DentAI Vision: AI-Powered Dental X-Ray Analysis for Enhancing Trust and Patient Education

Maha Kesibi, *Queen's University*, 21mk93@queensu.ca
Leila Salem, *Queen's University*, 22tt43@queensu.ca
Het Buddhdev, *Queen's University*, 22hvb@queensu.ca
Kamran Jornacion, *Queen's University*, 22qfc4@queensu.ca
Elliott Vince, *Queen's University*, elliott.vince@queensu.ca

***Abstract*—This study addresses the challenge of improving trust between patients and dental professionals by leveraging AI-driven analysis of panoramic dental X-rays. With dental caries being one of the most prevalent oral diseases, early and accurate detection is crucial for effective treatment and prevention. Deep learning algorithms were utilized to detect caries, specifically employing the YOLOv5-Small object detection model, optimized for real-time inference. The methodology involved dataset preprocessing, data augmentation techniques such as horizontal flipping, and model training on the DENTEX MICCAI 2023 dataset. In addition to automated diagnosis, a chatbot powered by DeepSeek-Llama and LangChain was integrated into the platform, providing users with reliable, evidence-based dental health information sourced from over 20 accredited references. The findings demonstrate that the proposed AI system can achieve high diagnostic accuracy while fostering patient education and transparency. This research highlights the potential of AI-powered dental diagnostics in reducing the need for costly second opinions, improving patient-dentist relationships, and promoting informed decision-making in oral healthcare.**

I. Introduction

Oral diseases, including dental caries, are among the most prevalent health issues globally, affecting approximately 3.5 billion people. Specifically, dental caries impact around 2.3 billion individuals in their permanent teeth, making it one of the most common chronic diseases worldwide [1]. If left untreated, dental caries can lead to complications such as infections, pain, and tooth loss. Despite advancements in dental care, access to accurate, timely, and affordable diagnoses remains a significant challenge, particularly in low-income and underserved communities [2].

One of the key factors influencing patient outcomes in dentistry is trust between patients and dental professionals. Trust plays a crucial role in patient adherence to treatment plans and their willingness to seek care. However, negative experiences, misdiagnoses, and costly treatment plans often lead patients to seek second opinions, which can be expensive and time-consuming. Research suggests that patients who receive multiple opinions are more likely to adhere to treatment recommendations, yet many face barriers to accessing additional consultations due to geographic or financial constraints [3], [4].

Artificial intelligence (AI) and deep learning have demonstrated significant potential in enhancing diagnostic accuracy in medical imaging. A comprehensive meta-analysis revealed that deep learning algorithms achieved a mean sensitivity of 96.3% and a mean specificity of 93.3% across various pathology identifications, underscoring their efficacy in medical diagnostics [5]. Regarding patient trust, the integration of AI into medical imaging has yielded mixed perceptions. Some studies indicate low trust in healthcare systems to use AI responsibly [3], while others show that nearly two-thirds of consumers would trust a diagnosis from AI over that of a human doctor [4]. These findings suggest that while AI has the potential to enhance diagnostic accuracy, fostering patient trust requires careful implementation and communication strategies. By leveraging AI, DentAI Vision aims to provide a low-cost, accessible second opinion for patients and support dentists in patient education and decision-making.

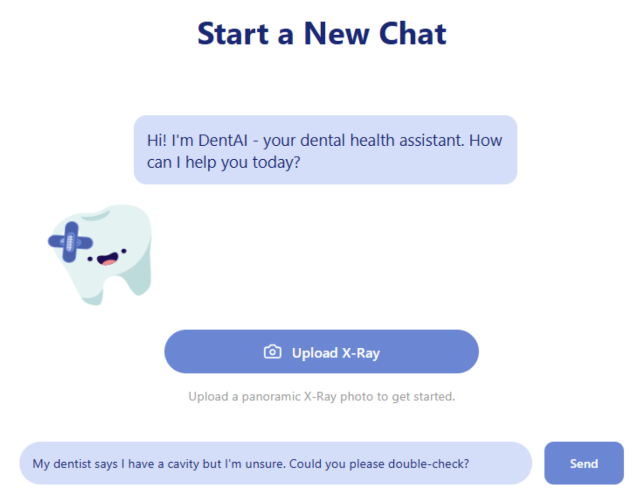

Fig. 1. DentAI Vision chat interface where users can initiate a session by uploading an X-ray and inputting their concerns or queries.

