DigiHealth-Asia

D2.6 - Pilot use case: Remote consultation of patients – Early childhood caries detection and prevention

Work package	WP2
Task	Task 2.6
Due date	
Submission date	3rd December 2024
Deliverable lead	ULL and NUM
Version	2
Authors	ULL team (Giacomo Kahn, Aicha Sekhari), NUM team (Mend-Amar Majig, Enkhtuvshin Bat-Ochir) and MNUMS team (Ariuntuul Garidkhuu, Bat-Erdene Myagmar, Delgerbaigali Munkhjargal)
Reviewers	Adnan Shahid (UGent)

Table of Contents

List of Figures

Executive Summary

This pilot case aimed to explore how artificial intelligence (AI) could be used to remotely diagnose early childhood dental caries. Tooth decay is a significant health concern in Mongolia, where 2023 statistics from the Ministry of Health revealed an average prevalence of 70% among children. This high rate of dental issues affects both rural and urban areas, with contributing factors such as poor oral hygiene education, high consumption of sugary foods, and socio-economic challenges. The pilot project sought to leverage AI technology to address these issues, especially in rural areas where access to dental care is limited.

Led by a joint team from the University of Lumiere Lyon 2, Mongolian National University of Medical Science and the National University of Mongolia, the project aimed to develop an AI-based system capable of diagnosing dental caries in children. Its objectives included building a comprehensive database of dental images, designing a disease recognition algorithm, and integrating this system into a cloud-based platform for remote consultation. Starting in March 2023, the project collected data from 650 children across kindergartens and preschools in Ulaanbaatar city as well as in rural areas, using intraoral cameras to capture high-resolution images of their teeth.

The data collection process involved several steps. First, informed consent was obtained from parents, followed by a detailed questionnaire on the children's oral health. The medical examinations were then carried out, with dentists using intraoral cameras to take images of each child's teeth. These images were processed using the YOLOv8 machine learning model, which was trained to detect dental conditions such as caries, pulpitis, and fluorosis. Despite some data imbalance that initially affected the model's accuracy, particularly for less common conditions, the AI system performed well in identifying and categorizing dental issues. Dentists also validated the AI’s predictions to ensure its reliability.

One of the pilot's key innovations was the development of a web-based platform (@www.digihealth-mongolia.com), which allowed dentists and healthcare professionals to upload intraoral images for AI analysis. This platform created a feedback loop where the AI system’s diagnoses could be reviewed by experts, and any discrepancies were corrected. This iterative process enabled continuous improvement of the model, helping it become more accurate over time.

The pilot case demonstrated the potential for AI to transform healthcare in Mongolia, especially in underserved rural areas where dental services are scarce. By incorporating AI technology, the healthcare system could ease the burden on dentists, provide quicker and more accurate diagnoses, and ultimately improve patient outcomes. The successful completion of this pilot established a foundation for further advancements in AI-driven healthcare, with plans to expand the system’s diagnostic capabilities to other areas, including mucosal diseases and X-ray interpretation.

In conclusion, this pilot case showcased the transformative power of AI in addressing healthcare disparities in Mongolia. As the AI system continues to evolve and more data is collected, its accuracy will improve, offering scalable solutions for healthcare delivery. The project also highlighted the importance of interdisciplinary collaboration between medical professionals and AI researchers. With

continued development, AI-driven diagnostic systems could significantly enhance healthcare access and quality, particularly for children in rural and underserved areas of Mongolia.

I. Introduction

According to the 2023 statistics of the Ministry of Health of Mongolia, the prevalence of tooth decay in children is, on average, 70% [1]. Recent studies indicate this high prevalence in urban and rural settings. Several factors contribute to this problem, including inadequate parental education on oral health and a lack of emphasis on oral hygiene in preschool, primary, and secondary education curriculums. Additionally, the increase in the consumption of sugary foods and drinks exacerbates this issue, alongside broader socio-economic challenges, population migration, and limited access to healthcare infrastructure.

Mongolia's healthcare system faces critical challenges, especially in rural areas where healthcare services are limited. The sparsity of medical centers, including dental services, in rural Mongolia makes it difficult for residents to access even basic healthcare. Many remote areas lack specialist doctors and necessary healthcare personnel, creating unequal healthcare access between rural and urban populations. In rural areas, there is a shortage of dental treatment facilities and specialist personnel, while in more central regions, the lack of dental assistants puts additional strain on dentists. This disparity in healthcare access emphasizes the need for innovative solutions, including using artificial intelligence (AI) in healthcare.

Oral health is a key component of overall well-being, and ensuring proper dental hygiene from an early age can prevent future complications and improve quality of life. As such, today's oral healthcare systems increasingly require interdisciplinary approaches that combine medical expertise with technological innovations, such as AI. The growing use of AI in healthcare is evident globally, and it is beginning to be incorporated into educational curricula in technical fields and medicine. In Mongolia, some private universities and secondary education institutions are introducing AI into medical education, recognizing its potential to solve healthcare challenges in collaboration with other professional fields.

This study explores how AI can be used to remotely diagnose early-stage dental caries in children, educate parents and caregivers about oral hygiene, and provide timely treatment advice. This research is part of the European Union Erasmus+ DigiHealth-Asia project, where a joint team from the Mongolian National University of Medical Science (MNUMS) and the National University of Mongolia (NUM) is collaborating with AI researchers from the University of Lumiere Lyon 2 (ULL) to develop AI-based solutions for addressing dental health issues in children. By leveraging AI, this project seeks to improve early detection and treatment of dental caries, especially in underserved rural areas.

II. Objectives of the pilot case

Purpose:

The pilot case aims to create a system for evaluating and providing remote diagnosis of oral diseases in preschool children with the help of artificial intelligence (AI).

Objectives:

1. Create a database of dental caries and healthy images using an intraoral camera.
2. Design a dental disease recognition algorithm by building a database from preschool children's mouth photos.
3. Implement labeling and machine learning techniques to train the AI model.
4. The AI's diagnostic performance is validated by having professional dentists randomly review pre-diagnosed cases.
5. Evaluate the effectiveness of the AI system when integrated into dental practice using a cloud-based platform.

2.1. System architecture

The general workflow of the pilot case is illustrated in Figure 1, highlighting how technology and healthcare can work hand in hand. It all starts when a medical doctor or nurse takes images of the patient's mouth using an intraoral camera. Once the image is taken, it can be uploaded to the online server at www.digihealth-mongolia.com using either USB or Bluetooth.

This online platform houses a pre-trained machine-learning model the ULL and the NUM developed. Before we made this service available, the model was carefully trained with sufficient data to ensure it could provide accurate and useful analysis. After uploading the image, the model generates a preliminary diagnosis of the patient's oral health.

