinical concordance of 94.3% with expert oncologist recommendations while reducing treatment planning time by 67%. The proposed multi-modal framework achieved superior conformity indices and enhanced organ-at-risk protection during radiotherapy planning. Furthermore, federated learning enabled privacy-preserving collaborative model development across multiple institutions with minimal performance degradation.

Keywords: Generative Artificial Intelligence, Precision Oncology, Dental Cancer, Multi-modal Learning, Federated Learning, Explainable AI, Deep Learning

1. INTRODUCTION:

Dental cancer, particularly oral squamous cell carcinoma (OSCC), gingival carcinoma, tongue cancer, salivary gland malignancies, and other maxillofacial neoplasms, represents one of the most aggressive and clinically challenging forms of head and neck cancers. Globally, oral cancer has emerged as a major public health concern due to its increasing incidence, high mortality rate, and substantial impact on patients' quality of life. According to recent global cancer statistics, more than 377,000 new cases of oral and dental cancers are diagnosed annually, with a disproportionately higher prevalence observed in developing countries due to tobacco consumption, alcohol abuse, betel nut chewing, poor oral hygiene, viral infections, and genetic predisposition. Despite considerable advancements in surgical oncology, radiotherapy techniques, chemotherapy protocols, immunotherapeutic interventions, and targeted molecular therapies, the overall five-year survival rate of dental cancer patients remains unsatisfactory, ranging between 50% and 60% in many regions of the world [1] [2]. One of the primary reasons for these suboptimal outcomes is the highly heterogeneous nature of dental malignancies. Tumors that appear clinically similar often exhibit significantly different biological behaviors, molecular characteristics, therapeutic responses, and prognostic outcomes.



Traditional treatment planning approaches predominantly rely upon TNM staging systems, histopathological grading, clinician expertise, and standardized population-based treatment guidelines. Although these approaches have been instrumental in improving cancer management over the past decades, they frequently fail to capture the complex inter-individual variations present among patients [3] 4].

Conventional clinical decision-making processes generally emphasize anatomical characteristics of tumors while often overlooking several critical factors that substantially influence therapeutic effectiveness. These factors include genomic mutations, epigenetic alterations, tumor microenvironment characteristics, immune infiltration signatures, radiomic features extracted from advanced imaging modalities, metabolic activity patterns, and individual responses to previous therapies. Consequently, many patients receive generalized treatment regimens that may not be optimally tailored to their specific biological and clinical profiles, resulting in treatment resistance, disease recurrence, unnecessary toxicities, and reduced survival outcomes [5].

In recent years, the rapid development of Artificial Intelligence (AI), particularly deep learning and generative modeling techniques, has revolutionized numerous domains of healthcare and biomedical sciences. Among these emerging technologies, Generative Artificial Intelligence (Generative AI or GenAI) has demonstrated exceptional potential for transforming precision oncology. Unlike conventional discriminative models that primarily focus on classification and prediction tasks, generative models possess the capability to learn underlying data distributions and generate entirely new representations, simulations, and recommendations based on complex multi-dimensional datasets [6].

Generative AI techniques, including Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs), Transformer architectures, and Denoising Diffusion Probabilistic Models (DDPMs), have shown remarkable success in medical image synthesis, drug discovery, disease progression modeling, radiotherapy optimization, and clinical decision support systems. These advanced models are capable of integrating heterogeneous patient information from multiple sources and identifying latent relationships that may remain undetected by traditional analytical approaches.

The concept of precision oncology has emerged as a transformative paradigm in cancer care, emphasizing individualized treatment strategies based on the unique molecular and clinical characteristics of each patient. Precision medicine aims to move beyond the traditional "one-size-fits-all" approach by leveraging genomic sequencing, radiological imaging, histopathological analyses, electronic health records, and real-time treatment response monitoring to develop highly personalized therapeutic interventions. The integration of Generative AI with precision oncology presents an unprecedented opportunity to enhance treatment planning for dental cancer patients. By utilizing multi-modal learning frameworks, AI systems can simultaneously process radiological imaging data, genomic profiles, histopathological features, and longitudinal clinical records to generate patient-specific therapeutic recommendations. Furthermore, advances in Federated Learning enable collaborative model development across multiple healthcare institutions while preserving patient privacy and complying with stringent regulatory requirements such as HIPAA and GDPR. Another critical challenge in the adoption of AI-based clinical decision support systems is the lack of interpretability and transparency. Healthcare professionals often hesitate to rely on "black-box" algorithms when making life-critical decisions. Therefore, Explainable Artificial Intelligence (XAI) has become an essential component of modern medical AI systems. Techniques such as SHapley Additive exPlanations (SHAP) provide clinically meaningful interpretations regarding the contribution of various features toward specific treatment recommendations, thereby increasing clinician trust and facilitating practical implementation in healthcare settings [7] [8] [9].

The proposed framework is designed to process heterogeneous patient information including computed tomography (CT), magnetic resonance imaging (MRI), positron emission tomography (PET), genomic sequencing data, digitized histopathological images, and electronic health records. By learning complex interactions among these modalities, the system generates individualized recommendations concerning surgical interventions, radiotherapy dosimetry, chemotherapy regimens, immunotherapy suitability, and supportive care strategies. The proposed research

therefore represents a significant advancement toward next-generation AI-driven precision oncology systems and highlights the transformative potential of Generative AI in improving personalized cancer care, reducing treatment planning time, enhancing therapeutic outcomes, and supporting clinicians in making informed and evidence-based decisions in dental oncology [10].

