International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 11 | Nov - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

Flexible AI Assistance in Dentistry: Real-Time Disease Detection in Intraoral Video Streams Using YOLOv8

Prof. Dilip S[1], Raghvendra N O[2], Thejas M P[3], Roshan D R[4], Rahul S P[5]
[1]Project Guide, Department of ISE, East West Institute of Technology, Bengaluru, Karnataka, India
[2]UG Scholar (1EW23IS073), Department of ISE, East West Institute of Technology, Bengaluru, Karnataka, India
[3]UG Scholar (1EW23IS108), Department of ISE, East West Institute of Technology, Bengaluru, Karnataka, India
[4]UG Scholar (1EW23IS083), Department of ISE, East West Institute of Technology, Bengaluru, Karnataka, India
[5]UG Scholar (1EW23IS075), Department of ISE, East West Institute of Technology, Bengaluru, Karnataka, India

***Abstract*—Dental diagnosis relies on visual examination of the oral cavity, and in busy clinics the diagnostic task is separated from registration, scheduling and record keeping, which are often handled on paper or in unconnected tools. This paper presents DentalAI, a web-based dental diagnosis and clinic management system that places AI-assisted intraoral image analysis, performed by a YOLOv8n object-detection model, in the same workflow as patient management. Receptionists and doctors sign in through role-based authentication. The receptionist registers patients, assigns a doctor, and sets the appointment date, time and consultation status. The assigned doctor sees the patient on a personal dashboard, uploads an intraoral image (camera capture is offered by the interface), and receives preliminary findings from the YOLOv8n model, which locates abnormalities in the image and assigns each detection one of the documented classes: calculus, cancer, caries, gingivitis and ulcer. Each finding is presented with the model's confidence score, a bounding box marking the detected region, an explanation, a recommended next step, treatment, prescription and precautions. The doctor adds clinician notes and selects a follow-up date or "No Need", and the interface offers PDF download and saving of the report to the patient record. Appointment details are delivered as Telegram bot messages. The project stack combines React.js, Firebase, a Python FastAPI AI service and the Telegram Bot API. The main contribution is the integration of these functions in a working prototype. This paper separates implemented from proposed functionality, reports no accuracy, precision, recall or training figures, and treats every AI finding as a preliminary result that a qualified dentist must verify.**

***Keywords*—Dental AI, Artificial Intelligence, Dental Diagnosis, Deep Learning, Convolutional Neural Networks, YOLOv8n, Object Detection, Medical Image Analysis, Clinic Management, Patient Records.**

I. INTRODUCTION

Oral health is part of general health, and oral diseases remain among the most widespread health conditions worldwide; the World Health Organization estimates that they affect close to 3.5 billion people [1]. Early identification of conditions such as caries, calculus deposits, gingival inflammation and mucosal lesions allows simpler and less costly treatment, yet many patients reach a dentist only when symptoms such as toothache have already developed.

Traditional dental diagnosis depends on the clinician's visual and tactile examination, supplemented by radiographs where required. This approach is effective in skilled hands, but it is time-consuming, it is subjective, and its consistency varies between examiners, particularly for early or subtle findings. Documentation is a second burden. Patient registration, doctor allocation, appointment booking, diagnosis notes and prescriptions are frequently recorded on paper or across separate tools, which makes it hard to retrieve a patient's history, to see who is waiting for which doctor, and to remind patients about follow-up visits.

Artificial intelligence is increasingly used in healthcare to support, not replace, clinical judgment [2]. Deep learning, and convolutional neural networks (CNNs) in particular, has become the dominant approach for medical image analysis [3], and reviews of dental applications describe growing use of CNNs for diagnostic tasks on dental images [4], [5]. At the same time, a model that is not embedded in the clinic's daily workflow is unlikely to be used consistently. This motivates a system in which image analysis, patient records, scheduling and reporting share one interface.

Problem addressed. Given a dental clinic with receptionists and doctors, how can a single application register patients and assign them to doctors, let the assigned doctor capture or upload an intraoral image and obtain preliminary AI-detected findings with confidence values and detection regions, produce a structured diagnostic report that is saved to the patient record, and schedule and announce follow-up appointments, while keeping the clinician responsible for the final decision?

The objectives of DentalAI are to:

1) provide role-based access for doctors and receptionists;
2) support patient registration, editing, search, doctor assignment and consultation-status tracking;
3) accept intraoral images by camera capture or file upload;
4) detect potential dental abnormalities with a YOLOv8n model and display class, confidence and detection region;
5) generate a structured report with treatment, prescription, precautions and clinician notes, downloadable as PDF and savable to the patient record;
6) support follow-up scheduling; and
7) send appointment details through a Telegram bot.

The main contributions of this work are: (i) an integrated prototype that combines AI-assisted intraoral image analysis with patient management, scheduling, reporting and notification; (ii) a doctor-in-the-loop report in which AI findings are explicitly labelled as not a confirmed diagnosis and are accompanied by editable clinician notes; and (iii) an explicit separation between implemented, partially evidenced and

proposed functionality, with an evaluation that reports only what the project evidence supports. These contributions are not presented as unprecedented. Individual capabilities exist elsewhere; the contribution lies in their integration at prototype level.