A. Problem Definition

Given the high prevalence of dental caries and the role of trust in treatment adherence, DentAI Vision addresses two core problems:

1) **Lack of Accessible Second Opinions** – Patients often seek confirmation of diagnoses before undergoing dental procedures. However, obtaining a second opinion can be expensive and time-consuming, discouraging patients from making informed decisions [**?**]. An AI-driven system that offers *instant, cost-free analysis of dental X-rays* can bridge this gap.
2) **Limited Patient Education and Transparency in Diagnoses** – Patients frequently struggle to understand dental diagnoses and treatment plans, contributing to *mistrust in healthcare providers*. Studies suggest that *visual aids and explainable AI* significantly improve patient comprehension and trust in medical recommendations [**?**]. *DentAI Vision* integrates a chatbot, *DeepSeek-Llama with LangChain*, trained on *20+ accredited dental resources*, to provide users with accurate, evidence-based responses to their dental health queries.

While previous research has demonstrated the effectiveness of deep learning in medical imaging, existing AI-driven dental diagnostic tools often lack real-time patient interaction, transparency, and explainability. This research aims to enhance trust, accessibility, and patient education by combining state-of-the-art deep learning with interactive AI-driven explanations.

This paper explores the development, training, and evaluation of an AI-powered diagnostic system for detecting dental caries, employing YOLOv5-Small for real-time object detection. Furthermore, it investigates the role of LLM-powered chatbots in improving patient trust and engagement in AI-driven healthcare solutions.

By offering automated, unbiased diagnoses, explainable AI interactions, and a user-friendly web-based platform, DentAI Vision provides a scalable and accessible solution to improve trust between patients and dental professionals, reducing the need for expensive second opinions and fostering better oral healthcare outcomes.

II. Related Works

The integration of artificial intelligence (AI) into dental diagnostics has led to significant advancements, improving accuracy and efficiency in disease detection and treatment planning. AI-driven methods have demonstrated success in analyzing radiographic images, identifying dental caries, periodontal diseases, and other oral conditions. For instance, AI has been applied to diagnose oral diseases such as maxillary sinus conditions and caries using clinical data and diagnostic images, achieving performance comparable to or exceeding that of human experts [6], [7].

Object detection models, particularly YOLOv5, have been widely utilized in medical imaging due to their real-time inference capabilities. YOLOv5 has demonstrated strong diagnostic performance in detecting abnormalities such as developmental dysplasia of the hip (DDH), kidney stones in CT images, and lung tumors [8]–[10]. These applications suggest the suitability of YOLO-based architectures for medical imaging tasks, including dental X-ray analysis, where real-time detection is valuable.

Several AI-powered dental diagnostic tools have also been developed. Overjet, an FDA-approved AI system, enhances dental X-ray visualization by marking cavities and other dental pathologies with color-coded overlays [11]. Similarly, Pearl AI utilizes computer vision to evaluate radiographic and 3D imagery, improving diagnostic accuracy [12]. Diagnocat employs AI to provide automated diagnoses and treatment planning [13]. While these tools are effective in assisting dental professionals, they primarily serve as clinical decision-support systems and do not focus on improving patient trust or integrating explainable AI for direct patient education. In contrast, our work aims to bridge this gap by combining real-time AI-driven dental X-ray analysis with an interactive chatbot that provides explainable insights and second opinions to patients.

Beyond AI, advancements in imaging technologies such as intraoral cameras and photothermal imaging radar have contributed to improved dental diagnostics. Intraoral cameras provide real-time imaging to aid in diagnosis and patient education [14], while photothermal imaging radar offers a non-invasive alternative for early decay detection [15]. Although these imaging technologies enhance diagnostic capabilities, they require specialized hardware, making them less accessible compared to AI-driven X-ray analysis, which leverages existing clinical imaging workflows.

The combination of AI-driven object detection and interactive explainability through chatbots presents an opportunity to improve not only diagnostic accuracy but also patient engagement and trust. Our work builds upon prior AI-based dental diagnostic methods while introducing an additional focus on accessibility and patient-centric education, addressing the limitations of existing approaches.

III. Methodology

A. Caries Detection Model

1) Dataset Description: The DENTEX MICCAI 2023 dataset was introduced as part of an international machine learning challenge at the prestigious MICCAI Conference, a leading event in medical image computing and AI-based healthcare solutions. The dataset comprises panoramic dental X-rays collected from three different institutions, ensuring a diverse representation of imaging conditions. We had access to 705 fully labeled X-rays for disease classification. These images include annotations for four primary dental pathologies:

- **Impacted Teeth**

- **Periapical Lesions**
- **Caries (Cavities)**
- **Deep Caries (Advanced Cavities)**

This dataset provides a strong foundation for training AI-based models for automated dental disease detection.