Research Contributions:

This study presents a novel and comprehensive contribution to the emerging field of Artificial Intelligence-driven precision oncology by proposing an advanced Generative AI framework specifically designed for personalized treatment planning in dental cancer management. The research addresses several limitations associated with conventional oncology decision-support systems and introduces an integrated multi-modal computational framework capable of generating clinically meaningful and individualized therapeutic recommendations. The first major contribution of this work lies in the development of a novel Generative Artificial Intelligence architecture tailored for dental oncology applications. While previous studies have primarily focused on cancer diagnosis and prognosis prediction, relatively limited attention has been devoted toward AI-driven treatment synthesis and personalized therapeutic recommendation generation.

The proposed framework extends the capabilities of existing AI systems by employing advanced generative models that can simulate and optimize patient-specific treatment pathways based on complex clinical data distributions.

Another significant contribution of this research is the integration of diverse multi-modal datasets into a unified precision oncology platform. The proposed framework simultaneously processes radiological imaging modalities including CT, MRI, and PET scans, digitized histopathological slides, genomic sequencing information, and longitudinal electronic health records. Such comprehensive integration enables the system to capture both phenotypic and genotypic characteristics of dental tumors, thereby facilitating a deeper understanding of tumor heterogeneity and individualized disease progression patterns. Furthermore, this study introduces an innovative combination of multiple state-of-the-art deep learning methodologies within a single clinical decision-support framework. Vision Transformers are employed for extracting high-dimensional radiomic features from volumetric medical images, while Graph Neural Networks model genomic interaction networks and identify actionable mutations and molecular pathways. Generative Adversarial Networks and Denoising Diffusion Models are utilized for treatment plan synthesis and radiotherapy optimization, whereas Reinforcement Learning mechanisms enable dynamic adaptation of treatment strategies according to patient responses and disease evolution [11] [12] [13]

The integration of Federated Learning constitutes another important contribution of this work. Medical data sharing among institutions often faces substantial regulatory, ethical, and privacy-related challenges. The proposed federated framework enables collaborative model training across multiple healthcare institutions without requiring direct exchange of sensitive patient information. This privacy-preserving mechanism significantly enhances the scalability and generalizability of the developed AI system while ensuring compliance with international healthcare regulations. A further contribution of the present research is the incorporation of Explainable Artificial Intelligence mechanisms through SHapley Additive exPlanations (SHAP). Unlike conventional black-box deep learning models, the proposed framework provides transparent explanations regarding the contribution of various clinical features toward specific treatment recommendations. Such interpretability is essential for increasing clinician confidence and promoting practical adoption of AI systems within clinical environments.

The proposed framework has been rigorously validated using a large-scale dataset comprising 1,847 real-world dental cancer patient cases collected from multiple medical institutions. Experimental findings demonstrate that the developed system achieves expert-level performance with a concordance rate exceeding 94% when compared with multidisciplinary oncologist recommendations. Moreover, the framework significantly reduces treatment planning time while simultaneously improving personalization capabilities and treatment optimization. The present study contributes to the growing body of knowledge in precision oncology by demonstrating the feasibility, effectiveness, and clinical applicability of Generative AI technologies for personalized dental cancer management. The proposed

framework establishes a strong foundation for future research in AI-assisted oncology and represents an important step toward the realization of next-generation intelligent healthcare systems capable of delivering individualized, data-driven, and clinically interpretable cancer treatment strategies [13] [14] [15].

2. LITERATURE REVIEW:

Artificial Intelligence has emerged as one of the most transformative technologies in modern oncology, significantly influencing cancer diagnosis, prognosis prediction, treatment planning, and precision medicine. The integration of machine learning, deep learning, and generative models into healthcare systems has enabled researchers and clinicians to process large-scale heterogeneous datasets and identify complex relationships that are often beyond the capability of conventional statistical methods. Over the last decade, numerous studies have demonstrated the potential of AI-based techniques in improving cancer management; however, most of these investigations have primarily focused on disease detection and classification rather than individualized treatment generation. One of the earliest and most significant applications of deep learning in dental oncology involved the utilization of Convolutional Neural Networks (CNNs) for oral cancer detection and histopathological image analysis. Wang et al. (2023) developed a CNN-based automated diagnostic framework for detecting oral squamous cell carcinoma from digitized histopathological images. Their proposed system achieved an impressive diagnostic accuracy of approximately 96.7%, demonstrating the ability of deep learning models to recognize complex cellular and tissue patterns that may be difficult to identify through conventional pathological examination. The study highlighted the effectiveness of artificial intelligence in reducing diagnostic variability and supporting early cancer detection. Nevertheless, the model was primarily limited to classification tasks and did not address therapeutic planning or personalized treatment recommendations.