The remainder of the paper is organized as follows. Section II reviews related work, Section III describes the proposed system, Section IV presents the architecture, Section V explains the methodology and the underlying CNN and YOLO concepts, Section VI details the implementation, Section VII reports the functional evaluation, Section VIII discusses results, Section IX lists limitations, Section X outlines future scope, and Section XI concludes.

II. RELATED WORK / LITERATURE SURVEY

This survey covers (1) AI-based dental image analysis, (2) CNN-based dental disease classification and detection, (3) YOLO-based object detection and its use in medical imaging, (4) deep learning in oral healthcare, (5) digital dental clinic management, and (6) AI-assisted clinical decision support. It is limited to sources that could be identified with complete bibliographic details; where the available literature is limited or could not be verified, this is stated explicitly.

A. AI-Based Dental Image Analysis and Deep Learning in Oral Healthcare

Hwang *et al.* [5] give an overview of deep learning in dentistry, explaining core concepts and surveying applications such as image-based detection and classification. Schwendicke *et al.* [4] conducted a scoping review of CNN studies on dental image diagnostics, and later discussed both the opportunities and the challenges of AI in dentistry, including data quality, validation and responsible use [6]. These works establish that dental imaging is a mature target for deep learning and that rigorous validation is essential. They are largely concerned with models rather than with the surrounding clinical workflow.

B. CNN-Based Dental Disease Detection

Lee *et al.* [7] evaluated a deep CNN for detecting and diagnosing dental caries on periapical radiographs. Krois *et al.* [8] applied CNNs to periodontal bone loss on panoramic radiographs and compared the model with dentists' assessments. Both address radiographic images and a single disease family. They do not provide patient registration, scheduling, report generation or notification, and they do not use intraoral photographs, which is the input DentalAI accepts.

C. YOLO-Based Object Detection

The foundational CNN architectures that underlie modern detectors include LeNet [9] and the deeper network of Krizhevsky *et al.* [10]. Two-stage detectors such as Faster R-CNN [11] first propose regions and then classify them. YOLO, introduced by Redmon *et al.* [12], reframed detection as a single regression pass that predicts boxes and class probabilities directly from the whole image, which supports fast inference. The Ultralytics implementation of YOLOv8 [13] provides the family of models, including the compact "n" variant, on which DentalAI is based. Peer-reviewed work on YOLOv8 for intraoral photographs is limited. The closest verified study identified is that of Jahan *et al.* [21], who trained and compared YOLOv8 with two vision–language models on a small set of annotated intraoral photographs covering staining, micro-cavitation, cavitation, calculus and non-carious lesions; it concerns different lesion classes and a different task setting from DentalAI. No published study of YOLOv8n covering the five DentalAI classes (calculus, cancer, caries, gingivitis and ulcer) was identified, and none is claimed here.

D. Digital Clinic Management and Clinical Decision Support

Bates *et al.* [14] formulated widely cited principles for effective clinical decision support, among them that support should fit the user's workflow and be delivered at the time and place where decisions are made. This principle motivates placing AI output inside the same record that the clinician already uses. Peer-reviewed evaluations of dental clinic-management software could not be matched to sources with complete, verifiable bibliographic details during preparation of this paper; this subsection therefore makes no claim about them, and this is acknowledged as a limitation of the literature coverage.

Table I summarizes the surveyed works and the limitation of each relative to DentalAI. The literature shows strong work on image-level models and on the principles of decision support, but the reviewed works do not combine intraoral image detection with registration, doctor assignment, reporting, follow-up and notification in one system. This integration is the gap addressed here. DentalAI's scope is more modest than the research systems in one respect: it has not been quantitatively benchmarked against them.

TABLE I COMPARISON OF RELATED DENTAL AI WORKS

Ref.	Year	Approach	Task / scope	Key finding / merit	Limitation vs. DentalAI
[7]	2018	Deep CNN	Caries detection and diagnosis, periapical radiographs	Shows CNN feasibility for caries	Radiographs only; no workflow or reporting
[8]	2019	CNN, compared with dentists	Periodontal bone loss, panoramic radiographs	Compares model with dentist assessment	Single condition; no clinic management
[4]	2019	Scoping review	CNNs for dental image diagnostics	Maps emerging evidence on CNN diagnostics	Review; no system or workflow
[5]	2019	Overview	Deep learning in dentistry	Explains concepts and applications	Conceptual; no implementation
[6]	2020	Perspective	AI in dentistry: chances and challenges	Identifies opportunities and challenges	No implemented system
[12]	2016	YOLO detector	General real-time object detection	Single-pass box and class prediction	Not dental-specific; no clinical use
[14]	2003	CDS principles	Clinical decision support design	Support should fit the clinical workflow	Not dental; no imaging
Proposed	2026	YOLOv8n + web clinic system	Intraoral detection, patient and appointment management, reports, Telegram	Integrated prototype workflow	Prototype; no accuracy or clinical validation

III. PROPOSED SYSTEM

A. System Overview

DentalAI is a web application that unifies the receptionist's administrative tasks and the doctor's diagnostic tasks. A user signs in with a selected role. The receptionist registers a patient, records the chief complaint, assigns a doctor, and sets the appointment date, time and status. The assigned doctor sees that patient on a personal dashboard, proceeds to the scan page, and uploads or captures an intraoral image. The image is analysed by a YOLOv8n model and the findings open in a report in which the doctor adds notes, selects a follow-up date, downloads a PDF and saves the report to the patient record. Appointment details are sent to a Telegram chat. Table II lists the modules and their evidence status.