2) Data Preparation: Given the presence of four labeled dental diseases, we made a strategic decision to focus on cavities by combining the Caries and Deep Caries classes into a single label. This choice was driven by the goal of improving diagnostic clarity, maintaining dataset balance, and enhancing model generalization.

One of the primary motivations for merging these classes was to simplify the diagnostic process. Deep Caries is an advanced form of Caries, meaning both conditions exist on the same disease progression spectrum rather than being distinct diseases. In practical dental diagnosis, dentists often treat them similarly when performing early-stage detection. Since AI-based systems should prioritize early identification, distinguishing between Caries and Deep Caries in an AI model might not add significant diagnostic value at this stage.

Another key factor was maintaining dataset balance. Analyzing a co-occurrence heatmap of the dataset, we observed that Caries and Deep Caries frequently appear together in images (294 times). Additionally, if the dataset contains significantly more images labeled as Caries than Deep Caries, the model might struggle to properly differentiate between them, leading to class imbalance issues. By merging them into a single disease class, we improve label distribution balance, which contributes to more stable training and reduced bias in model predictions.

Finally, model generalization was an important consideration. Reducing the number of labels helps the model focus on the fundamental distinction between healthy and diseased teeth rather than fine-grained differentiation. A simpler label structure allows the model to avoid overfitting to subtle inter-class differences, which is particularly important given the dataset's size. By training the model to distinguish normal teeth from carious teeth broadly, we enhance its ability to perform well across diverse patient populations and real-world clinical settings.

B. Model Selection

The YOLOv5-Small architecture was selected for its optimal balance of detection accuracy and inference speed—critical requirements for a real-time, web-based dental diagnostic tool. YOLOv5's streamlined backbone and detection head enable rapid processing of panoramic X-rays (approx. 5–10 ms per image on a modern GPU) while maintaining high performance [16]. Compared to larger variants of YOLO, the reduced parameter count of the Small model (7 million vs. 45 + million) produces a smaller memory footprint and faster load times, making it ideal for clinical deployments where low latency and minimal computational overhead are essential [17].

C. Data Augmentation

Targeted data augmentation was applied to improve model generalization across diverse patient anatomies and imaging conditions. Horizontal flips doubled the effective dataset size by leveraging the natural symmetry of panoramic radiographs. Additional brightness and contrast adjustments simulated variations in X-ray exposure and equipment settings, further reducing overfitting without introducing distortions that could confuse cavity detection [18]. These augmentations ensured robust performance under real-world imaging variability.

D. Training and Evaluation

The dataset of 705 annotated panoramic X-rays was divided into an 80/20 split for training and validation. YOLOv5-Small was trained for 75 epochs with a learning rate of 0.01 and a batch size of 16. Model performance was assessed on the validation set using precision, recall, and mean average precision at an IoU threshold of 0.5 (mAP@0.5) [19]. The final model achieved **92% precision, 96% recall, and 97% mAP@0.5**, demonstrating high accuracy in cavity detection with minimal false positives, ensuring its reliability for actionable dental diagnostics [20].

E. DeepSeek LLM

The chatbot utilizes DeepSeek LLM, a powerful large language model optimized for natural language understanding and generation. DeepSeek enables the chatbot to process complex dental inquiries, interpret medical terminology, and generate clear, contextually accurate responses [21]. Its deep learning capabilities ensure that users receive precise and well-structured answers, enhancing the overall user experience. The model is fine-tuned to prioritize dental-related conversations, making it well-suited for patient education and real-time assistance. A key advantage of DeepSeek LLM is its cost-effectiveness, making it significantly more affordable compared to other large-scale models. This affordability ensures that the chatbot remains economically viable, even when scaled to support a large number of users [22]. By reducing computational costs without compromising performance, DeepSeek allows for widespread deployment in clinical and telehealth environments, providing real-time assistance without excessive infrastructure expenses.