Similarly, Kumar and Singh (2024) investigated the application of multi-modal deep learning approaches for improving oral cancer staging and prognosis assessment. Their research integrated radiological imaging features with clinical and demographic data to enhance staging accuracy and risk stratification. The study reported a substantial improvement in predictive performance compared with traditional staging methods, indicating that combining multiple data modalities can significantly improve disease characterization. Although their findings emphasized the importance of multi-modal learning in precision oncology, the developed framework remained primarily focused on disease prediction and did not extend toward individualized therapeutic strategy generation [15] [16] [17] [18].

The emergence of Generative Artificial Intelligence has further expanded the capabilities of medical AI systems. Chen et al. (2023) introduced Conditional Generative Adversarial Networks (CGANs) for automated radiotherapy planning in head and neck cancers. Their proposed methodology significantly reduced treatment planning time while maintaining dosimetric quality comparable to expert-generated plans. The study demonstrated that generative models could effectively synthesize radiation dose distributions and optimize treatment parameters based on patient-specific anatomical characteristics. Moreover, the AI-generated plans exhibited improved organ-at-risk sparing without compromising tumor coverage. Despite these promising outcomes, the study primarily focused on radiotherapy optimization and did not incorporate other important modalities such as genomics, histopathology, and longitudinal treatment response data. Recent developments in diffusion models have also shown considerable promise in healthcare applications. Rahman and Thompson (2024) proposed Denoising Diffusion Probabilistic Models (DDPMs) for inverse radiotherapy planning and demonstrated superior optimization performance compared with conventional algorithms. Their findings suggested that diffusion-based approaches could generate smoother and clinically deliverable dose distributions while reducing computational complexity. Nevertheless, these approaches were mainly designed for specific therapeutic tasks and lacked comprehensive multi-modal integration required for holistic precision oncology systems. Transformer-based architectures have similarly gained significant attention in cancer research due to their exceptional capability for modeling long-range dependencies and heterogeneous data relationships. Zhang and Martinez (2023) utilized Transformer-based natural language processing models to extract clinically relevant information from unstructured electronic health records in oncology settings. Their study demonstrated that attention mechanisms could effectively identify critical clinical features and improve the quality of patient datasets available for predictive modeling. Furthermore, Vision Transformers (ViTs) have recently shown

remarkable performance in radiological image analysis and treatment response prediction due to their superior feature extraction capabilities [17] [18].

Another major advancement in AI-driven oncology is the application of Graph Neural Networks (GNNs) for modeling genomic interactions and molecular signaling pathways. Zhao et al. (2024) demonstrated that GNNs could effectively identify actionable mutations and pathway dysregulations associated with therapeutic responses in precision oncology. Their study emphasized the importance of integrating genomic information into clinical decision-support systems to facilitate individualized treatment recommendations and improve patient outcomes.

A critical challenge in medical AI research involves patient privacy and data sharing restrictions across healthcare institutions. To address these concerns, Federated Learning has emerged as a promising paradigm for collaborative AI development. Patel et al. (2024) proposed a privacy-preserving federated learning framework involving multiple cancer centers for developing predictive oncology models. Their study demonstrated that decentralized learning approaches could achieve predictive performances comparable to centralized training while ensuring compliance with healthcare privacy regulations such as HIPAA and GDPR. The framework enabled collaborative model development without directly sharing sensitive patient data, thereby overcoming one of the major limitations associated with large-scale medical AI applications.

Despite the impressive predictive capabilities of deep learning models, their lack of transparency and interpretability has remained a significant barrier to clinical adoption. Explainable Artificial Intelligence (XAI) techniques have therefore received increasing attention in recent years. Anderson et al. (2023) introduced SHAP-based explainable oncology models that provided interpretable feature attribution mechanisms for clinical decision support. Their findings indicated that the incorporation of explainability substantially increased clinician trust and improved acceptance of AI-assisted recommendations. By identifying the relative contributions of individual biomarkers and clinical variables, SHAP-based systems facilitated more transparent and clinically meaningful decision-making processes.

Reinforcement Learning has also been increasingly explored in adaptive cancer treatment optimization. Yoshida et al. (2024) developed reinforcement learning agents capable of dynamically modifying treatment sequences according to real-time patient responses.

Their research demonstrated significant improvements in progression-free survival through adaptive therapy strategies that continuously adjusted treatment protocols based on evolving clinical conditions. Such approaches represent an important step toward fully personalized and adaptive precision oncology systems.

Although the aforementioned studies have significantly advanced the application of Artificial Intelligence in oncology, several important limitations remain. Most existing investigations predominantly concentrate on diagnostic classification, prognostic prediction, image segmentation, or isolated treatment optimization tasks. Very few studies have attempted to develop comprehensive frameworks capable of integrating heterogeneous multi-modal datasets for generating individualized treatment recommendations across the entire spectrum of cancer management.

Furthermore, the majority of current AI models lack simultaneous integration of radiological imaging, genomic information, histopathological characteristics, clinical records, and longitudinal treatment response data. Similarly, limited attention has been given to combining advanced generative models, federated learning mechanisms, explainable AI techniques, and adaptive reinforcement learning strategies within a unified precision oncology framework [18] [19] [20].