Status definitions. Implemented means the function is visible in the supplied screenshots. *Interface shown* means the control is visible but its outcome is not evidenced. *Project-provided* means the technology is part of the stated project stack but is not observable in the screenshots and requires verification. *Proposed* means the function is future work.

TABLE II DENTALAI FUNCTIONAL MODULES AND STATUS

Module	Description	Status
Role-based login	Receptionist / doctor selection, e-mail and password (Fig. 4)	Implemented
User registration	"Register here" link on login screen	Interface shown
Receptionist dashboard	Patient list, status counts, search, Register Patient, Edit / Assign (Fig. 5)	Implemented
Patient edit and assign	Edit details, assign doctor, set date, time, status (Fig. 6)	Implemented
Doctor dashboard	Assigned patients, status filters, search, scan and history buttons (Fig. 7)	Implemented
Intraoral image input	Start Camera and Upload Image (Fig. 8a)	Upload: Implemented; camera: Interface shown
AI image analysis	YOLOv8n detection; "Analysis Complete" then report (Fig. 8b)	Implemented
AI findings display	Condition, confidence, severity label, region, treatment, prescription (Fig. 10)	Implemented
Diagnostic report	Patient data, scan, findings, notes, doctor name (Fig. 9)	Implemented
PDF download, save record	Download as PDF and Save to Patient Record buttons	Interface shown
Follow-up appointment	Date picker or No Need (Fig. 9b)	Interface shown
Telegram notification	"Appointment Scheduled" messages (Fig. 11)	Implemented
Automated reminders	Reminder before next check-up	Not evidenced
Patient portal	Patient-facing login and access	Proposed

B. Roles and Access

The login screen offers "Receptionist Login" and "Doctor Login" (Fig. 4). The receptionist role covers registration, assignment and scheduling; the doctor role covers consultation, scanning and reporting. The screenshots do not show a patient-facing login, so a patient portal is treated as proposed work.

C. Technology Roles

Each technology performs a different task, as summarized in Table III. Items that cannot be established from the screenshots are labelled project-provided and require verification; no implementation detail is assumed for them.

TABLE III TECHNOLOGY STACK AND ROLE OF EACH COMPONENT

Technology	Role in DentalAI	Evidence
React.js, JavaScript, HTML, CSS	Browser interface: login, dashboards, scan page, report view	Screens
Firebase Authentication	E-mail and password sign-in; role-based access [15]	Login screen; internals to verify
Cloud Firestore	Document database for patient, appointment and report data [15]	Project-provided; schema not documented
Firebase Storage	File storage, e.g. intraoral images [15]	Project-provided; verify
Python, FastAPI	HTTP service for model inference [16]	Project-provided; verify
PyTorch, Ultralytics, YOLOv8n	Detection model and runtime [17], [13]	Report header "YOLOv8n"
CNN	Convolutional backbone inside YOLOv8n; no separate classifier evidenced	Verify
Node.js, Express.js	Additional backend integration	Project-provided; verify
Anthropic SDK / Claude API	Where applicable; exact role not established	Project-provided; verify
Telegram Bot API	Sends appointment messages [18]	Fig. 11

IV. SYSTEM ARCHITECTURE

Fig. 1 shows the layered architecture of DentalAI. Solid borders mark layers whose function is visible in the screenshots; dashed borders mark project-provided components whose internal integration requires verification. The numbering of layers below matches the figure.

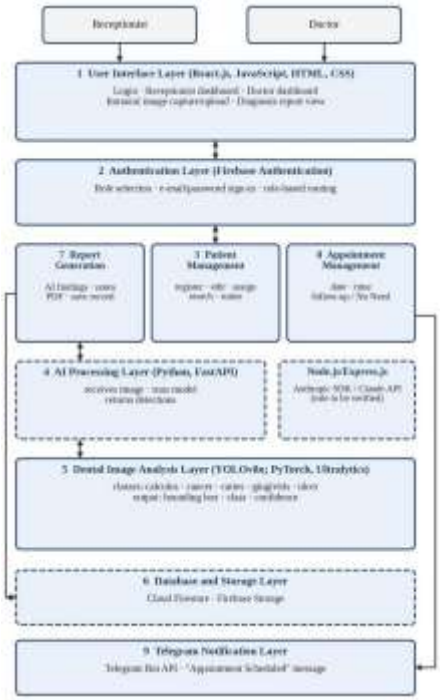


Fig. 1. Layered architecture of DentalAI. Solid: function visible in screenshots; dashed: project-provided, to be verified.

1) User Interface Layer. A React.js single-page interface presents the login screen, the two dashboards, the intraoral capture page and the report view. It is the only layer that users interact with directly.

2) Authentication Layer. Firebase Authentication verifies credentials, and the selected role routes the user to the

receptionist or doctor dashboard, so each role sees only the functions relevant to it.

3) Patient Management Layer. Registration, editing, search, doctor assignment and status updates operate on patient records held in the database layer. The doctor dashboard reads the subset of records assigned to that doctor.

4) AI Processing Layer. A Python FastAPI service is the stated AI backend. It receives the image from the interface, invokes the model and returns detections. A Node.js/Express.js integration with the Anthropic SDK is also listed in the project stack; its exact role is not established.