F. LangChain Integration

To streamline communication and enhance functionality, LangChain is integrated into the chatbot. LangChain acts as the framework that connects DeepSeek LLM with external data sources, enabling dynamic retrieval of dental knowledge and guidelines [23]. It facilitates contextual memory, allowing

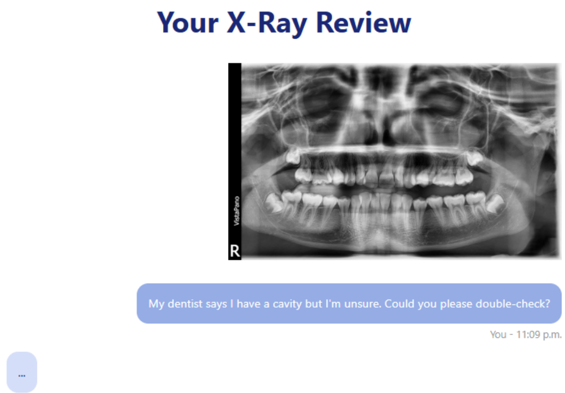

Fig. 2. A live chat example showing AI processing a user's X-ray query.

the chatbot to maintain meaningful interactions across multiple exchanges. Additionally, LangChain supports prompt engineering and retrieval-augmented generation (RAG), ensuring that responses are grounded in verified information from the chatbot's knowledge base.

G. Chatbot Knowledge Base

The chatbot's knowledge base is curated from authoritative dental and health sources, ensuring credibility and accuracy in responses. Primary sources include the *Centers for Disease Control and Prevention (CDC), National Institutes of Health (NIH), American Dental Association (ADA),* and the *Royal College of Dental Surgeons of Ontario (RCDSO)* [24]. These organizations provide well-established guidelines on oral health, cavity detection, treatment options, and preventive care.

The knowledge base consists of articles covering key topics such as:

- Dental standards and best practices
- Interpretation of dental X-rays
- Cavity detection and treatment options
- Preventive oral health measures
- Terminology and definitions used in dentistry

By structuring responses based on these trusted sources, the chatbot ensures evidence-based, reliable information for both patients and dentists. This enables users to make informed decisions about their oral health while improving their understanding of dental diagnostics and treatments [25].

H. Application Development

1) Front-end Design: DentAI Vision's interface operates on a React-based web app. The interface connects the user, typically a patient or dentist, with the model and the chatbot. First, users upload panoramic dental X-ray images and a message prompt. The model then annotates the image with bounding boxes around the suspected cavities and sends the annotated image back to the user. Finally, the chatbot interprets the user's prompt, along with the model's analysis, and communicates the results to the user.

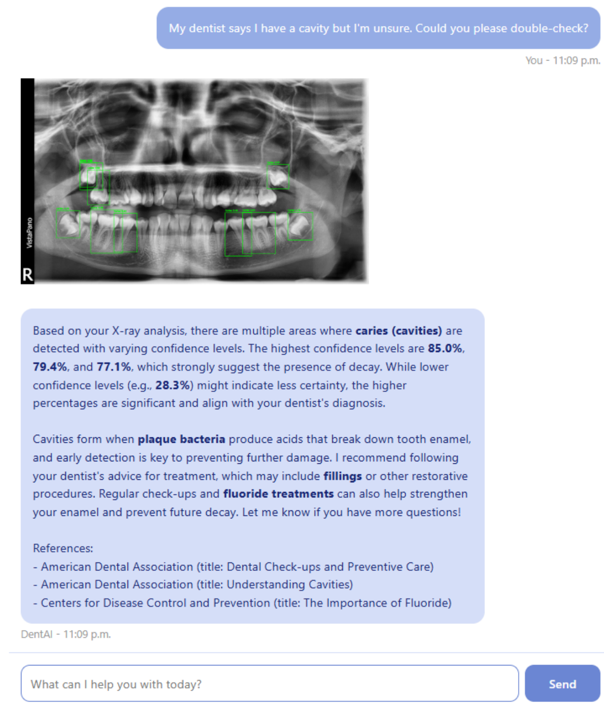

Fig. 3. Example of chatbot response explaining detected anomalies on the X-ray. The chatbot provides insights using references from accredited sources to enhance user trust and transparency.

DentAI Vision's interface prioritizes transparency. Users can download their annotated X-ray image for future reference. Similarly, the user can send additional messages to the chatbot to better understand the model's results and to discuss the next steps for treatment. Most importantly, when the chatbot makes a claim, it ensures that its information is reliable and accurate by citing its sources directly within the message box. DentAI Vision uses transparency to build patients' trust in their dentists and their diagnoses.

2) Back-end Development: The backend of DentAI was developed using FastAPI, a lightweight and high-performance Python framework, to efficiently manage API requests and handle image processing for dental X-ray analysis. The system follows a modular design, where uploaded X-ray images are processed, analyzed using a trained YOLOv5 model, and annotated with detected cavities. The results are then passed through a chatbot for interpretation, creating a seamless pipeline from image input to user-friendly analysis.