Therefore, a substantial research gap exists in the development of intelligent systems capable of synthesizing comprehensive, patient-specific treatment plans for dental cancer management. The present study seeks to address this critical gap by proposing a Generative AI-Enabled Precision Oncology Framework for Multi-Modal Clinical Decision Support in Dental Cancer, integrating state-of-the-art deep learning architectures, privacy-preserving collaborative learning techniques, and explainable decision-support mechanisms to facilitate highly personalized and clinically interpretable treatment recommendations.

3. Proposed Framework




Figure 1: Architecture of Multi-Modal Healthcare Data Processing and Feature Extraction Framework

Proposed Multi-Modal Precision Oncology Framework

The proposed Generative AI-enabled precision oncology framework has been designed as a comprehensive and integrated architecture capable of processing heterogeneous clinical information and generating personalized treatment recommendations for dental cancer patients. Unlike conventional decision-support systems that rely on isolated data modalities or predefined clinical rules, the proposed framework employs advanced deep learning and generative intelligence techniques to model the complex interactions among radiological, molecular, pathological, and clinical characteristics of individual patients.

The entire framework consists of five interconnected layers, namely: data acquisition and preprocessing, multi-modal feature extraction, latent representation learning, generative treatment synthesis, and adaptive treatment optimization. These layers work collaboratively to transform raw patient information into clinically actionable therapeutic recommendations. The architecture has been specifically designed to support precision oncology by considering both phenotypic and genotypic heterogeneity that exists among dental cancer patients.

3.1 Data Acquisition and Preprocessing Layer

The first stage of the proposed framework focuses on acquiring and harmonizing diverse patient information collected from multiple healthcare institutions. Since dental cancer is a highly heterogeneous disease, relying solely on a single source of information often leads to incomplete understanding of disease progression and therapeutic response. Therefore, the framework incorporates multiple modalities of patient information to achieve a comprehensive characterization of individual tumors. The imaging component of the dataset includes Computed Tomography (CT), Magnetic Resonance Imaging (MRI), and Positron Emission Tomography (PET) scans. CT imaging provides detailed anatomical information regarding tumor location, bone invasion, and structural abnormalities. MRI contributes superior soft tissue characterization and enables the identification of tumor boundaries and infiltration patterns. PET imaging provides valuable functional information by quantifying metabolic activity and identifying potential metastatic regions.

In addition to radiological imaging, the framework incorporates digitized histopathological whole-slide images obtained from biopsy specimens. Histopathology remains the gold standard for cancer diagnosis and provides essential information regarding cellular morphology, tumor differentiation, mitotic activity, necrosis, and immune infiltration patterns. The genomic component of the framework consists of whole exome sequencing and transcriptomic profiling data. Genomic information enables the identification of driver mutations, molecular biomarkers, signaling pathway dysregulations, and therapeutic targets that are critically important for personalized treatment planning. Furthermore, electronic health records containing demographic information, comorbidities, previous therapeutic interventions, laboratory findings, and longitudinal follow-up data are incorporated into the framework to facilitate comprehensive clinical understanding. The treatment outcome database further enriches the dataset by providing historical information regarding therapeutic responses, recurrence patterns, toxicity profiles, and survival outcomes [20] [21] [22].

The experimental dataset utilized in the present study consists of 1,847 dental cancer patient cases collected from three tertiary healthcare institutions. To ensure robust model development and unbiased evaluation, the dataset was divided into training, validation, and testing subsets. Approximately 70% of the data were used for model training, 15% for hyperparameter tuning and validation, and the remaining 15% for independent performance evaluation. Such partitioning enables effective model generalization while minimizing overfitting. Before model training, extensive preprocessing procedures were performed. Medical images underwent normalization, intensity correction, resampling, and segmentation procedures. Histopathological slides were standardized through stain normalization and patch extraction methods. Genomic data were subjected to quality control measures, variant calling pipelines, and functional annotation using publicly available molecular databases. Missing clinical information was handled through imputation techniques, and all datasets were subsequently standardized and transformed into machine-readable representations.

3.2 Multi-Modal Feature Extraction

Following data acquisition and preprocessing, the framework performs advanced feature extraction from each modality through specialized deep learning architectures. Since different data types possess unique structural characteristics, dedicated neural models are employed to capture modality-specific information while preserving clinically relevant relationships.

3.3 Imaging Feature Extraction Using Vision Transformers

Medical imaging data constitute one of the most important components in dental cancer management due to their ability to reveal tumor morphology, invasion patterns, and anatomical relationships. Traditional convolutional neural networks often face limitations in capturing long-range spatial dependencies present in volumetric imaging data. To address this challenge, the proposed framework utilizes a three-dimensional Vision Transformer (ViT-3D) architecture. The ViT-3D model consists of twelve transformer encoder blocks with a hidden embedding dimension of 768. Volumetric medical images are partitioned into patches of size $16 \times 16 \times 16$, which are subsequently embedded into high-dimensional feature representations. The self-attention mechanism of the transformer architecture enables effective modeling of global contextual relationships across the entire imaging volume. The imaging module extracts several clinically meaningful features, including tumor morphology, volumetric characteristics, invasion depth, tissue heterogeneity, vascularity patterns, and spatial relationships between tumors and surrounding anatomical structures. Such information is critically important for surgical planning and radiotherapy optimization.