5) Dental Image Analysis Layer. The YOLOv8n model, run through PyTorch and Ultralytics, returns a bounding box, class and confidence for each detection. The documented classes are calculus, cancer, caries, gingivitis and ulcer.

6) Database and Storage Layer. Cloud Firestore holds structured records and Firebase Storage is the stated store for files. **7) Report Generation Layer** merges patient data, image, AI findings and clinician notes into the report, which can be exported as PDF and saved. **8) Appointment Management Layer** stores the appointment date and time set by the receptionist and the follow-up date chosen by the doctor. **9) Telegram Notification Layer** formats the appointment data and sends it through the Telegram Bot API.

Data flow. Credentials flow from the interface to the authentication layer, which returns the role. Patient and appointment data flow between the interface, the management layers and the database. For a scan, the image flows from the interface to the AI processing layer and the YOLOv8n model, and the detections flow back to the interface and into the report, which is then written to the database. Appointment data flow from the management layers to the Telegram layer.

V. METHODOLOGY

A. Workflow Stages

The methodology follows the 17 stages in Fig. 2.

1) Registration and authentication. A user signs up or signs in with an e-mail address and a password (minimum six characters, as shown on the login screen).

2) Role identification. The Receptionist or Doctor tab determines which dashboard is opened.

3) Patient registration. The receptionist records name, age, gender, a 10-digit phone number and chief complaint.

4) Doctor assignment. The receptionist selects a doctor from a list (e.g. "Dr.Dinesh (General Dentist)") and sets date, time and status.

5) Patient retrieval. The assigned doctor sees the patient with complaint and appointment, and can open the profile and history.

6) Image acquisition. The doctor starts a camera or uploads an existing image.

7) Preprocessing. Input resizing is performed by the detection pipeline; project-specific normalization steps are not documented and are not claimed.

8) Model processing. The image is passed to the YOLOv8n model.

9) Abnormality detection. The model proposes boxes and classes.

10) Confidence generation. Each detection carries a confidence value.

11) Findings presentation. The report shows condition, confidence, region, explanation and next steps.

12) Doctor review. The doctor examines the AI output against the clinical picture.

13) Treatment and notes. Treatment, prescription and precautions are shown, and the doctor adds clinician notes.

14) Report generation. A structured report is assembled.

15) Saving. The doctor downloads a PDF and saves the report to the patient record.

16) Follow-up. A next appointment date is selected, or "No Need" is chosen.

17) Notification. A Telegram message states patient, doctor, date and time.



Fig. 2. Methodology flowchart of the DentalAI workflow.

B. CNN Image Classification

A convolutional neural network (CNN) is a family of architectures built from convolutional layers. A CNN can be used as a classifier, which assigns one label to a whole image, or as the feature-extracting backbone of an object detector, which also locates objects (Section V-C). Its basic operations are as follows [9], [10]. A convolutional layer slides a kernel over the input and computes, for the k-th output channel,

$$y_k(i,j) = b_k + \Sigma_c \Sigma_m \Sigma_n \, w_k(c,m,n) \cdot x_c(i+m, j+n) \quad (1)$$

where x_c is input channel c, w_k the kernel weights and b_k the bias. The rectified linear unit (ReLU) [19] introduces non-linearity:

$$f(z) = \max(0, z) \quad (2)$$

Max pooling over a window of size s reduces spatial resolution while keeping the strongest response:

$$p_k(i,j) = \max_{0\le m,n<s} y_k(si+m, sj+n) \quad (3)$$

Stacked convolution, ReLU and pooling stages perform feature extraction, from edges and textures to higher-level shapes. The resulting features are flattened and passed to a fully connected layer and then to a softmax:

$$z = W a + b \quad (4)$$

$$P(c \mid x) = \exp(z_c) / \Sigma_{j=1,\dots,C} \exp(z_j) \quad (5)$$

where C is the number of classes. The output is one probability vector for the image; it states which class is likely but not where the abnormality is.

C. YOLO Object Detection

Object detection answers both what and where (Fig. 3 contrasts the two pipelines). In the original YOLO formulation [12], each predicted box carries centre coordinates, width, height and a confidence defined as

$$\text{Conf} = Pr(\text{Object}) \times \text{IoU} \quad (6)$$

and class-specific scores are obtained by combining conditional class probability with that confidence:

$$S_i = Pr(\text{Class}_i \mid \text{Object}) \times Pr(\text{Object}) \times \text{IoU} \quad (7)$$

Intersection over Union (IoU) measures the overlap of a predicted box B_p and a reference box B_g [20]:

$$\text{IoU} = area(B_p \cap B_g) / area(B_p \cup B_g) \quad (8)$$

Overlapping duplicate boxes are removed by non-maximum suppression (NMS) [12]: a lower-scoring box B_j is discarded if its IoU with a higher-scoring box B_i of the same class exceeds a threshold τ,

$$\text{IoU}(B_i, B_j) > \tau \quad (9)$$

and detections below a confidence threshold are not shown. IoU therefore appears in two distinct roles. In the original formulation it is part of the confidence target in Eq. (6), where the reference box is the ground-truth box [12]; in Eq. (9) it measures the overlap between two predicted boxes during post-processing. The detection confidence is a per-box score attached to each predicted detection; it is neither the IoU itself nor a measure of model accuracy.

