

OP 02

AI-Powered Dental Prosthetic and Recommendation System

Nilupul H.M.[#], Fonseka N., Akalanka A.U.

School of Computing, ESOF Metro Campus, Galle, Sri Lanka

[#]Corresponding author: akilanilupul3@gmail.com

Abstract

Dental prosthetic selection is largely subjective, exhibiting inter-practitioner variability and typically requiring multiple consultation visits. Existing dental artificial intelligence (AI) solutions concentrate on radiographic diagnosis and stop short of producing actionable, patient-specific treatment recommendations that integrate clinical, demographic, and financial factors. This study developed an integrated, explainable web-based decision-support prototype that combines automated dental radiograph analysis with structured patient data to generate ranked prosthetic treatment recommendations supporting consistent clinical workflow. A three-tier system (Next.js, FastAPI, MongoDB, Flask reporting service) was implemented. A YOLOv8-Nano detector was fine-tuned for 20 epochs on the public Rob flow Dental X-Ray Panoramic v3 dataset (16,995 annotated radiographs, ten conditions, CC BY 4.0). A hybrid rule-based and heuristic scoring engine over a nine-treatment knowledge base produces ranked top-five recommendations with rationale and indicative costs. On the validation split, the fine-tuned detector achieved Precision 0.60, Recall 0.56, mAP@0.5 0.55, and mAP@0.5:0.95 0.28 across ten dental classes. The end-to-end pipeline executed within seconds on commodity hardware, returning annotated radiographs alongside ranked recommendations and downloadable PDF reports compiling patient data, detected conditions, suitability scores, rationales and indicative costs for clinical handover. The prototype demonstrates the feasibility of an integrated, explainable, web-deployable decision-support pipeline linking radiograph analysis with patient-specific data for prosthetic planning. Detection accuracy is moderate at this prototype stage. Expanded training data, held-out clinical validation and electronic health record integration are identified as priorities for future work.

Keywords: *clinical decision support, dental prosthetics, hybrid recommendation engine, object detection, YOLOv8*

Introduction

Dental prosthetic treatment planning is a complex clinical activity that requires the integration of radiographic findings, demographic factors, the patient's medical history and financial considerations. In conventional practice, the selection of a suitable prosthetic restoration is largely a subjective decision dependent on the experience of the individual practitioner, which leads to inter-practitioner variability and consultation pathways that may extend over multiple visits before a treatment plan is finalised. This variability is particularly pronounced in resource-constrained settings where access to specialist consultation is limited.

Artificial intelligence (AI) has emerged as a transformative force in dentistry over the past decade [1]. Deep-learning approaches have produced strong performance across a range of diagnostic tasks, including the automated detection of caries, periapical lesions and other dental pathologies on panoramic and periapical radiographs [2,3]. Comprehensive scoping reviews of machine-learning applications in dentistry confirm rapid growth in the number and scope of such studies [4].

Object-detection architectures, and convolutional neural networks (CNNs) more broadly, have featured prominently in this body of work. Chen demonstrated automatic tooth detection and numbering on periapical radiographs using deep object-detection methods [5]. Londono evaluated CNN algorithms for landmark detection on cephalometric images [6]. Earlier work also explored hybrid multi-convolution networks for dental X-ray analysis [7]. A growing body of literature confirms that deep learning can deliver clinically meaningful detection performance on dental imagery [8].

In parallel, commercial AI-driven dental imaging platforms have emerged. Systems such as Overjet [9], Denti.AI [10], WeDiagnostiX [11] and Pearl AI [12] provide automated radiograph analysis, often integrated with practice-management software. These offerings emphasise the detection and annotation of pathology, but generally stop short of producing structured treatment recommendations that are tailored to the individual patient. The system presented in this paper is explicitly distinguished from these platforms because, rather than confining AI to image diagnosis, it integrates explainable, patient-specific treatment recommendations into the same workflow. As Aldahabani [13] and Najeeds [14] highlight, the application of AI to dental prosthetics and restorative planning, as opposed to diagnostic imaging, remains comparatively under-explored.

This gap motivates the present study. While detection-only systems can flag conditions, the clinical question of which treatment to perform still requires the practitioner to manually map detected findings to suitable interventions and to weigh patient-specific factors such as age, comorbidities, expected outcomes and cost. An integrated decision-support system that performs this mapping explicitly has the potential to standardise reasoning, document the rationale behind a recommendation and reduce the number of consultation iterations needed to converge on a plan [15].

The aim of this work is therefore to design, implement and demonstrate an end-to-end, web-deployable decision-support prototype that (i) detects dental conditions from panoramic radiographs using a fine-tuned YOLOv8-Nano object detector, (ii) combines the detection output with structured patient information through a hybrid rule-based and heuristic scoring engine and (iii) produces a ranked, explainable list of prosthetic treatment recommendations together with a downloadable clinical-handover report.