3.4 Genomic Feature Extraction Using Graph Neural Networks

Cancer progression and therapeutic responses are strongly influenced by complex molecular interactions occurring at the genomic level. Therefore, genomic information plays a fundamental role in precision oncology. To model intricate gene interaction networks, the proposed framework employs Graph Neural Networks (GNNs). In this representation, genes are modeled as nodes, while molecular interactions and signaling pathways are represented as edges within a biological graph structure. The GNN architecture enables the identification of actionable genetic mutations, dysregulated signaling pathways, gene-expression signatures, and predictive biomarkers associated with

therapeutic outcomes. The model further identifies molecular determinants of treatment resistance and immunotherapy responsiveness.

3.5 Histopathological Feature Extraction

Histopathological examination provides microscopic insights into tumor biology and remains indispensable in dental oncology. To extract meaningful pathological representations, the framework utilizes a deep ResNet-152 architecture pre-trained on large-scale pathology datasets. The histopathology module automatically identifies cellular morphology patterns, nuclear atypia, mitotic figures, stromal architecture, immune cell infiltration, angiogenesis patterns, and tumor microenvironment characteristics.

3.6 Multi-Modal Feature Fusion

Although each data modality provides valuable information independently, true precision medicine requires simultaneous integration of all available patient information. Therefore, the extracted imaging, genomic, histopathological, and clinical features are combined through an advanced feature fusion strategy. The proposed framework employs a Variational Autoencoder (VAE) to project heterogeneous information into a unified latent representation space. The VAE effectively learns the underlying probability distributions governing patient-specific disease characteristics and compresses multi-modal information into a 512-dimensional latent embedding space. This latent representation captures complex relationships among anatomical, molecular, pathological, and clinical variables while reducing redundancy and preserving clinically relevant information. The resulting embedding serves as the foundation for subsequent generative treatment synthesis and optimization [22] [23].

4. Generative AI Architecture

The core innovation of the proposed framework lies in its advanced Generative Artificial Intelligence architecture, which synthesizes personalized treatment plans by learning complex relationships between patient characteristics and therapeutic outcomes.

Unlike traditional predictive models that merely estimate risk scores or classify disease stages, the proposed generative framework actively creates individualized treatment recommendations by simulating potential therapeutic strategies and optimizing them according to patient-specific constraints.

4.1 Conditional Diffusion-Based Treatment Generation

The primary treatment synthesis engine is based on Conditional Denoising Diffusion Probabilistic Models (DDPMs). Diffusion models have recently emerged as one of the most powerful generative paradigms due to their superior stability, diversity, and generation quality. The proposed diffusion model utilizes a U-Net backbone architecture and performs iterative denoising through one thousand diffusion steps. During training, random noise is gradually introduced into treatment representations, and the network learns the reverse process of reconstructing optimal treatment embeddings conditioned on patient-specific features.

Classifier-free guidance mechanisms are incorporated to improve generation quality and enable flexible conditioning on clinical information.

The generated treatment recommendations include:

- Surgical resection margins,
- Radiation dose distributions,
- Fractionation schedules,

By learning complex relationships from historical treatment outcomes, the diffusion model generates clinically meaningful and individualized therapeutic plans that closely resemble expert oncologist recommendations.

4.2 Generative Adversarial Network Module

To further enhance treatment simulation and optimization, the framework incorporates Generative Adversarial Networks. GANs consist of two competing neural networks: a generator and a discriminator. The generator attempts to synthesize realistic treatment plans, whereas the discriminator evaluates their clinical plausibility.

Within the proposed framework, GANs are utilized for several important tasks including treatment response simulation, radiotherapy field generation, dose distribution optimization, and virtual patient modeling. Particularly in radiotherapy planning, GANs enable rapid generation of clinically acceptable dose distributions while minimizing exposure to surrounding healthy tissues. The generated treatment plans demonstrate superior conformity indices and improved organ-at-risk protection compared with traditional planning approaches.

4.3 Reinforcement Learning-Based Adaptive Optimization

Cancer treatment is inherently dynamic because tumor characteristics and patient responses continuously evolve during therapy. Consequently, static treatment protocols may not always provide optimal outcomes. To address this limitation, the proposed framework incorporates Reinforcement Learning (RL) mechanisms for adaptive treatment optimization. The RL agent continuously monitors patient responses, disease progression patterns, treatment toxicities, and biomarker dynamics. Based on these observations, the agent dynamically modifies therapeutic recommendations throughout the treatment course. The optimization process is guided by clinically meaningful reward functions, primarily including:

4.4 Tumor Control Probability (TCP)

TCP estimates the probability of achieving complete tumor eradication following treatment. Maximizing TCP ensures improved local disease control and survival outcomes.

4.5 Normal Tissue Complication Probability (NTCP)

NTCP quantifies the likelihood of treatment-induced toxicity and adverse effects in surrounding healthy tissues. Minimization of NTCP contributes to improved quality of life and reduced treatment-related morbidity.

By simultaneously maximizing TCP and minimizing NTCP, the reinforcement learning module achieves an optimal balance between therapeutic efficacy and patient safety.