YOLOv8 departs from the original grid-based formulation by using an anchor-free split detection head [13], [22]. Ultralytics has not published a formal research paper for YOLOv8 [22], so the exact computation of the YOLOv8n detection confidence is not restated here: Eqs. (6)–(7) are presented as the conceptual formulation of the original YOLO [12] and are not asserted to be the YOLOv8n computation. The weights, training data and thresholds of the project's model are not documented here, so no architecture details beyond the model name are claimed.

In summary, classification and detection differ in what they output. A classifier returns one probability vector for the whole image (Eqs. (4)–(5)), whereas a detector such as YOLOv8n returns zero or more boxes, each with a class and a confidence score. They do not differ in the use of convolution: YOLOv8n is itself a convolutional network whose backbone extracts features, and only its detection head and post-processing (score filtering and NMS) produce boxes. Eqs. (1)–(5) describe generic CNN operations, and the individual layer types of YOLOv8n differ from this textbook sequence.

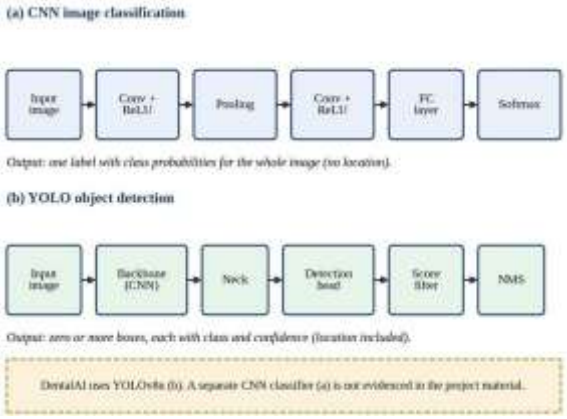


Fig. 3. CNN image classification versus YOLO object detection.

D. Relevance to DentalAI

A dentist needs to know where a suspicious region lies, so detection is more informative than whole-image classification. DentalAI accordingly uses a single YOLOv8n model as an object detector. The paper does not claim that a second, separate CNN classifier was implemented, and none is evidenced in the project material (Fig. 3, Table III). The report therefore displays the box returned by YOLOv8n (Fig. 10). For that example the interface reports X_1 = 35.6, Y_1 = 123.5, X_2 = 158.0 and Y_2 = 249.0, so the box measures

$$w = X_2 - X_1 = 122.4,\ h = Y_2 - Y_1 = 125.5 \quad (10)$$

in the interface's coordinate units. This is illustrative arithmetic on one displayed example, not a measure of accuracy. The 18% shown in Fig. 10 is the confidence score the model attached to this single detection of the caries class. It expresses how strongly the model scores that box. It is not the accuracy of the model, it is not the probability that the tooth is carious, and no calibration against clinical outcomes is reported. The confidence and NMS thresholds of the deployed pipeline are not documented, and no threshold value is assumed here.

VI. SYSTEM IMPLEMENTATION

Frontend. The interface is implemented with React.js, JavaScript, HTML and CSS, with a consistent navy and cyan theme, status badges and role-specific headers. **Backend and AI.** The AI backend is a Python FastAPI service around a YOLOv8n model (PyTorch, Ultralytics); a Node.js/Express.js integration and the Anthropic SDK are also in the stack, and neither can be characterized from the screenshots. **Firebase.** Firebase Authentication, Cloud Firestore and Firebase Storage are the project's authentication, database and file-storage services. The report header reads "DentalAI YOLOv8n", which is the only direct on-screen evidence of the model. Figs. 4–11 are screenshots of the implemented interface.

A. Role-Based Login



Fig. 4. Role-based login interface with Receptionist and Doctor tabs.

Fig. 4 shows the login card. Two tabs select the role; the form takes an e-mail and a password with a six-character minimum and offers "Register here" for new users. The note beneath states that access is role-based. The completed registration flow is not shown.

B. Receptionist Dashboard

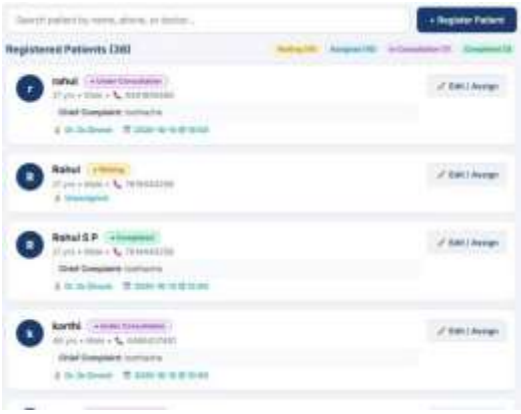


Fig. 5. Receptionist dashboard listing registered patients with status badges.

In Fig. 5 the dashboard lists 36 registered patients with counts of Waiting (16), Assigned (10), In Consultation (7) and Completed (3). A search box accepts a patient name, phone number or doctor, and "Register Patient" opens the registration form. Each card shows initial, name, status badge, age, gender, phone, chief complaint, assigned doctor and appointment date and time, with an "Edit / Assign" button. Unassigned patients show "Unassigned". The same screen therefore gives the receptionist a live view of the queue.

C. Patient Registration, Assignment and Scheduling



Fig. 6. Edit Patient Record form for details, doctor assignment, schedule and status.