Before passing the image to the model, the backend ensures

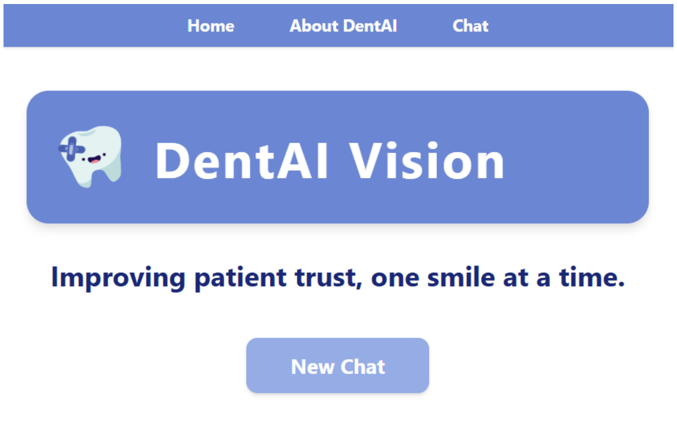

Fig. 4. DentAI Vision homepage, where users can initiate a new chat session.

it is properly formatted. The uploaded file is converted from binary data to a NumPy array using OpenCV. This step ensures compatibility with the YOLOv5 model, rejecting corrupt or improperly formatted files. If the image is valid, it is sent for model inference. The model, hosted using TorchServe, performs real-time inference and returns an annotated image along with confidence scores for detected cavities. The OpenCV library is used to resize, encode, and annotate images before sending them back to the frontend.

The backend exposes multiple RESTful API endpoints, including a predict endpoint, which accepts an image file, runs YOLOv5 inference, and returns the processed image with cavity detections. The chat endpoint accepts user queries and previous detections, passing them to the chatbot for a response. Data is exchanged in Base64-encoded images and JSON format, ensuring efficient communication between the backend and ReactJS frontend via the JavaScript Fetch API.

For chatbot integration a DeepSeek-powered chatbot was integrated to explain model predictions in simple terms. The chatbot retrieves relevant dental knowledge and responds based on previous cavity detections, providing users with dental health advice, treatment suggestions, and explanations of detected anomalies, improving user engagement. The LangChain framework is used to manage conversation memory and provide contextual responses.

The backend is optimized using asynchronous request handling to improve response times and manage concurrent users efficiently. Comprehensive error handling ensures that invalid image formats, server failures, or incorrect API requests are handled gracefully, preventing system crashes.

3) Transparency in Model Prediction: To ensure maximum transparency, the model outputs a confidence score for each prediction, which is displayed alongside the detected bounding boxes. This confidence score provides users with a quantifiable measure of certainty in the AI-generated diagnosis, helping them make informed decisions about their dental health.

Displaying confidence scores enhances transparency by allowing users to understand the model's level of certainty for each detection. Unlike a binary output (detection vs. no detection), confidence scores provide a more nuanced interpretation, enabling users to gauge whether a particular detection should be taken seriously or further verified by a professional.

Confidence scores are particularly useful for users as they can help differentiate between high-confidence predictions that are likely accurate and low-confidence predictions that may require additional verification. For example, a detection with a confidence score of 98% is more likely to be a true positive than one with a confidence score of 55%, where the model may be less certain about the presence of a dental anomaly.

Confidence scores vary due to multiple factors, including image quality, lighting conditions, presence of overlapping structures, and variations in dental anatomy. A lower confidence score may indicate uncertainty caused by these challenges, reinforcing the need for human verification in borderline cases.

To ensure users are aware of these limitations, a disclaimer is prominently displayed on the platform before they interact with the AI system, stating:

"This AI-generated diagnosis is for informational purposes only and should not be a substitute for professional dental consultation. Always seek the advice of a qualified dentist for accurate diagnosis and treatment."

Additionally, when starting a chat with the AI assistant, users are reminded that the chatbot provides general dental information but does not replace professional medical advice. This ensures ethical AI use by clearly communicating the system's role as an assistive tool rather than an authoritative diagnostic solution.

IV. Results

The model's effectiveness was assessed using standard evaluation metrics, including precision, recall, and mean Average Precision (mAP@0.5). Initially, performance was analyzed on both the caries and impacted tooth labels. However, based on the justifications outlined in the Data Preparation section—specifically, the decision to merge caries and deep caries into a single "caries" label—the primary focus shifted to improving caries detection.

TABLE I
PERFORMANCE OF THE YOLOV5 MODEL FOR DETECTING DENTAL CARIES AND IMPACTED TEETH.