Upon completion of the analysis, the findings are shared with both the doctor and the patient. In addition to the diagnosis, the system provides helpful suggestions and recommendations, enabling the healthcare provider and patient to make informed decisions about next steps. By streamlining the diagnostic process, this workflow enhances efficiency and strengthens the partnership between patients and healthcare providers. Furthermore, the system diagram is adaptable not only to intraoral camera images but also to other types of labeled images, as demonstrated in research studies [3, 4, 5].

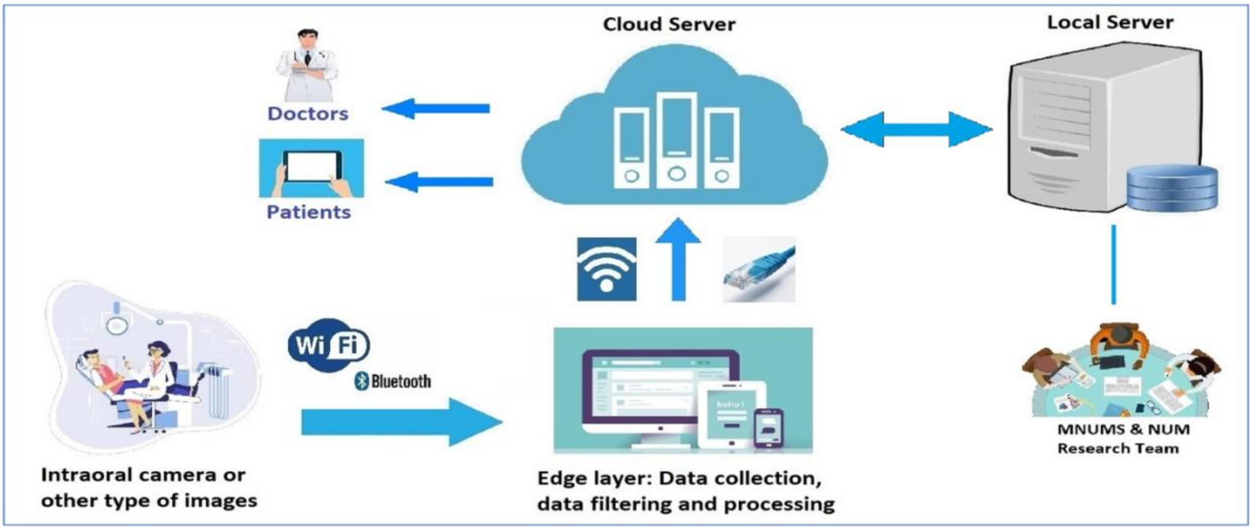

Figure 1. System Architecture

III. Data collection and preparation

Data collection for this study commenced in March 2023, following the formal approval from the Mongolian National University of Medical Sciences Ethics Control Committee. Over the course of 7 days, a comprehensive examination was conducted on 50 children from Khotont Soum in Arkhangai Province and Orkhon Soum in Selenge Province (Figure 3), as well as in the capital city of Ulaanbaatar, where data was collected from kindergartens and preschools. The data collection step was continued in the city of Ulaanbaatar (Figure 2), and by December 2023, the research team gathered data from a total of 650 children, contributing to a robust dataset.

The data collection team comprised a multidisciplinary group of professionals, including 2 experienced oral surgeons, 7 dedicated doctors from the Department of Oral and Maxillofacial Surgery, and 3 specialists in oral and maxillofacial technology. In addition, the team included 15 skilled dental hygienists, 12 other medical practitioners, and 8 nurses affiliated with Megadent Dental.

Data collection also involved the distribution of a detailed survey questionnaire to parents of the 650 participating children, aimed at gathering comprehensive information regarding dental health and hygiene practices. Moreover, approximately 10 intra-oral camera images were captured for each child, culminating in around 6,000 high-resolution images. This extensive visual documentation complements the survey data, providing a thorough basis for analysis and insights into the dental health status of children in these provinces and Ulaanbaatar.

3.1. Data collection procedure

The data collection process involved the following three key stages.

Stage 1. Informed consent forms were distributed to the parents mid-week, and a review and information session were held to explain the study. Consent forms, officially signed by the parents or guardians, were collected by the end of the week.

Stage 2. Survey Administration Parents and guardians were asked to complete a detailed questionnaire over the weekend, which was collected by group teachers on Monday. In rural areas, the informed consent and questionnaire process was streamlined, with consent obtained and information completed on the first day of the inspection to save time.

Stage 3. Medical Examination and Data Collection Medical examinations were conducted based on the capacity of the kindergarten and the available time. Examination rooms were typically set up in advance, either in designated doctor's offices or institutional halls. The process involved close collaboration with the group of teacher and assistant teachers.

To begin, the children received oral hygiene lessons to help them understand the importance of dental care, and special attention was given to creating a positive, team-oriented environment for the children. Following this, the teachers and the children demonstrated the proper use of toothpaste and toothbrushes. Finally, the children were provided professional advice on brushing techniques, focusing on brushing their teeth both in the morning and after meals (guided by the Oral Hygiene Team).

Four doctors conducted oral examinations on the children, carefully documenting their findings on individual medical records. At the same time, another team of four doctors used an intraoral camera to capture detailed images of the children's teeth, ensuring a comprehensive record of each child's dental health.

Figure 2. 164th Kindergarten of Ulaanbaatar City Railway

During the rural inspection, the collection team collaborated with the Megadent Hospital as a medical supporter. In examining the residents of remote regions, they provided crucial assistance in oral treatment, oral surgery, and pediatric dental care while also addressing urgent cases and taking necessary emergency measures.

Figure 3. Oral examination of citizens of Orkhon soum, Selenge province

The examinations focused primarily on identifying hard dental tissue diseases in children through oral assessments. To minimize the children's fear of doctors visiting their kindergarten, the medical team avoided wearing traditional doctor's attire as much as possible. The World Health Organization's dental disease classification system was used to document diagnoses in the medical records, and any changes in symptoms were also carefully noted. For instance, black caries were classified as either hard or soft based on their characteristics.

The oral mucosa, labial frenulum, and tongue frenulum were also examined and recorded. Using the SNAP camera (Figure 4) produced by the Korean Osstem company, oral images were captured, and a data archive was created with each child assigned a unique numeric code in the software (e.g., K164P002, where K stands for the kindergarten name and P represents the child's number). Before taking the photos, the oral cavity was dried with a sterile bandage. Disposable gloves and protective covers were used to shield both diseased and healthy teeth surfaces. Photos were taken at the highest resolution to ensure clarity, with an average of 10 images captured per child.