Methods

Study approach

This study followed a design-science research approach in which a software artefact was built and evaluated rather than a clinical hypothesis tested. The research progressed through six phases. These were problem identification (subjectivity and inter-practitioner variability in prosthetic selection); a structured literature review covering deep learning in dentistry [1-4], object detection for dental radiographs [5-8] and commercial dental AI products [9-12], which confirmed that recommendation-oriented systems remain under-explored [13,14]; requirements definition covering radiograph upload, automated condition detection, patient-data capture, ranked recommendation, rationale explainability and downloadable reporting; design of a three-tier web architecture and the recommendation engine; implementation; and evaluation, measuring detection performance on the validation split using Precision, Recall and mean average precision (mAP) at intersection-over-union (IoU) thresholds of 0.5 and 0.5:0.95, with the end-to-end pipeline reviewed qualitatively on sample radiographs.

The study did not involve human participants. All images came from a publicly available, CC BY 4.0 licensed collection on Roboflow Universe. No patient-identifiable information was processed, and no ethics approval was required.

Implementation was managed as a single-developer Scrum across seven two-week sprints, each closing with a working increment, covering dataset acquisition and training-pipeline assembly (sprints 1-2), backend, persistence layer, and recommendation engine (3-4), frontend, reporting microservice and integration (5-6), and validation, documentation, and packaging (7). Version control was managed through GitHub, with sprint tasks tracked as Issues.

Dataset

Training and validation images were obtained from the Dental X-Ray Panoramic (version 3) dataset distributed through the Roboflow Universe platform under a Creative Commons Attribution 4.0 (CC BY 4.0) licence. The release contained 16,995 panoramic radiographs annotated for ten clinically relevant classes, namely Bone Loss, Caries, Crown, Filling, Implant, Missing Teeth, Periapical Lesion, Root Canal Treatment, Root Piece and Impacted Tooth. Pre-processing included auto-orientation with EXIF-orientation stripping, resizing to 640×640 pixels (stretch), conversion to greyscale, and histogram-equalisation auto-contrast. Online augmentation introduced 50 % horizontal flip, 50 % vertical flip, $\pm 15^\circ$ rotation, and ± 25 % brightness adjustment. Annotations were supplied in YOLOv8 format, and the train/validation split followed the dataset distributor's published split.

Detection Model

Object detection was performed using YOLOv8-Nano, the smallest variant of the YOLOv8 family (≈ 3.01 M parameters). The model was initialised from the publicly distributed yolov8n.pt COCO-pre-trained weights and fine-tuned end-to-end on the dental dataset. An initial 10-epoch schedule was extended by further fine-tuning to a total of 20 epochs. The Ultralytics framework's "auto" optimiser selection was used with an initial learning rate $\text{lr}_0 = 0.01$, momentum 0.937, weight decay 5×10^{-4} , and a three-epoch warm-up. Training was performed in a Google Colab environment using cloud-provisioned GPU resources, with the best-performing checkpoint (selected by validation fitness) retained for inference.

Recommendation Engine

Following detection, an explainable recommendation engine maps the detected conditions to candidate treatments and ranks them according to a composite suitability score. The engine maintains a knowledge base of nine prosthetic and restorative treatments, namely Dental Implant, Dental Filling, Crown, Bridge, Denture, Root Canal Treatment, Veneer, Inlay/Onlay and Composite Bonding. Each treatment record contains the conditions it can address, an indicative cost range (in Sri Lankan Rupees), an expected duration, a documented success-rate estimate, age-suitability bounds, contraindications, a risk-level label, and qualitative descriptions of pros and cons. The knowledge base was compiled by the authors from published prosthodontic reference material and indicative cost information for Sri Lankan clinical practice; its parameter values are heuristic estimates selected by the researcher and have not yet been calibrated by an expert panel.

For each radiograph, the engine first extracts the set of detected condition labels from the detector output and selects the treatments whose condition lists include any detected label. Each candidate receives a base suitability score of 80 on a 0-100 scale, adjusted additively by five factors, (i) an age-suitability adjustment, awarding a bonus when the patient's age lies within the treatment's optimal window and a penalty when it falls outside the acceptable bounds; (ii) a budget-fit adjustment, strongly penalising treatments whose indicative cost exceeds the patient's stated budget and rewarding those well within it; (iii) medical-history modifiers for recorded comorbidities such as diabetes, smoking and bleeding disorders, including explicit contraindication penalties for implant and other surgical procedures; (iv) a success-rate bonus favouring treatments with higher documented success rates; and (v) a condition-severity weight reflecting the seriousness of the detected condition. A lightweight heuristic adjustment layer then refines the score using normalised features (age, gender, budget, treatment cost, success rate, risk level and the number of detected conditions), and the final score is clamped to the 0-100 range.