Fig. 6 shows the "Update Patient & Schedule" form: full name, age, gender, a 10-digit phone number, chief complaint, an "Assign Doctor" list, appointment date and time pickers and a patient-status list (here, Under Consultation), with Cancel and Update Patient buttons. Updating this record is what makes the patient visible to the chosen doctor and what supplies the date and time for the notification in Section VI-G.

D. Doctor Dashboard



Fig. 7. Doctor dashboard with consultation counters, search and patient card.

Fig. 7 shows a welcome banner with the doctor's name and specialty, the number of assigned patients, three counters (Assigned, Under Consultation, Completed), a search box for name, ID or phone, and filter tabs. The patient card shows an ID of the form name_phone, status, complaint and appointment, with "Proceed to AI Scan & Analysis" and "View Profile & History" buttons. The doctor in this capture (Dr.Dhanush) differs from the doctor in Fig. 5 (Dr.Dinesh) because the screenshots were taken in different sessions; the paper uses them only to show interface elements.

E. Intraoral Image Acquisition and AI Analysis



Fig. 8. (a) Intraoral Image Capture page. (b) "Analysis Complete" screen before redirection to the report.

Fig. 8(a) offers "Start Camera", which is intended for a stomatoscope or another camera, and "Upload Image". In the captured state the page shows "No image yet" and the message "Camera access is not supported by this browser", so camera capture is not demonstrated; image upload is inferred from Fig. 8(b), which shows the analysed image and "Analysis Complete! Redirecting to report…". The image shown is a sample intraoral photograph used for demonstration; the screenshots do not state its source.

F. Diagnostic Report and AI Findings

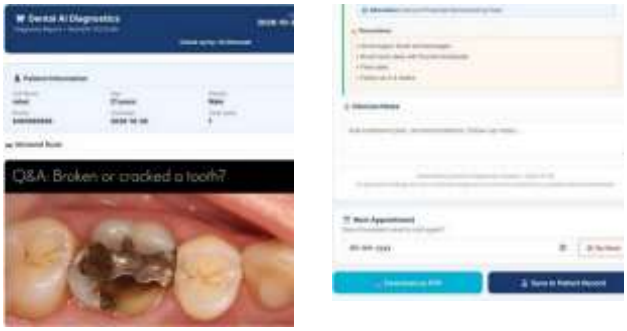


Fig. 9. AI-generated diagnostic report. (a) Header, patient information and intraoral scan. (b) Clinician notes, next appointment and actions.

Fig. 9(a) shows the report header ("Dental AI Diagnostics – Diagnosis Report • DentalAI YOLOv8n"), the date, the checking doctor ("Check up by: Dr.Dhanush") and a patient panel with full name, age, gender, phone, visit date and total visits, followed by the intraoral scan. Fig. 9(b) shows the end of the findings, the Clinician Notes field, a footer stating that AI-generated findings are not a confirmed diagnosis and must be reviewed by a qualified dental professional, the Next Appointment control (date picker and "No Need"), and the "Download as PDF" and "Save to Patient Record" buttons.



Fig. 10. AI Analysis Findings section of the report.

Fig. 10 shows one finding, "Possible Tooth Decay", with a "Mild" label and a confidence of 18%. An explanation box states that the model classified the region as caries and that the prediction cannot establish the presence or severity of decay. The detected region is given as X_1 = 35.6, Y_1 = 123.5, X_2 = 158.0, Y_2 = 249.0, with the note that these coordinates identify the model's detection box, not a confirmed anatomical location. The recommended next step is that a dentist examine the tooth. Treatment (fluoride remineralization therapy with monitoring), a prescription card (sodium fluoride gel 1.1%, topical, after brushing, once daily on the affected tooth for 4 weeks, with a warning and an alternative) and precautions follow. This is an illustrative example from the interface and is not evidence of model accuracy; the low confidence is itself a reason for clinician review. How the severity label and the treatment text are produced is not documented and requires verification.

G. Telegram Notification



Fig. 11. Telegram bot "DentalAI Appointment Assistant" showing appointment notifications.

Fig. 11 shows a Telegram bot chat. A first message reports that the integration is working as a test notification. Subsequent messages titled "DentalAI – Appointment Scheduled" list patient, doctor, date and time. One of them (Rahul S P, Dr.Dinesh, 2026-10-13, 12:00) matches a patient card in Fig. 5, which links the dashboard record to the message. Other entries,

such as "Test Patient", are test data. Messages before the next check-up, as described in the project brief, are not visible in the evidence, so automated reminders are not claimed. The bot interface follows the Telegram Bot API [18].

H. Patient Record Management

Patient records are created at registration and updated through the edit form. Each record holds demographic data, complaint, assigned doctor, schedule and status, and a unique ID derived from name and phone. The doctor dashboard retrieves records for the signed-in doctor and offers "View Profile & History", and the report shows "Total Visits". The storage of saved reports and the content of the history page are not shown, so these are marked Interface shown.

VII. FUNCTIONAL / EXPERIMENTAL EVALUATION

DentalAI is a student prototype and the project material contains no benchmark dataset or test log, so no accuracy, precision, recall, F1-score, response-time or user-adoption figures are reported. Table IV lists functional checks and the evidence available. "Evidence Pending" means the interface element exists but a screenshot or test record of the outcome was not supplied.