3.2. Data labeling

Following data collection, the process of cleaning and pre-processing the data was carried out in a structured, multi-stage approach to ensure accuracy and usability. The first stage involved extracting and organizing image data using the SNAP intraoral camera, developed by the Korean company Osstem, to collect high-quality images of the children's teeth (Figure 4). These images, captured in JPG format with a resolution of 640x480 pixels and file sizes ranging from 55 to 64 KB, were systematically stored in a structured library. Each child's data was organized using a unique code representing the nursery and the child number, ensuring easy retrieval and traceability.

Figure 4. SNAP Intraoral camera

A designated doctor exported the image data, categorized it into clearly labeled folders, and made it readily available for analysis. Pediatric orthodontists then worked collaboratively to identify and label 16 types of dental conditions within the images. These included diagnoses ranging from healthy teeth to various stages of caries, pulpitis, fillings, and other conditions, ensuring consistent labeling across the dataset. This thorough organization of the images was crucial for facilitating subsequent data analysis and diagnostic review.

In the second stage, six doctors collaborated to process the questionnaire responses. The doctors divided the workload by distributing the completed questionnaires among themselves. Each doctor was responsible for manually entering the data into a carefully designed Excel spreadsheet, ensuring that every field from the questionnaire was accurately represented in the database. The doctors cross-checked each other's entries to minimize human error and ensure consistency in the data entry process. Once all the responses were recorded, the individual spreadsheets were consolidated into a single, unified dataset.

Finally, in the third stage, the comprehensive dataset, now consisting of both image data and questionnaire responses, was compiled into a summary report. This report provided a complete overview of the data collected, including statistical summaries, key findings, and metadata documentation. This report served as the foundation for the next phase of the project: data labeling.

The team used the Python Label Studio package to classify each image into one of the 16 predefined categories for the data labeling process. This tool made it easier for everyone to manage and organize the dental images effectively. To ensure all team members were comfortable with the software, the ULL team provided hands-on training to resident doctors, guiding them through the step-by-step labeling process (Figure 5). In addition to the in-person sessions, they also created a video tutorial demonstrating how to use another annotation tool, the VGG Image Annotator (VIA), so the doctors could practice independently whenever needed.

A crucial part of the project was securing ethical approval from the Mongolian National University of Medical Sciences research boards. This approval ensured that the data was handled responsibly and paved the way for the project's open data initiative. By getting the green light from the ethics board, the team ensured that the data could be accessible for future research, fostering transparency and encouraging collaboration within the scientific community.

The data labeling process ran slowly but smoothly, and everyone was properly trained and had ethical approval. The combination of thorough preparation, clear communication, and the right tools helped ensure the data was classified accurately, setting the stage for meaningful analysis of the dental images.

Figure 5. Training sessions for image labelling

This collaboration ensured consistent labeling across different diagnoses, such as healthy teeth, various stages of caries, pulpitis, fillings, and more. Each diagnosis was carefully marked by outlining the relevant areas in the image, ensuring that the condition was accurately highlighted.

If there were any uncertainties during the diagnosis, the team cross-referenced the patient's medical record to confirm the findings. The dental conditions were labeled in detail, including categories like healthy enamel, staining, superficial cavities, deep caries, pulpitis, fillings, fluorosis, hypoplasia, and other specific dental issues. This methodical approach helped maintain precision throughout the labeling process. The 16 categories chosen for the label are described below.

So - Healthy

The tooth enamel is intact, with a white, glossy surface. The borders are marked by fully enclosing the tooth ridge.

C1 – Stain

The enamel of the teeth is intact, but in certain areas, its color has changed from yellowish white to dark brown or black. These stains are primarily found in the grooves of the teeth and on the labial and lingual collar surfaces. The affected areas are marked by encircling the regions where color changes are detected (Figure 6).

So - Healthy

C1 - Stain

C - Caries

Figure 6. Categories (So, C1, and C)

C2 - Superficial Perforation (Superficialis)

Tooth enamel has discolored (ranging from white to black) and developed a rough surface. Cavities have formed at the enamel level, with superficial cavities potentially occurring on any tooth surface.

C3 - Caries Media

In cases of acute caries progression, the enamel appears white and dead, even without forming a cavity. When a cavity is present, it is surrounded by yellow-white dead enamel and filled with pale yellow softened cusps. Chronically progressing caries in this stage are characterized by a dark brown hard crust.

C4 - Deep Caries

The cavity in this stage is filled with dark, softened crimson material. The carious cavity is large, covering several tooth surfaces and deepening significantly.

P – Pulpitis

Pulpitis involves a large perforation cavity, with inflammation detected in the surrounding soft tissues. Multiple tooth surfaces may be affected by the carious cavity.

F – Filling

Cement fillings in children are usually vibrant and colorful, standing out against the natural tooth. Light-cured fillings are subtler, with borders close to the tooth's natural color.

E – Extracted

The diagnosis of extraction is based on the child’s age, the condition of the adjacent teeth, and the corresponding markings on the medical card.

F1 – Fluorosis

Fluorosis is identified by pale yellow and dark spots, with shallow grooves forming on the tooth's surface (Figure 7).

Figure 7. Categories (P, F and F1))

H – Hypoplasia

Hypoplasia manifests as linear and punctate defects on the tooth root surfaces that erupt simultaneously. These defects are also seen at the center of the molar’s chewing surface.

A1 – Attrition

In this stage, the anatomical structure of the tooth’s chewing surface is lost, appearing flattened with a glossy yellow hue.

T - Tetracycline Staining

Tetracycline-induced staining changes the enamel’s color from pale yellow to thick yellow, with the drug’s effect depending on the dosage used (Figure 8).

H - Hypoplasia A1 - Attrition T - Tetracycline

Figure 8. Categories (H, A1, and T)

R – Root
A portion of the tooth is missing.

Pq – Plaque

Plaque can form on any tooth surface, presenting a fluffy, rough texture. The black plaque is usually located in the collar area of the tooth ridge.

Ca - Calculus (Stone)

Calculus, or dental stone, forms as a hard, light yellow, or dark brown substance, typically in the collar area of the tooth.

S - Saforide
Saforide, a black substance that stops caries in primary teeth, forms a hard cavity (Figure 9).

Pq - Plaque Ca - Calculus S - Saforide

Figure 9. Categories (Pq, Ca, and S)

After multiple labeling sessions, a total of 1,959 images were annotated within these categories to train the machine-learning model.

IV. AI Models

Existing AI applications in healthcare are broadly categorized into diagnosis, clinical decision-making, and prognosis [2, 6]. The joint team from ULL and NUM focused on developing a machine-learning model capable of detecting and predicting tooth cavities in children based on the data collected and labeled by the MNUMS team.