The adaptive capabilities of reinforcement learning further facilitate personalized treatment modifications in response to emerging clinical information, thereby supporting the implementation of truly individualized precision oncology strategies [23] [24] [25].

Federated Learning Framework




Figure 2: Secure Federated Learning Architecture for Privacy-Preserving Artificial Intelligence in Healthcare

4.6 Federated Learning Framework

The application of Artificial Intelligence in healthcare frequently encounters significant challenges related to data privacy, security, and regulatory compliance. Medical institutions often possess highly valuable patient datasets; however, legal restrictions and ethical concerns prevent the direct sharing of sensitive healthcare information across organizations. These limitations considerably reduce the availability of large-scale datasets required for developing robust and generalizable deep learning models. To address these challenges, the proposed framework incorporates a Federated Learning (FL) architecture that enables collaborative model development while preserving patient privacy and institutional data ownership. Federated Learning represents a decentralized machine learning paradigm in which participating institutions train models locally using their own datasets without transferring raw patient information to a central server. Instead of sharing sensitive clinical records, only model parameters and encrypted gradient updates are communicated during the training process. This strategy significantly reduces privacy risks and facilitates large-scale collaboration among multiple healthcare centers.

In the proposed precision oncology framework, data from three tertiary cancer institutions were utilized through a federated learning infrastructure. Each participating center independently trained local versions of the deep learning models using its respective imaging, genomic, histopathological, and clinical datasets. After local training, model parameters were securely transmitted to a central aggregation server, where global model updates were generated and redistributed for subsequent training rounds.

The framework employs the Federated Averaging (FedAvg) algorithm as the primary aggregation mechanism. FedAvg combines local model parameters through weighted averaging, thereby enabling efficient convergence while accounting for differences in dataset sizes across institutions. The use of FedAvg significantly enhances model generalization by incorporating diverse patient populations and reducing institution-specific biases.

In addition, Secure Multi-Party Computation (SMPC) techniques were integrated into the aggregation process. SMPC enables multiple institutions to collaboratively compute global model updates without exposing their individual model parameters to other participants. The encrypted aggregation strategy provides an additional security layer against adversarial attacks and unauthorized access to sensitive healthcare information.

The federated architecture was specifically designed to comply with major healthcare regulations, including the Health Insurance Portability and Accountability Act (HIPAA), General Data Protection Regulation (GDPR), and other international standards governing patient data protection. Consequently, the proposed framework offers a practical and legally compliant solution for large-scale AI deployment in oncology settings. Experimental results demonstrated that the federated framework achieved predictive performances remarkably close to those obtained through centralized learning approaches. The proposed federated model achieved approximately 93.1% concordance with expert recommendations compared with 94.3% for centralized training, corresponding to a performance degradation of only 1.2%. This minimal reduction indicates that privacy-preserving collaborative learning can be successfully implemented without substantially compromising predictive accuracy [25] [26] [27].

5. Explainable Artificial Intelligence Module

Despite the remarkable predictive capabilities of modern deep learning models, their widespread clinical adoption remains limited due to concerns regarding transparency and interpretability. Most advanced neural networks operate as "black-box" systems, producing predictions without providing clear explanations regarding the underlying reasoning processes. In healthcare, where clinical decisions directly influence patient outcomes, explainability is an essential requirement rather than an optional feature. To address this challenge, the proposed framework integrates an Explainable Artificial Intelligence (XAI) module capable of providing interpretable and clinically meaningful explanations for every generated treatment recommendation. The primary objective of this module is to enhance clinician trust, improve model transparency, and facilitate evidence-based decision-making.

The explainability framework is based on SHapley Additive exPlanations (SHAP), which originates from cooperative game theory. SHAP values quantify the contribution of each input feature toward the final prediction generated by the AI model. By decomposing model outputs into feature-specific importance scores, SHAP enables clinicians to understand why particular treatment recommendations have been proposed. For each patient, the system generates detailed explanations describing the relative influence of genomic biomarkers, radiological characteristics, histopathological findings, and clinical variables on treatment planning decisions. Such explanations provide valuable insights into the biological mechanisms underlying AI-generated recommendations and facilitate greater confidence among healthcare professionals.

The explainability module primarily provides three categories of information:

1. **Feature Importance Analysis:** Identification of the most influential variables affecting treatment decisions.
2. **Treatment Rationale Interpretation:** Biological and clinical reasoning supporting specific therapeutic recommendations.
3. **Confidence Estimation:** Quantification of model certainty and prediction reliability.

Experimental findings indicated that several biomarkers consistently exhibited strong influences on treatment recommendations. Among these, Programmed Death Ligand-1 (PD-L1) expression emerged as one of the most significant determinants of immunotherapy suitability. Similarly, Tumor Mutation Burden (TMB) demonstrated considerable importance in predicting responsiveness to immune checkpoint inhibitors.

Other highly influential biomarkers included Epidermal Growth Factor Receptor (EGFR) mutations, genomic pathway alterations, immune infiltration signatures, and radiomic indicators of tumor aggressiveness [27] [28].

6. EXPERIMENTAL RESULTS AND PERFORMANCE EVALUATION

The proposed framework was extensively evaluated using an independent testing cohort comprising 277 previously unseen dental cancer cases derived from the original dataset of 1,847 patients. Multiple performance indicators were assessed, including concordance with expert recommendations, computational efficiency, radiotherapy quality, personalization capability, and explainability.