TABLE IV FUNCTIONAL EVALUATION OF DENTALAI

Test case	Input / action	Expected result	Observed result	Status
User login	Open login, choose role	Role login form shown; dashboard opens	Login screen (Fig. 4); both dashboards shown (Figs. 5, 7)	Pass (screenshots)
User registration	Click "Register here", submit	Account created	Link visible; outcome not shown	Evidence Pending
Patient registration	Register Patient	Record appears in list	36 registered records listed (Fig. 5); form not shown	Pass (indirect)
Doctor assignment	Choose doctor, update	Doctor name on patient card	Assign form (Fig. 6); doctor names on cards	Pass
Patient search	Type name / phone / doctor	Filtered list	Search boxes shown; result not shown	Evidence Pending
Image upload / capture	Upload Image; Start Camera	Image loaded	Uploaded image shown (Fig. 8b); camera unsupported in captured browser	Partial
AI image analysis	Submit image	Analysis completes, report opens	"Analysis Complete" (Fig. 8b); report (Fig. 9)	Pass
Condition detection	Run on sample image	Class and region returned	One caries finding with box (Fig. 10); other four classes not shown	Pass (single example)
Confidence display	View findings	Confidence value shown	18% shown (Fig. 10)	Pass
Report generation	Complete analysis	Structured report	Patient data, scan, findings, notes (Fig. 9)	Pass
PDF download	Download as PDF	PDF file	Button visible; file not supplied	Evidence Pending
Save patient record	Save to Patient Record	Report stored	Button visible; stored record not shown	Evidence Pending
Follow-up selection	Pick date or No Need	Choice recorded	Control visible; outcome not shown	Evidence Pending
Appointment scheduling	Set date and time	Date and time on patient card	Dates on cards (Figs. 5, 7)	Pass
Telegram notification	Schedule appointment	Message in bot chat	Messages with name, doctor, date, time (Fig. 11)	Pass (screenshots)

The rows marked Pass rest on screenshots, not on automated tests. The only AI example supplied is a single low-confidence caries finding, so detection of calculus, cancer, gingivitis and ulcer is not demonstrated and nothing is claimed about them.

VIII. RESULTS AND DISCUSSION

Functional workflow. Figs. 4–11 show that one application can carry a patient from registration to a saved report and a notification without leaving the interface: the receptionist creates and schedules the record, the doctor opens it, scans, reviews and reports, and Telegram announces the appointment.

AI-assisted analysis. The report presents a model output together with a region, a confidence value, an explanation and next steps. Equally important, it states the limits of that output: the finding box says the prediction cannot establish presence or severity, and the footer says the findings are not a confirmed diagnosis. This framing keeps the clinician in control, consistent with the decision-support principle of supporting rather than replacing judgment [14]. The 18% example does not demonstrate diagnostic ability.

Doctor–receptionist coordination. Because assignment, date, time and status are edited on one record, both roles see consistent information. The receptionist's status counters and the doctor's counters are two views of the same underlying state, which reduces manual hand-over.

Reporting, appointments and notification. A structured report with notes and a follow-up control standardizes documentation, and Telegram messages provide a lightweight channel that needs no separate patient application. *Benefits of integration.* Placing AI output in the same record as scheduling and notes avoids switching tools and keeps image, findings and decisions together for later retrieval. These are design benefits observed at prototype level; no gain in diagnostic accuracy, clinic efficiency or patient outcome has been measured.

IX. LIMITATIONS

Confirmed project limitations (from the evidence):

- No accuracy, precision, recall, F1-score, dataset size or training details are documented, so model quality cannot be assessed.
- Only one AI finding (caries, 18% confidence) is evidenced; the other four classes are not demonstrated.
- Camera capture was reported unsupported in the captured browser; upload is the evidenced path.
- PDF output, record saving, patient search results, user registration and follow-up storage lack outcome evidence.
- Automated reminders before the next check-up and a patient portal are not evidenced.
- The internal roles of FastAPI, Node.js/Express.js, the Anthropic SDK and Firebase Storage are not documented.
- Test entries (e.g. "Test Patient") coexist with sample data, and the screenshots come from different sessions.
- Operation depends on internet connectivity and external services (Firebase, Telegram).
- The literature survey is limited: no peer-reviewed evaluation of dental clinic-management software could be verified, and published work on YOLOv8 for intraoral photographs is sparse (Section II).

General limitations of AI-based dental systems that also apply here:

- Dependence on image quality, lighting, focus and angle.
- Limited dataset size and diversity, and weak generalization to other devices, populations and lighting [4], [6].
- False positives and false negatives, which are possible for any detector.
- The need for clinical validation by qualified dental professionals before any clinical reliance.
- AI findings require professional verification; they do not replace the dentist.

X. FUTURE SCOPE

1) Train with larger, more diverse and clinically annotated dental datasets.
2) Improve detection accuracy through model optimization and systematic evaluation.
3) Support additional imaging modalities, such as radiographs.
4) Improve real-time image processing.
5) Extend to more advanced dental disease classification.
6) Add a patient-facing portal.
7) Implement automated appointment reminders before the next check-up.
8) Add multilingual support.
9) Provide advanced dental analytics for clinic management.
10) Integrate with digital dental equipment such as intraoral cameras.
11) Deploy on cloud infrastructure for scalability.
12) Strengthen security and privacy of patient data.
13) Conduct clinical validation with qualified dental professionals.