To achieve this, the team first prepared training data by examining children aged 3 to 6. Each child underwent a full oral examination using an intraoral camera, which produced around ten images per child. These images and diagnoses made by pediatric dentists formed the core of the dataset. In addition to the image data, the children's guardians were asked to complete a 47-question survey covering five categories, including health information and oral care habits.

Several kindergartens volunteered to participate in the study, allowing the team to gather a larger pool of data. As part of this effort, free dental examinations were conducted in dental offices, ensuring a thorough and accessible data collection process. By combining intraoral images, diagnoses, and survey responses, the team created a comprehensive dataset that not only helped train the machine learning model but also provided valuable insights into children's oral health.

After labeling the initial batch of data, the ULL team immediately trained a prediction model. During this phase, we encountered a significant imbalance in the dataset, as certain category labels were underrepresented, resulting in the model's inability to detect them. This imbalance occurred because the initial batch contained fewer samples for some categories, which affected the model's performance by skewing predictions toward the more common categories (Figure 10).

In the second phase of data labeling, we focused on resolving this issue. We took a targeted approach to gather additional data specifically for the underrepresented categories, aiming to create a more balanced dataset. While this effort improved the overall distribution (Figure 10), some categories remained imbalanced due to their rarity in general cases, although they are still crucial to classify.

However, this remaining imbalance is manageable and can be addressed post-deployment. Once the model is deployed online, continuous data collection and periodic reviews will allow us to refine and rebalance the dataset further. This iterative process will help improve the model's accuracy and ensure reliable predictions across all categories.

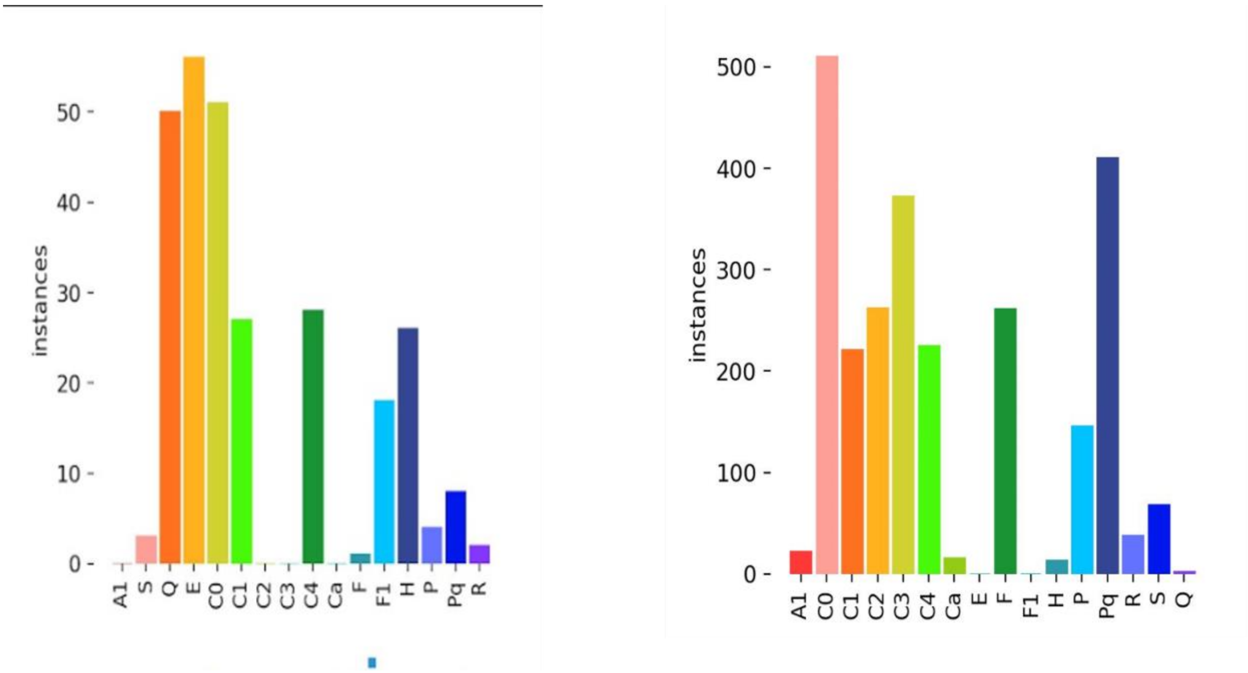

Figure 10. Data cleaning and preprocessing

To train our model, we chose to use the YOLOv8 (You Only Look Once, Version 8) framework, which is well-known for its impressive efficiency and accuracy in object detection tasks. As shown in Figure 11, YOLOv8 is a state-of-the-art deep learning model designed specifically for real-time object detection and classification, making it a perfect fit for analyzing medical images [7].

What makes YOLOv8 stand out is its architecture, which processes the entire image in a single pass. This means it can identify regions of interest and categorize them all at once, resulting in faster and more precise detection than traditional methods requiring multiple processing steps. In our project, we used YOLOv8 to detect tooth cavities in intraoral images, taking full advantage of its ability to quickly analyze visual data and make accurate predictions.

We trained the model using our carefully labeled dataset, which was processed and balanced to create a comprehensive collection of images for the model to learn from. We randomly split the dataset into a training set of 1,602 images and a validation set of 357. After completing the training, we also utilized an unlabeled dataset to test the model's performance. Although we still faced some data imbalance in certain categories, as illustrated in Figure 10, YOLOv8's robustness helped it effectively manage these challenges.

Another reason we chose this model was its flexibility. It allows us to fine-tune its performance even after deployment through online updates and revisions to the dataset. This means we can continuously improve the model's accuracy and effectiveness over time, ensuring that it remains a reliable tool for detecting tooth cavities and contributing to better oral health outcomes.