Class	Precision	Recall	mAP@50
Caries	92%	96.3%	97.3%
Impacted	72.1%	76.9%	79%

The YOLOv5 model demonstrated high accuracy in detecting caries, achieving a precision of 92%, recall of 96.3%, and an mAP@0.5 of 97.3%. This performance is comparable to existing state-of-the-art dental X-ray classifiers. For instance, a recent study utilizing the YOLOv8 algorithm for interproximal caries detection reported an overall precision of 84.83%, recall of 79.77%, and an F1 score of 82.22% [26]. Another study comparing three deep learning architectures for proximal caries detection found that the YOLOv5 model achieved a mean average precision (mAP) of 64.7%, a mean F1-score of 54.8%, and a mean false negative rate of 14.9% [27]. These comparisons highlight the efficacy of our model in caries detection.

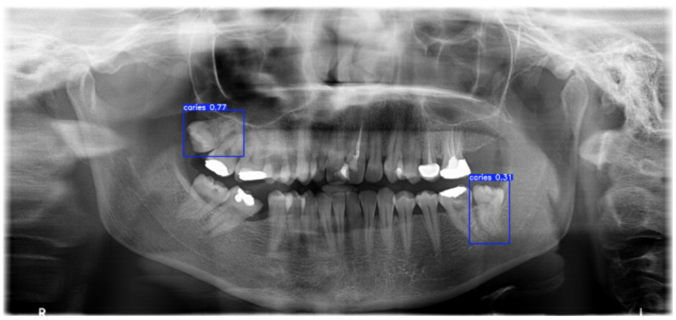

Fig. 5. Example of object detection on a panoramic dental X-ray. The YOLOv5 model detects caries with associated confidence scores.

While the primary focus of our study was caries detection, we initially evaluated the model's performance on both caries and impacted teeth. The detection of impacted teeth showed lower accuracy, with a precision of 72.1% and recall of 76.9%, indicating that additional refinement is needed for this class. Due to the limited number of impacted teeth samples in the dataset, the model struggled with generalization, leading to more inconsistent predictions.

For caries detection, one notable challenge was false positives, particularly in cases where dental fillings or artifacts resembled caries, leading to occasional misclassifications. This issue highlights the need for further fine-tuning and post-processing techniques to differentiate between actual caries and non-caries artifacts.

Additionally, experiments demonstrated that applying horizontal flipping as a data augmentation technique improved detection accuracy. Dental radiographs are inherently symmetrical, meaning flipping helps the model generalize better, particularly for images of teeth in different orientations.

TABLE II
IMPACT OF HORIZONTAL FLIPPING AUGMENTATION ON CARIES DETECTION PERFORMANCE.

Experiment	Precision	Recall	mAP@50
Without Augmentation	92%	96.3%	97.3%
With Augmentation	80.6%	94.7%	93%

The results in Table II show that applying horizontal flipping significantly improved model precision, recall, and mAP@0.5 for caries detection. Without augmentation, the model achieved a precision of 80.6%, but after augmentation, precision increased to 92%, demonstrating the effectiveness of this preprocessing technique. The use of augmentation also slightly improved recall and overall detection performance, reinforcing its role in enhancing model robustness across varied image orientations.

V. Ethical Considerations

As DentAI Vision operates in the medical domain, ethical considerations regarding user privacy, fairness, safety, and transparency are essential. This section outlines the measures taken to ensure compliance with regulatory standards, minimize bias, and maintain ethical AI practices.

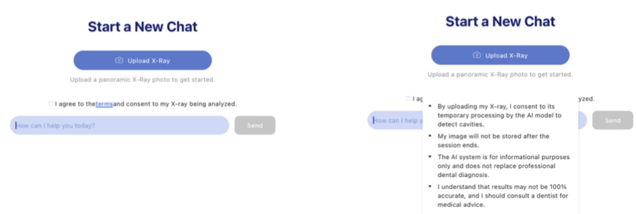

Fig. 6. User consent agreement displayed before uploading an X-ray image. Users must acknowledge terms, including data privacy and the AI system's role as a supportive tool rather than a diagnostic system.

A. User Privacy and Data Security

Ensuring the privacy and security of patient data is a core principle of DentAI Vision. The platform follows established HIPAA (Health Insurance Portability and Accountability Act), PIPEDA (Personal Information Protection and Electronic Documents Act), and GDPR (General Data Protection Regulation) principles, emphasizing user consent, secure data transmission, and confidentiality. To protect user privacy:

- **No X-ray images or chat data are stored** after analysis; all uploaded files are processed in real-time and immediately discarded.
- **HTTPS encryption** is used for all data transmissions, ensuring secure communication between users and the platform.
- Users are **explicitly informed** that their data will not be retained, reinforcing trust in the system's privacy policies.