Because all adjustments are additive and bounded, conflicting clinical factors- for example, a clinically ideal treatment that exceeds the stated budget- lower a candidate's ranking rather than eliminate it, keeping alternatives visible to the practitioner. Treatments are sorted by final score and the top five are returned, each accompanied by a structured rationale, indicative cost range, expected duration, and success-rate estimate. The scoring logic is deterministic, so every ranking is fully traceable; the output is advisory, and the final treatment decision remains with the clinician.

System Architecture

The prototype is organised in three tiers (Figure 1). The presentation tier is a Next.js (TypeScript) single-page application that supports radiograph upload, patient-data entry, recommendation review and report download. The interface uses Tailwind CSS and is responsive across desktop, tablet and mobile form factors. The application tier

is a FastAPI service that exposes REST (Representational State Transfer) endpoints for radiograph inference (invoking the YOLOv8-Nano model), patient management and recommendation generation. A separate Flask microservice handles PDF report generation using ReportLab, isolating heavy rendering work from the main inference pipeline. The data is MongoDB and stores patient records, detection results and generated recommendations. As the prototype is intended for use exclusively by dental technicians and dentists in a single-tenant environment, no end-user authentication is implemented, and access control is deferred to deployment-environment safeguards.




Figure 1. Three-tier architecture of the prototype (presentation, application and data tiers, with the reporting microservice).

Implementation environment

Development used Python 3.11, FastAPI 0.110, Flask 3.0, Ultralytics 8.3, Next.js 14, MongoDB 6.0 and PyTorch 2.x. The trained model checkpoint (best.pt) is bundled with the backend service. Inference runs on CPU or GPU, and no specialised hardware is required.

Results

Detection performance

The fine-tuned YOLOv8-Nano detector was evaluated on the dataset's validation split. Aggregate performance across the ten dental classes is summarised in Table 1.

Table 1. Validation-set detection performance of the fine-tuned YOLOv8-Nano detector.

Metric	Value
Precision	0.60
Recall	0.56
mAP@0.5	0.55
mAP@0.5:0.95	0.28

These figures reflect the performance of the model after 20 epochs of fine-tuning from a COCO-pre-trained backbone. They should be interpreted as prototype-stage rather than production-stage results, and deeper training schedules and larger backbones are expected to improve them further.

Qualitatively, the detector localised larger and visually distinctive structures, such as crowns, dental implants and root canal treatments, more reliably than fine pathologies such as small caries lesions and periapical lesions, which are more sensitive to the confidence threshold. Figure 2 shows a representative inference output in which the detector has placed bounding boxes around several visible features of a validation-set panoramic radiograph.




Figure 2. Representative annotated radiograph showing YOLOv8-Nano detection of dental conditions on a validation-set panoramic image.

End-to-end pipeline behaviour

The end-to-end pipeline ran without modification on commodity hardware. The radiograph is uploaded through the Next.js client, the FastAPI service runs inference and returns the detected conditions, and the recommendation engine scores candidates against the nine-treatment knowledge base. The frontend presents the top-five ranked recommendations alongside the patient-data form and an annotated copy of the radiograph. On request, the Flask reporting service compiles a multi-section PDF containing the patient details, an annotated radiograph thumbnail, the detected conditions and the ranked recommendations with rationale and indicative cost ranges.

In informal local benchmarks on a mid-range laptop with CPU-only inference, per-radiograph processing (including detection, scoring and rendering) completed within a few seconds, and PDF generation added on the order of one second.

Recommendation output

For a representative test radiograph containing detected “Caries” and “Missing Teeth” labels, the engine returned (in descending suitability order) Dental Filling, Composite Bonding, Crown, Dental Implant and Bridge. Each item in the output carried an explicit rationale string linking the detected condition(s) to the chosen treatment, an indicative cost range, an expected duration and a documented success-rate estimate sourced from the knowledge base.

Discussion

The prototype described in this paper differs from the current generation of dental AI tools in three respects. First, where systems such as Overjet [9], Denti.AI [10], WeDiagnostiX [11] and Pearl AI [12] focus on detection and annotation, the present system goes further and produces an explicit ranked treatment recommendation. Second, the recommendation logic combines radiographic findings with patient-specific demographic, medical and budget factors, addressing the multifactorial nature of prosthetic decision-making that purely image-driven systems cannot capture. Third, the rule-based scoring engine makes every recommendation explainable, since the rationale for any ranking can be inspected directly without looking inside opaque model internals.