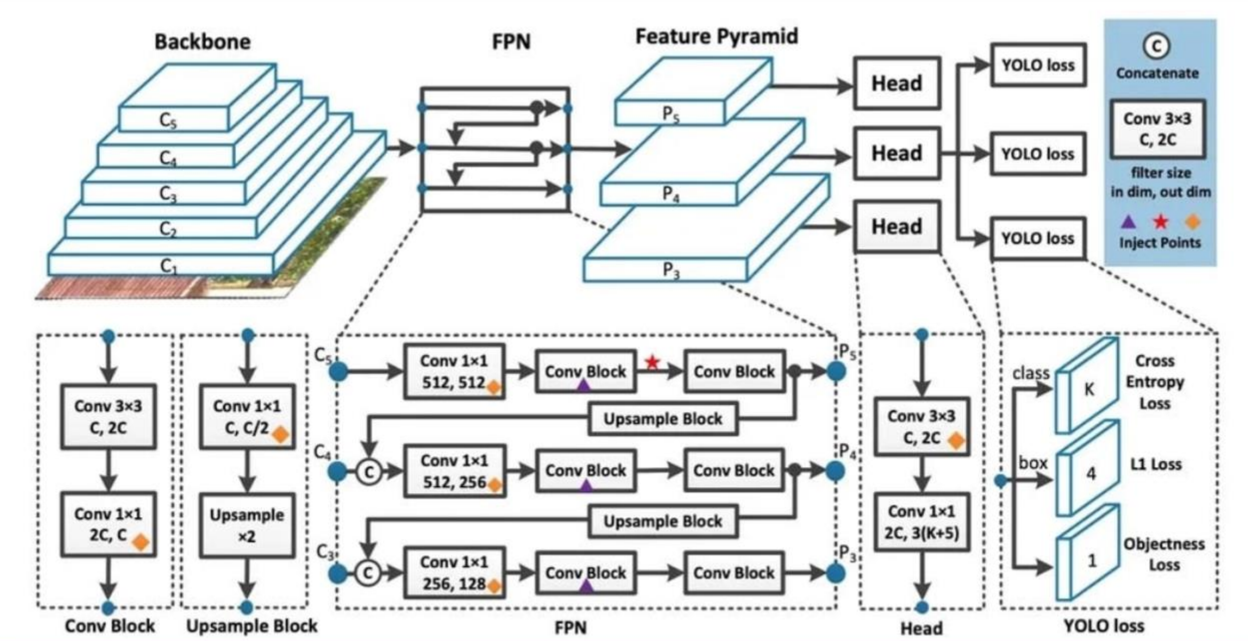

Figure 11. Training model (YOLOv8 architecture)

The results of our model training are illustrated in Figure 12. The model demonstrated strong performance in detecting and classifying certain tooth categories, while its effectiveness varied for others. This variation was primarily attributed to the imbalance in the dataset, where some categories were underrepresented, resulting in less accurate predictions for those specific cases. However, we anticipate significant improvements in the model's overall performance as the dataset expands and we gather more balanced data. By continuously acquiring and refining this data, the model would be able to learn more accurate representations of the underrepresented categories.

Furthermore, after deployment, we planned to fine-tune the model through ongoing updates and retraining with a more diverse and comprehensive dataset. This iterative approach aimed to address the current shortcomings and enhance the model's predictive accuracy across all tooth categories, ensuring it remained a reliable tool for dental diagnostics.

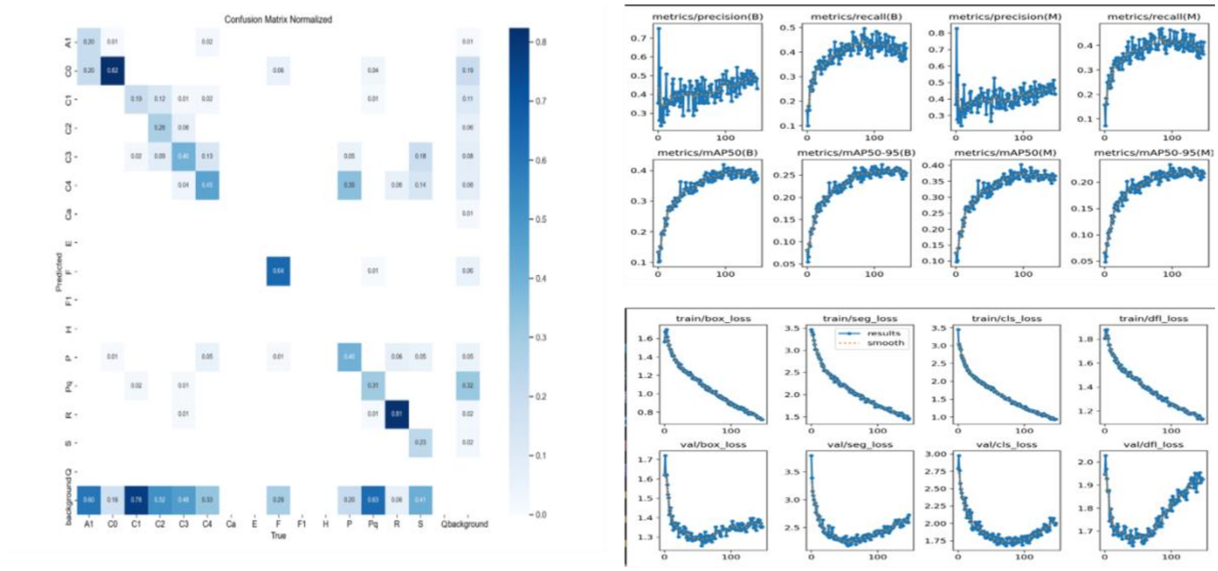

Figure 12. Training results

V. Results and Analysis

Through model training and validation, we tested our model in a real-life use case. Doctors uploaded intraoral camera images and received predictions, sharing feedback on the most critical applications. For instance, it is essential to differentiate between the two types of Saforide images used in the clinic. In one scenario, glossy black surfaces were identified as Saforide, with recognition rates ranging from 0.22 to 0.8, which aligned with our expectations. Parents were advised to continue monitoring their child's teeth even after successful treatment.

However, if the black color weakened or lightened, the recognition module accurately identified it as disease progression rather than Saforide. In the image below, the AI misinterpreted the painted surface of Saforide as C3 or C4 caries (Figure 13).