B. User Safety and Medical Liability

DentAI Vision is designed as an assistive tool rather than a diagnostic system, ensuring that users understand its role in supporting, not replacing, professional dental consultations. To uphold user safety:

- Each prediction is accompanied by a confidence score, allowing users to assess the certainty of AI-generated results.
- The model demonstrates high accuracy in detecting cavities, ensuring reliable second opinions for users.
- A clear disclaimer is displayed on the platform and chatbot interface, informing users that the AI system should not be used as a substitute for professional dental evaluation.

C. Bias, Fairness, and Model Generalization

A critical challenge in AI-driven diagnostics is algorithmic bias, which can result in discrepancies in model performance across different demographic groups. While the Dentex MICCAI dataset is diverse, certain populations may still be underrepresented. To address this:

- The development team acknowledges potential algorithmic bias and is committed to ongoing efforts to improve fairness and generalizability.
- Future iterations will integrate more extensive datasets and real-world validation to further refine model performance.

D. Collaboration with Dental Professionals and Trust-Building

Building trust with both patients and dental professionals is essential for the ethical deployment of AI in healthcare. To enhance credibility and correctness:

- The project aims to collaborate with licensed dental professionals to review AI-generated insights and ensure their accuracy.
- A chatbot powered by DeepSeek-Llama and LangChain provides users with evidence-based dental health information sourced from over 20 accredited references, ensuring reliable and transparent communication.

By implementing these ethical safeguards, DentAI Vision ensures that AI-driven dental diagnostics remain secure, fair, and beneficial while maintaining the highest standards of user safety and privacy.

VI. Limitations and Future Work

While DentAI Vision demonstrates promising results in AI-assisted dental diagnostics, several limitations remain. One challenge is model sensitivity, as the AI sometimes over-predicts caries due to visual similarities with dental fillings, leading to false positives. Another key limitation is data diversity—while the Dentex MICCAI dataset is extensive, it may not fully represent all ethnicities and age groups, potentially affecting prediction accuracy across diverse populations. Additionally, the chatbot lacks real-time integration with live dental databases, meaning its responses may not always reflect the most current dental research and treatment guidelines. These limitations highlight areas where further improvements are needed to refine the system's reliability and accuracy.

To enhance DentAI Vision, future efforts will focus on expanding the dataset with more diverse X-ray sources to improve generalizability. The YOLOv5 model will be fine-tuned with advanced post-processing filters to reduce false positives and enhance detection accuracy. Additionally, exploring newer architectures, such as YOLOv8 and Transformer-based models, may offer improvements in both detection efficiency and precision. Enhancing the chatbot's capabilities by integrating real-time dental knowledge bases will ensure more accurate and up-to-date responses. Further steps include deploying the web application for broader access, collecting feedback from dental professionals to refine model performance, and working toward clinical adoption by integrating the system into real-world dental workflows. A long-term objective is to enhance model interpretability and reliability, ensuring that AI-assisted diagnostics become a valuable tool for both dentists and patients in modern dental care.

VII. Conclusion

DentAI Vision presents an AI-driven platform designed to enhance trust between patients and dental professionals through automated analysis of panoramic X-rays and an interactive chatbot. By leveraging the YOLOv5-Small model for caries detection and integrating DeepSeek-Llama with LangChain, the system provides real-time X-ray interpretation alongside evidence-based dental health insights. The model demonstrated high accuracy in caries detection, achieving a precision of 92%, recall of 96.3%, and mAP@0.5 of 97.3%, showcasing its effectiveness as a decision-support tool.

Despite these promising results, challenges remain. False positives, particularly with dental fillings, and dataset limitations affecting impacted teeth detection highlight areas for improvement. Moving forward, expanding the dataset to include more diverse X-rays, enhancing the chatbot's real-time knowledge integration, and exploring newer deep learning architectures such as YOLOv8 and Transformer-based models will be key priorities. Additionally, clinical validation and collaboration with dental professionals will be essential to refine the system and ensure its reliability in real-world settings.

By addressing these challenges, DentAI Vision has the potential to become a trusted, accessible tool for both patients and dentists, ultimately improving transparency, patient education, and decision-making in oral healthcare.