6.1 Clinical Performance Analysis

The primary evaluation metric was concordance between AI-generated recommendations and multidisciplinary tumor board decisions.

Table 1. Clinical Performance Comparison

Method	Concordance (%)	Planning Time (hours)
Individual Clinician	78.2	4.2
Protocol-Based Planning	82.5	3.8
Early GAN Models	86.4	2.1
Proposed Framework	94.3	1.4



Figure 3. Comparative analysis of clinical concordance and treatment planning time for different treatment planning approaches

The proposed framework achieved a remarkable concordance rate of **94.3%**, significantly outperforming conventional clinical approaches and earlier generative models. Furthermore, treatment planning time was reduced from 4.2 hours to only 1.4 hours per case. The substantial reduction in planning time can be attributed to automated feature extraction, parallelized multi-modal processing, and efficient diffusion-based treatment synthesis mechanisms.

6.2 Multi-Modal Performance Analysis

To evaluate the contribution of each modality, ablation studies were conducted.

Table 2. Multi-modal Performance Comparison

Model	Concordance (%)
Imaging Only	87.6
Genomics Only	82.3
Histopathology Only	79.8
Proposed Multi-modal Framework	94.3

The results clearly demonstrate that integrating heterogeneous patient information significantly improves treatment personalization and predictive accuracy. The proposed multi-modal framework achieved an improvement of more than 6% compared with the best single-modality model. These findings strongly support the hypothesis that precision oncology requires comprehensive phenotypic and genotypic characterization rather than isolated analyses.

6.3 Radiotherapy Quality and Organ Protection

One of the major advantages of the proposed framework lies in its capability to optimize radiotherapy planning. The AI-generated treatment plans demonstrated superior dose conformity and substantially improved organ-at-risk protection. In particular, the framework achieved an average parotid gland dose reduction of 8.3 Gy, which may significantly decrease the incidence of xerostomia and improve post-treatment quality of life. Importantly, this enhanced organ sparing was achieved without compromising tumor coverage, thereby maintaining therapeutic effectiveness.

6.4 Computational Efficiency

The proposed framework exhibited remarkable computational efficiency. The average treatment planning time was reduced by approximately **67%**, enabling rapid treatment generation and facilitating same-day planning in urgent clinical scenarios.

Pure inference time required only a few minutes per patient case when executed using high-performance graphical processing units. Such computational efficiency may substantially reduce clinician workload and improve healthcare resource utilization.

6.5 Personalization Capability

A major objective of precision oncology is to move beyond generalized treatment protocols and generate individualized therapeutic strategies. Treatment diversity analysis demonstrated the remarkable personalization capabilities of the proposed framework. Among fifty patients sharing identical TNM staging characteristics, the AI system generated **47 distinct treatment plans**, varying in surgical approaches, radiotherapy doses, fractionation schedules, and systemic therapies. These findings clearly indicate that conventional staging systems alone are insufficient for comprehensive treatment planning and further emphasize the importance of integrating molecular and clinical heterogeneity into decision-making processes [27] [28] [29].

7. DISCUSSION:

The present study demonstrates the immense potential of Generative Artificial Intelligence for transforming precision oncology and personalized dental cancer management. By integrating radiological imaging, genomic information, histopathological characteristics, and longitudinal clinical records, the proposed framework provides a holistic representation of patient-specific disease characteristics.

One of the major strengths of the proposed framework is its ability to generate highly personalized treatment recommendations rather than merely predicting disease outcomes. The integration of diffusion models, GANs, reinforcement learning, and explainable AI mechanisms enables comprehensive treatment synthesis and adaptive optimization throughout the therapeutic process. Another important advantage lies in the substantial reduction of clinician workload and treatment planning time. Automated generation of clinically meaningful treatment recommendations may improve healthcare efficiency and allow oncologists to focus on complex clinical decision-making and patient care. The incorporation of Federated Learning further enhances the practical applicability of the framework by enabling privacy-preserving collaborative learning across multiple institutions. This capability is particularly important in oncology, where large and diverse datasets are essential for developing robust AI systems.

Despite these encouraging findings, several challenges remain.

The heterogeneous nature of medical data continues to present significant difficulties in data harmonization and model generalization. Variations in imaging protocols, sequencing technologies, and clinical practices may influence model performance.

Another limitation involves the relatively limited availability of data for rare dental cancer subtypes. Future studies should include larger multi-institutional datasets to improve generalizability and robustness. Regulatory approval and clinical implementation also represent important challenges. Comprehensive prospective clinical trials

are necessary to validate the safety, effectiveness, and ethical implications of AI-driven treatment planning systems [28] [29] [30].

8. CONCLUSION:

This study presents a novel Generative AI-Enabled Precision Oncology Framework for Multi-Modal Clinical Decision Support in Dental Cancer, representing a significant advancement toward intelligent and personalized cancer management. By integrating advanced deep learning architectures, generative modeling techniques, federated learning mechanisms, and explainable artificial intelligence, the proposed framework successfully generates individualized treatment recommendations with expert-level performance.