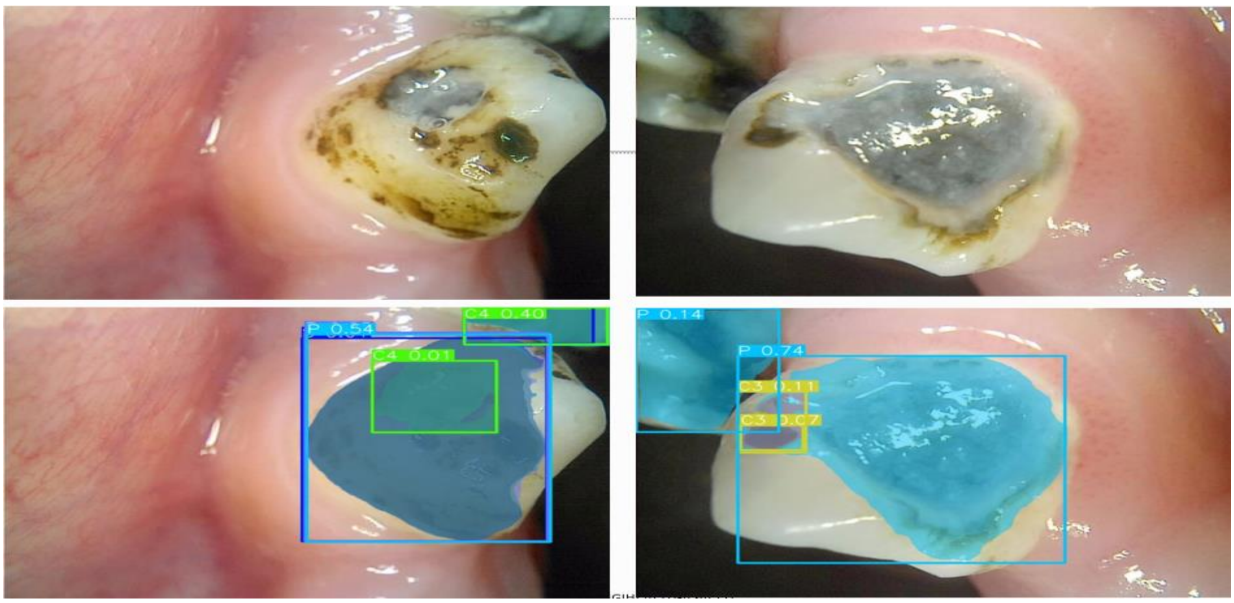

Figure 13. Progression of Caries After Application of Saforide

One challenge with Saforide application is distinguishing it from small black spaces between teeth, which are often mistakenly recognized as Saforide. To improve accuracy, we suggest adding this distinction to future image labeling protocols (Figure 14).

Figure 14. Recognition of Interdental Space

The doctors have observed that the detection of caries in the hard tissues of C2, C3, and C4 teeth, based on parameters such as the size of the caries and changes in color, has an accuracy rate exceeding 0.60. However, a notable challenge with the model arises in detecting tooth plaque (Pq). As seen in the training results (Figure 12), Pq is often mis-detected as background or soft tissues inside the mouth, such as gums, or classified incorrectly as C0 (healthy tissue). This misclassification lowers the overall confidence in detecting plaque accurately. While the model performs well in identifying more severe stages of caries, the dentist has expressed concern about this limitation, emphasizing the need to improve plaque detection. Addressing these misdiagnoses could significantly enhance the practical utility of the model.

To further enhance this detection, we plan to collect additional images using our oral cameras and finalize the web-based system by October 2024. This will involve discussions with several doctors to ensure optimal performance (Figure 15).

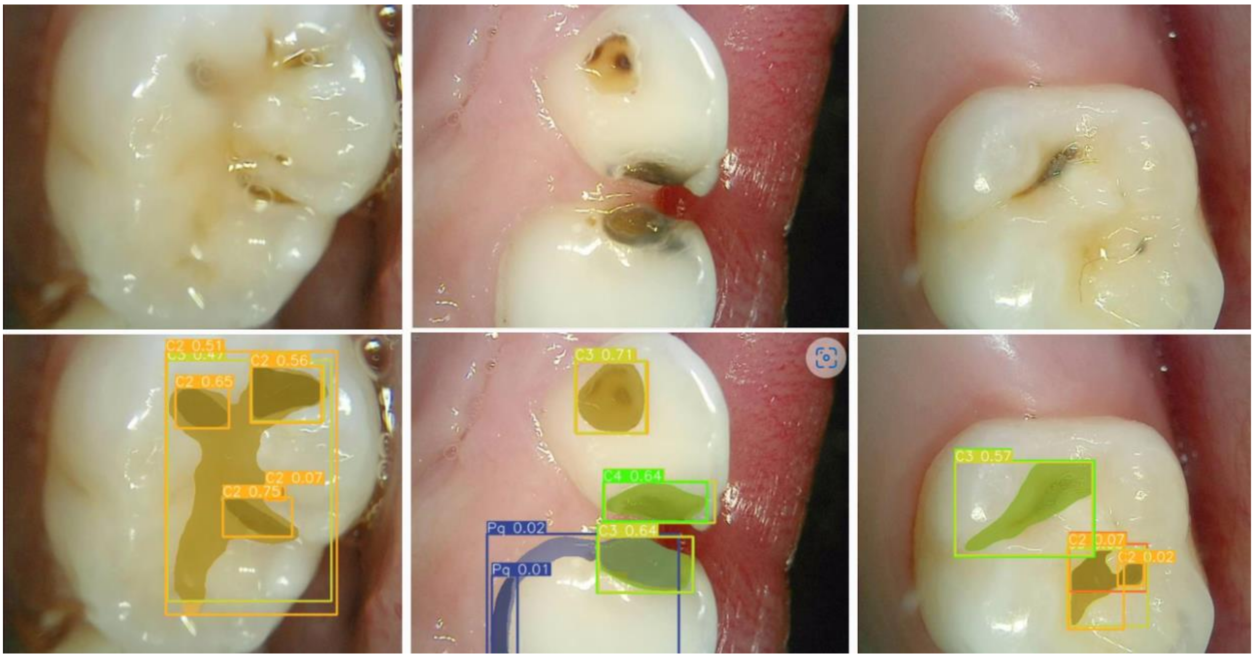

Figure 15. Cavity Type Recognition

For other conditions, although less data labeling has been conducted due to a lack of images of those categories, the identification index remains high. These conditions can be easily recognized by unique signs and indicators that set them apart from those with more extensive data. For example, dental fillings are distinguished by their different color from the natural teeth and their clearly defined borders. Additionally, more photos need to be collected and labeled to accurately identify changes in the color of old fillings and gaps formed between the teeth (Figure 16).