Each item is proposed work and is not part of the implemented system described in Sections III–VIII.

XI. CONCLUSION

This paper presented DentalAI, a web-based system that brings AI-assisted dental image analysis, patient management, digital reporting, appointment scheduling and Telegram notification into a single clinic workflow. Role-based access separates receptionist and doctor tasks; a YOLOv8n model supplies preliminary findings with confidence values and detection regions; and a structured report, with clinician notes and a follow-up control, supports documentation and continuity of care. The screenshots demonstrate a working interface across the workflow, while the evaluation remains functional: no accuracy or clinical effectiveness is claimed, several outcomes await evidence, and several backend integrations require verification. AI-generated findings are preliminary and must always be verified by a qualified dentist. Future work should concentrate on dataset development, quantitative evaluation and clinical validation.

ACKNOWLEDGMENT

The authors thank Prof. Dilip S, Department of Information Science and Engineering, East West Institute of Technology, Bengaluru, for guidance and supervision of this project, and the department and institution for their support.

REFERENCES

[1] World Health Organization, *Global Oral Health Status Report: Towards Universal Health Coverage for Oral Health by 2030*. Geneva, Switzerland: WHO, 2022.
[2] E. J. Topol, "High-performance medicine: The convergence of human and artificial intelligence," *Nature Med.*, vol. 25, pp. 44–56, 2019, doi: 10.1038/s41591-018-0300-7.
[3] G. Litjens *et al.*, "A survey on deep learning in medical image analysis," *Med. Image Anal.*, vol. 42, pp. 60–88, 2017, doi: 10.1016/j.media.2017.07.005.
[4] F. Schwendicke, T. Golla, M. Dreher, and J. Krois, "Convolutional neural networks for dental image diagnostics: A scoping review," *J. Dent.*, vol. 91, Art. no. 103226, 2019, doi: 10.1016/j.jdent.2019.103226.
[5] J.-J. Hwang, Y.-H. Jung, B.-H. Cho, and M.-S. Heo, "An overview of deep learning in the field of dentistry," *Imag. Sci. Dent.*, vol. 49, no. 1, pp. 1–7, 2019, doi: 10.5624/isd.2019.49.1.1.
[6] F. Schwendicke, W. Samek, and J. Krois, "Artificial intelligence in dentistry: Chances and challenges," *J. Dent. Res.*, vol. 99, no. 7, pp. 769–774, 2020, doi: 10.1177/0022034520915714.
[7] J.-H. Lee, D.-H. Kim, S.-N. Jeong, and S.-H. Choi, "Detection and diagnosis of dental caries using a deep learning-based convolutional neural network algorithm," *J. Dent.*, vol. 77, pp. 106–111, 2018, doi: 10.1016/j.jdent.2018.07.015.
[8] J. Krois *et al.*, "Deep learning for the radiographic detection of periodontal bone loss," *Sci. Rep.*, vol. 9, Art. no. 8495, 2019, doi: 10.1038/s41598-019-44839-3.
[9] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, Nov. 1998.
[10] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, vol. 25, 2012, pp. 1097–1105.
[11] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 6, pp. 1137–1149, Jun. 2017.
[12] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2016, pp. 779–788.
[13] G. Jocher, A. Chaurasia, and J. Qiu, "Ultralytics YOLOv8," ver. 8.0.0, 2023. [Online]. Available: https://github.com/ultralytics/ultralytics (accessed Nov. 2, 2026).
[14] D. W. Bates *et al.*, "Ten commandments for effective clinical decision support: Making the practice of evidence-based medicine a reality," *J. Amer. Med. Inform. Assoc.*, vol. 10, no. 6, pp. 523–530, 2003.
[15] Google, "Firebase documentation." [Online]. Available: https://firebase.google.com/docs (accessed Nov. 2, 2026).
[16] S. Ramírez, "FastAPI documentation." [Online]. Available: https://fastapi.tiangolo.com (accessed Nov. 2, 2026).
[17] A. Paszke *et al.*, "PyTorch: An imperative style, high-performance deep learning library," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, vol. 32, 2019, pp. 8024–8035.
[18] Telegram, "Telegram Bot API." [Online]. Available: https://core.telegram.org/bots/api (accessed Nov. 2, 2026).
[19] V. Nair and G. E. Hinton, "Rectified linear units improve restricted Boltzmann machines," in *Proc. 27th Int. Conf. Mach. Learn. (ICML)*, 2010, pp. 807–814.
[20] H. Rezatofighi *et al.*, "Generalized intersection over union: A metric and a loss for bounding box regression," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2019, pp. 658–666.
[21] M. Jahan, A. I. Siam, L. Z. Pronay, S. Ahmed, N. Mohammed, J. Dudley, and T. H. Farook, "Comparative evaluation of vision–language models for detecting and localizing dental lesions from intraoral images," *J. Imaging*, vol. 12, no. 1, Art. no. 22, 2026, doi: 10.3390/jimaging12010022.
[22] Ultralytics, "Ultralytics YOLOv8." [Online]. Available: https://docs.ultralytics.com/models/yolov8/ (accessed Nov. 2, 2026).

Note: Studies of dental clinic-management software could not be verified with complete bibliographic details and are not cited; see Section II-D.