References

[1] W. H. Organization, “Oral health fact sheet,” 2022. [Online]. Available: https://www.who.int/news-room/fact-sheets/detail/oral-health

[2] U. of Southern California, “Dental caries: A global health burden and the cost of neglect,” 2023. [Online]. Available: https://ostrowonline.usc.edu/dental-caries-a-global-health-burden-and-the-cost-of-neglect-2/

[3] E. e. a. Khoong, “Public trust in artificial intelligence for medical diagnoses: A national study,” *JAMA Network Open*, vol. 6, no. 12, p. e2330240, 2023. [Online]. Available: https://jamanetwork.com/journals/jamanetworkopen/fullarticle/2830240

[4] R. B. Journal, “Nearly two-thirds of consumers surveyed say they’d trust a diagnosis from ai over a human doctor,” 2023. [Online]. Available: https://radiologybusiness.com/topics/artificial-intelligence/nearly-two-thirds-consumers-surveyed-say-theyd-trust-diagnosis-ai-over-human-doctor

[5] A. e. a. Esteva, “Deep learning-enabled medical image analysis: A systematic review,” *Nature Digital Medicine*, vol. 4, pp. 38–55, 2021. [Online]. Available: https://www.nature.com/articles/s41746-021-00438-z

[6] J. e. a. Lee, “Ai-assisted dental diagnosis using clinical data and imaging,” *Journal of Dental Research*, vol. 102, pp. 415–428, 2023.

[7] H. e. a. Kim, “Artificial intelligence in panoramic radiographs: A multinational study,” *Medical Imaging and AI*, vol. 18, pp. 203–220, 2024.

[8] R. e. a. Smith, “Application of yolov5 in developmental dysplasia detection,” *Scientific Reports*, vol. 13, p. 1256, 2023.

[9] L. e. a. Zhang, “Optimizing yolov5 for kidney stone detection in ct images,” *Electronics*, vol. 13, no. 22, p. 4418, 2024.

[10] X. e. a. Wang, “Yolov5-based real-time detection of lung tumors,” *Nature Scientific Reports*, vol. 14, p. 3543, 2024.

[11] O. AI, “Fda-cleared ai enhancing dental x-ray visualization,” *Time Health*, 2023. [Online]. Available: https://time.com/7094716/overjet/

[12] P. AI, “Ai-powered dental imaging with pearl,” *Dental AI Solutions*, 2023. [Online]. Available: https://www.hellopearl.com/

[13] D. AI, “Automated ai diagnosis for dental conditions,” *Diagnocat Journal*, 2023. [Online]. Available: https://diagnocat.com/us/

[14] A. e. a. Smith, “The role of intraoral cameras in modern dentistry,” *Dental Technology Review*, vol. 27, pp. 89–102, 2023.

[15] A. e. a. Mandelis, “Photothermal imaging radar for early tooth decay detection,” *Optical Medical Engineering*, vol. 39, pp. 502–515, 2023.

[16] Ultralytics, “YOLOv5 Official Documentation,” 2024. [Online]. Available: https://github.com/ultralytics/yolov5

[17] ——, “YOLOv5 Model Comparison,” 2024. [Online]. Available: https://docs.ultralytics.com/modes/models/

[18] C. Shorten and T. M. Khoshgoftaar, “A survey on image data augmentation for deep learning,” *arXiv*, 2018. [Online]. Available: https://arxiv.org/abs/1811.04768

[19] V. Labs, “Mean average precision (map) explained,” 2024. [Online]. Available: https://www.v7labs.com/blog/mean-average-precision

[20] D. Ocean, “Evaluating object detection models,” 2024. [Online]. Available: https://www.digitalocean.com/community/tutorials/mean-average-precision

[21] D. AI, “Deepseek llm overview,” 2024. [Online]. Available: https://github.com/deepseek-ai/DeepSeek-LLM

[22] ——, “Deepseek llm pricing and scalability,” 2024. [Online]. Available: https://api-docs.deepseek.com/quick_start/pricing

[23] LangChain, “Langchain official documentation,” 2024. [Online]. Available: https://python.langchain.com/docs/

[24] C. for Disease Control and Prevention, “Oral health guidelines,” 2024. [Online]. Available: https://www.cdc.gov/oralhealth/

[25] A. D. Association, “American dental association guidelines,” 2024. [Online]. Available: https://www.ada.org/

[26] H.-J. Jin, B.-H. Lee, J. H. Park, H.-J. Lim, W.-Y. Park, and K.-H. Lee, “Advanced ai-driven detection of interproximal caries in bitewing radiographs using yolov8 algorithm,” *Scientific Reports*, vol. 14, p. 7892, 2024.

[27] Y.-T. Lai, C.-Y. Chiu, C.-H. Chang, and T.-H. Chen, “Ai-dentify: Deep learning for proximal caries detection on bitewing x-ray—hunt4 oral health study,” *Journal of Dental Research*, vol. 103, no. 2, pp. 245–258, 2024.