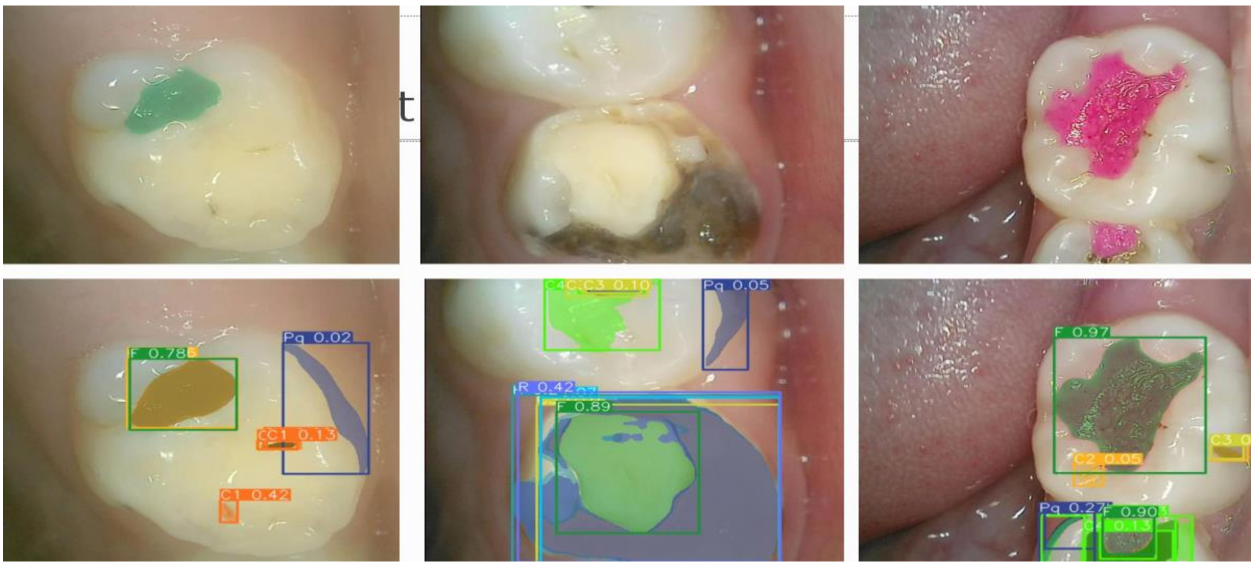

Figure 16. Recognition of Child Tooth Fillings

When evaluating our web-based machine learning system, we aimed to assess performance across various categories, including the ability to detect changes in the number of teeth. To do this, it is necessary to analyze images with multiple teeth, starting from single-tooth images, while increasing the size of the image as more data becomes available from the parents. For example, images captured on mobile phones and cameras larger than 2 MB cannot be processed by the system.

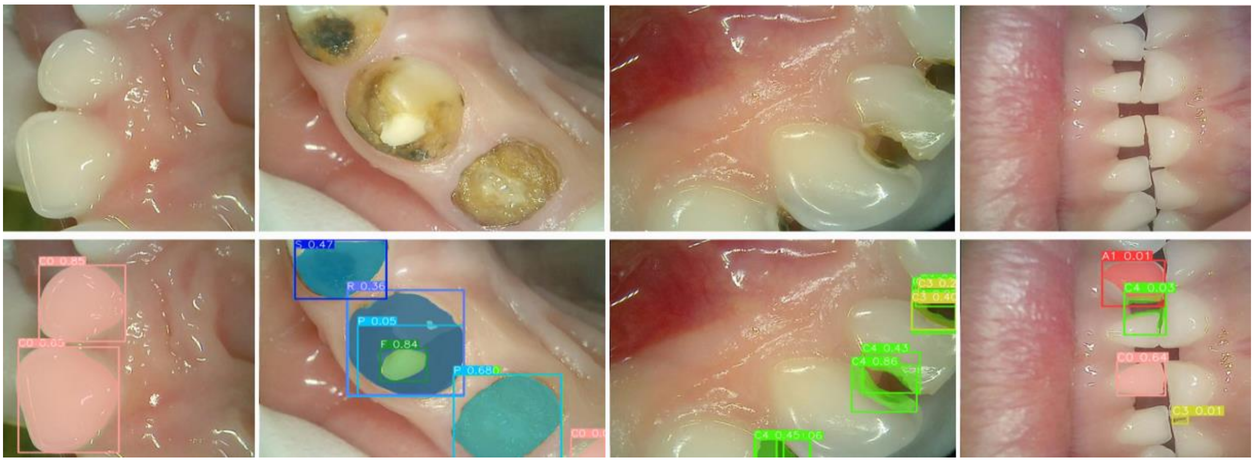

Figure 17. Image Recognition with Multiple Teeth

When analyzing new images with single or multiple teeth, it is possible to diagnose all detectable types of dental conditions present in the clinic (Figure 17). For identification purposes, labeling a single tooth may be sufficient in some cases, but for more complex or low-value cases, it is recommended that multiple images be labeled together to ensure a more accurate diagnosis.

VI. Web application

Our initial model was deployed on the website www.digihealth-mongolia.com in May 2024, marking a significant milestone in the project's progression. The website functions not only as a platform for users but also as a verification system for the model at an early stage of deployment. Medical experts are currently utilizing the platform to receive pre-diagnoses for intraoral images submitted through the system. If an expert disagrees with the classification provided by the model, they have the option to upload the image for further data collection. This feedback loop allows for continuous improvement of the model by incorporating expert-labeled data into future training stages.

In this manner, the website serves as a dynamic and interactive tool, contributing to both real-time diagnosis support and the ongoing refinement of the machine learning model. This process ensures that the model becomes more robust and accurate over time as more data is added and reviewed. The platform's design facilitates easy navigation and efficient use by medical professionals, helping to bridge the gap between automated analysis and human expertise. Figure 18 provides an overview of the website's workspace, highlighting its user-friendly interface, which is crucial for both rapid adoption by the experts and the seamless integration of additional data into the system. This iterative approach of model verification, correction, and expansion ensures that the system remains adaptive, allowing it to continually improve and evolve with each phase of development.