Experimental findings demonstrated a clinical concordance rate of 94.3%, a 67% reduction in treatment planning time, improved radiotherapy quality, enhanced organ protection, and highly interpretable decision-making capabilities. The incorporation of privacy-preserving federated learning further ensures secure and scalable deployment across multiple healthcare institutions, while explainable AI mechanisms substantially improve clinician trust and acceptance. The proposed framework highlights the transformative potential of Generative Artificial Intelligence in precision oncology and establishes a strong foundation for future AI-driven healthcare systems. With continued advancements in multi-modal learning, adaptive intelligence, and clinical validation, such technologies may fundamentally reshape the future of dental cancer treatment, enabling truly personalized, efficient, and data-driven cancer care.

Future research may involve the development of digital twin oncology systems, where virtual patient models continuously simulate disease progression and treatment responses in real time.

CONFLICT OF INTEREST:

The authors declare that there is no conflict of interest.

REFERENCES

1. Wang, S., Liu, J., & Chen, M. (2023). Deep learning-based automated detection and classification of oral squamous cell carcinoma in histopathological images. *IEEE Transactions on Medical Imaging*, 42(8), 2341–2353.
2. Kumar, R., & Singh, A. (2024). Multi-modal deep learning framework for accurate staging of oral cavity cancers using clinical and imaging data. *IEEE Journal of Biomedical and Health Informatics*, 28(3), 1456–1467.
3. Chen, L., Park, K., & Yamamoto, T. (2023). Conditional generative adversarial networks for automated radiation therapy planning in head and neck cancers. *IEEE Transactions on Radiation and Plasma Medical Sciences*, 7(6), 634–645.
4. Rahman, M., & Thompson, E. (2024). Diffusion probabilistic models for inverse treatment planning optimization in intensity-modulated radiation therapy. *IEEE Transactions on Biomedical Engineering*, 71(4), 1123–1135.
5. Liu, H., Zhang, Y., & Anderson, S. (2024). Vision transformer architectures for temporal analysis of cancer treatment response patterns. *IEEE Access*, 12, 45678–45691.
6. Zhang, Q., & Martinez, C. (2023). BERT-based natural language processing for clinical information extraction from oncology electronic health records. *IEEE Journal of Translational Engineering in Health and Medicine*, 11, 234–245.
7. Patel, N., Johnson, K., & Williams, R. (2024). Federated learning framework for privacy-preserving collaborative cancer prediction model development. *IEEE Transactions on Information Forensics and Security*, 19, 3456–3468.
8. Anderson, J., Murphy, L., & Foster, D. (2023). Explainable AI integration in oncology clinical decision support systems using SHAP analysis. *IEEE Intelligent Systems*, 38(5), 67–79.
9. Yoshida, T., & Kim, S. (2024). Reinforcement learning agents for adaptive cancer treatment sequencing and protocol optimization. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 32(2), 890–902.
10. Zhao, F., Lee, B., & Schmidt, K. (2024). Graph neural networks for modeling gene interaction pathways in precision oncology. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 21(1), 178–190.
11. Esteva, A., Robicquet, A., Ramsundar, B., et al. (2021). A guide to deep learning in healthcare. *Nature Medicine*, 27(1), 29–38.
12. Topol, E. (2019). High-performance medicine: The convergence of human and artificial intelligence. *Nature Medicine*, 25(1), 44–56.
13. Kather, J. N., Pearson, A. T., Halama, N., et al. (2020). Deep learning can predict microsatellite instability directly from histology in gastrointestinal cancer. *Nature Medicine*, 26(7), 1054–1056.
14. Jumper, J., Evans, R., Pritzel, A., et al. (2021). Highly accurate protein structure prediction with AlphaFold. *Nature*, 596, 583–589.
15. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521, 436–444.

16. Vaswani, A., Shazeer, N., Parmar, N., et al. (2017). Attention is all you need.
17. *Advances in Neural Information Processing Systems*, 30, 5998–6008.
18. Goodfellow, I., Pouget-Abadie, J., Mirza, M., et al. (2014). Generative adversarial nets. *Proceedings of NeurIPS*, 2672–2680.
19. Ho, J., Jain, A., & Abbeel, P. (2020). Denoising diffusion probabilistic models.
20. *Proceedings of NeurIPS*, 33, 6840–6851.
21. Dosovitskiy, A., Beyer, L., Kolesnikov, A., et al. (2021). An image is worth 16×16 words: Transformers for image recognition at scale. *International Conference on Learning Representations (ICLR)*.
22. Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. *MICCAI*, 234–241.
23. Litjens, G., Kooi, T., Bejnordi, B. E., et al. (2017). A survey on deep learning in medical image analysis. *Medical Image Analysis*, 42, 60–88.
24. Rajpurkar, P., Chen, E., Banerjee, O., & Topol, E. (2022). AI in health and medicine. *Nature Medicine*, 28, 31–38.
25. McKinney, S. M., Sieniek, M., Godbole, V., et al. (2020). International evaluation of an AI system for breast cancer screening. *Nature*, 577, 89–94.
26. Rieke, N., Hancox, J., Li, W., et al. (2020). The future of digital health with federated learning. *NPJ Digital Medicine*, 3, 119