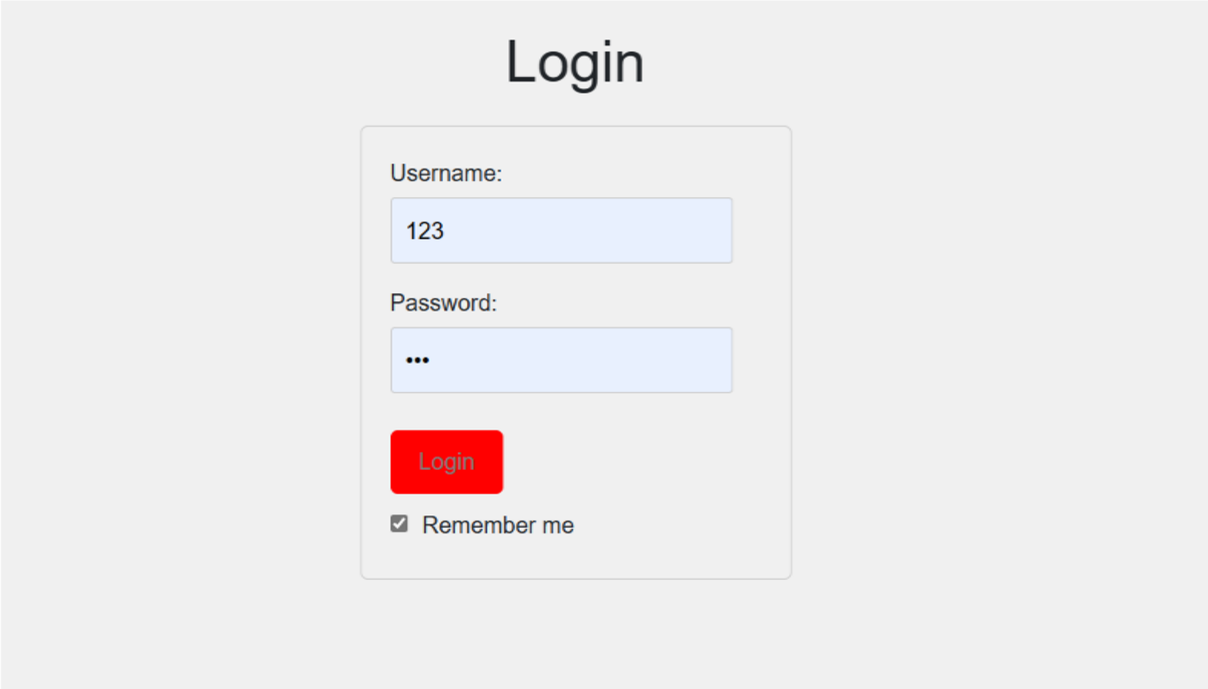

Figure 18. Pilot Case 3 website – Visualization

To further illustrate the functioning of the pilot case, we have included an example of the model's prediction in Figure 19. In this image, the model successfully identifies and segments the affected area of a tooth, providing an evaluation of the condition. On the right-hand side of the image, the model-generated segmentation clearly highlights the problematic region, offering a visual representation of its diagnostic capabilities. This segmentation is crucial for both medical professionals and patients, as it pinpoints the exact location of the dental issue, allowing for targeted treatment and a more informed decision-making process.

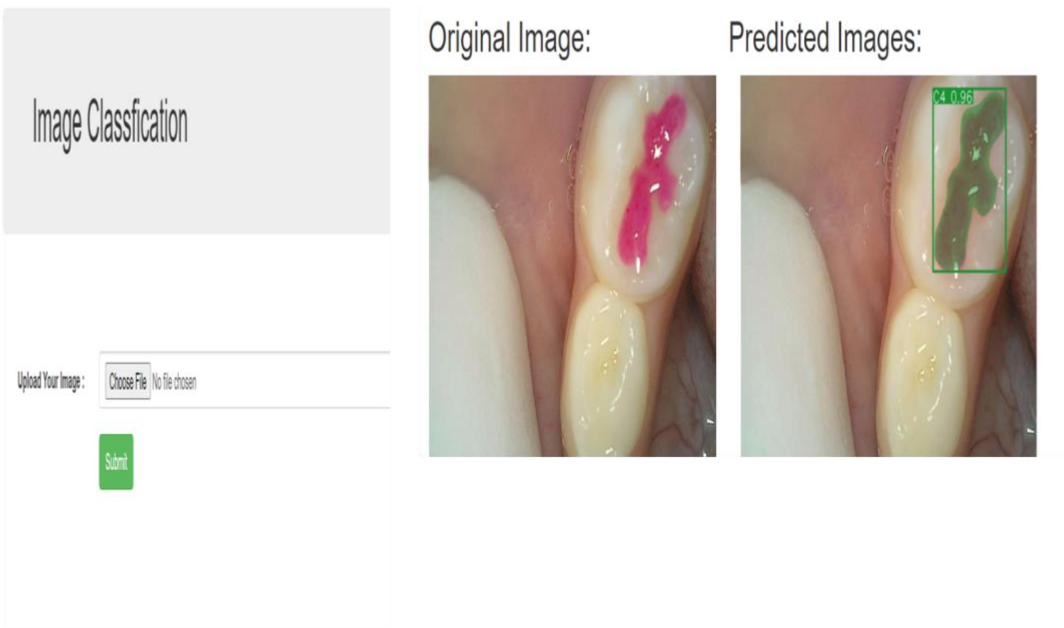

Figure 19. Prediction results

The model's ability to accurately detect and outline the affected tooth area underscores its potential as a reliable tool for early diagnosis. By integrating such precise visual feedback into the diagnostic workflow, the system enhances the efficiency of clinical decision-making. This feature is particularly valuable in remote or underserved regions, where access to specialized dental care may be limited. Over time, as the model is exposed to more data and undergoes further refinement, its predictions will become even more accurate and detailed, increasing its utility in real-world dental practice. Figure 13 showcases not only the current performance of the model but also its promise for future applications in the field of dental health.

VII. Conclusion and perspectives

As part of the project, the research team successfully gathered and labeled data from 650 children, resulting in 1,959 diagnostic (labeled) images representing 16 different dental conditions. By developing an AI-driven system, we aim to improve the early diagnosis and treatment of dental caries in children. This innovative approach holds significant potential to tackle healthcare disparities, especially in rural areas where access to dental services is limited.

Looking ahead, there is a need to expand the AI system to encompass more comprehensive diagnostic capabilities, such as the recognition of mucosal diseases and the interpretation of X-rays. Furthermore, additional research should explore how AI can be applied to other healthcare challenges in Mongolia, including public health initiatives and the provision of healthcare in remote regions. By continuing to refine AI-driven diagnostic systems, we can reduce the burden on healthcare professionals, ensure more timely diagnoses, and improve overall healthcare outcomes.

This project highlights the importance of interdisciplinary collaboration between medical professionals, AI researchers, and academic institutions. By integrating AI technologies into medical education and practice, Mongolia can strengthen its healthcare system and provide more effective care to its citizens. The success of this project also sets the foundation for future advancements in AI-driven healthcare solutions.

The ongoing development and application of AI technologies hold significant potential to transform healthcare delivery in Mongolia. As more data is collected and algorithms are further refined, AI systems will provide scalable solutions adaptable to various medical fields. With sustained research and investment, Mongolia can lead the way in utilizing AI to improve healthcare access and quality for its population.

References

1. Annual report of basic health indicators of Mongolia 2023. Ministry of Health of Mongolia. https://1212.mn/mn/statistic/statcate/48171317/table/48171317. Accessed September 20, 2024.
2. Devito KL, Barbosa FS, Felippe Filho WN. An artificial multilayer perceptron neural network for diagnosis of proximal dental caries. Oral Surg Oral Med Oral Pathol Oral Radiol Endod. 2008;106(6):879-884.
3. Talpur S, Azim F, Rashid M, Syed SA, Talpur BA, Khan SJ. Uses of different machine learning algorithms for diagnosis of dental caries. J Health Eng. 2022;Vol:2022.
4. You W, Hao A, Li S, Wang Y, Xia B. Deep learning-based dental plaque detection on primary teeth: a comparison with clinical assessments. BMC Oral Health. 2020;20:141.
5. Casalegno F, Newton T, Markram H. Caries detection with near-infrared transillumination using deep learning. J Dent Res. 2019;98(11):1227-1233.
6. Lee JH, Kim DH, Jeong SN, Choi SH. Detection and diagnosis of dental caries using a deep learning-based convolutional neural network algorithm. J Dent. 2018;77:106-111.
7. Jocher G. YOLOv8 [computer program]. GitHub. https://github.com/ultralytics/YOLOv8. Published 2022. Accessed September 20, 2024.