In the published literature, the closest contributions are the review-style papers of Aldahabani [13] and Najeehs [14], which discuss the future of AI in prosthetic and restorative dentistry but do not present an integrated end-to-end working system. Systematic reviews of AI in prosthodontics and implant dentistry similarly report that published models concentrate on laboratory, design and diagnostic tasks, such as automated restoration design, implant-type recognition and implant-success prediction, rather than on integrated, patient-specific treatment recommendation [16,17]. To the authors' knowledge, no published system yet combines panoramic-radiograph condition detection with demographic, medical and financial patient factors to produce ranked prosthetic recommendations, which precludes a direct quantitative comparison and underlines the gap the present prototype addresses as a concrete, inspectable reference implementation.

The detection performance achieved, mAP@0.5 of 0.55 across ten classes, is moderate; although it improves substantially on the initial short-schedule prototype, it remains below the level required for clinical use. Three factors plausibly explain the remaining gap. First, training still used a relatively short 20-epoch schedule from a COCO-pre-trained initialisation, and deeper schedules and larger backbones are expected to lift the mAP further. Second, the training data are community-contributed via the Roboflow Universe platform and have not been validated against a clinical gold-standard annotation protocol; annotation quality, class balance and the demographic and equipment diversity of the collection are not documented, so annotation noise and dataset bias cannot be excluded, and clinically validated datasets remain necessary before any pre-deployment evaluation. Third, the reported metrics are aggregated across the ten classes and are sensitive to class imbalance, and smaller or visually subtle classes (ex-periapical lesions) typically achieve lower per-class performance than larger structures.

The moderate detection performance has direct implications for the downstream recommendations and for patient safety. A missed detection (false negative) can remove an appropriate treatment from the candidate list entirely, while a false-positive detection can introduce candidates that are not clinically indicated; in both cases the ranking presented to the practitioner is degraded. Three design decisions mitigate this risk at the prototype stage. The annotated radiograph is displayed alongside the recommendations so the practitioner can visually verify every detection; the output is explicitly advisory, with the practitioner retaining full decision-making authority; and each recommendation carries a traceable rationale naming the detected condition it responds to, so an implausible detection is immediately visible rather than hidden inside an aggregate score. Nevertheless, substantially higher detection accuracy is a precondition for any clinical deployment.

Ethical, regulatory and medico-legal considerations also warrant discussion. The system is a clinical decision-support tool, not an autonomous diagnostic device, as clinical accountability for the selected treatment remains entirely with the treating clinician. Deployment beyond research would require conformity with the medical-device regulations applicable to software as a medical device (SaMD) in the target jurisdiction, together with documented risk management and transparent reporting of intended use and performance. The deterministic, rule-based design supports AI transparency, and the persistence of every recommendation with its rationale provides an audit trail relevant to medico-legal review. Patient-safety governance, informed consent regarding AI involvement and a clear delineation of responsibility between developer and practitioner would need to be established before clinical use.

Several limitations should be acknowledged. The evaluation reported here uses the dataset's published validation split rather than a held-out test set; a strict train/validation/test partition is a clear next step. No clinical validation has been performed, and the recommendation engine encodes deterministic heuristics whose cost ranges, age windows and severity weights have not yet been calibrated against published clinical guidelines or expert consensus. Likewise, the engine itself has been reviewed only qualitatively on sample cases; no quantitative measure of recommendation quality has yet been obtained. A planned evaluation will measure top-k recommendation accuracy against reference plans produced by practising dentists and prosthodontists, inter-clinician agreement and user satisfaction using standardised usability instruments. The system also lacks user authentication, acceptable for a controlled single-tenant deployment but required before any multi-tenant clinical use. Finally, the engine exposes a stable scoring interface deliberately designed so that the deterministic logic can later be replaced by a learned ranking component.

Regarding generalisability, the detection model was trained on Roboflow-distributed panoramic radiographs and performance on radiographs from substantially different imaging protocols, sensors or populations cannot be guaranteed without further validation. The recommendation knowledge base likewise encodes Sri Lankan Rupee cost ranges, and beyond costs, treatment availability, clinical conventions and patient expectations differ between healthcare systems; adapting the system to another setting would therefore require local clinicians to review its treatments, severity weights, cost data and contraindication rules. Despite these limitations, the prototype demonstrates the feasibility of an integrated, explainable, web-based pipeline that links radiograph analysis with

patient-specific data to deliver actionable treatment recommendations. Recent reviews [4,8] anticipate this class of integrated decision-support tool as a near-term direction for clinical AI in dentistry, and the present work contributes a working artefact in that direction.

Conclusions

This paper presented an integrated, web-deployable decision-support prototype for dental prosthetic treatment planning. The system combines a fine-tuned YOLOv8-Nano detector with a hybrid rule-based and heuristic scoring engine over a nine-treatment knowledge base and produces ranked, explainable recommendations together with a downloadable clinical-handover report. Detection performance on the validation split was moderate (mAP@0.5 0.55 across ten dental classes), reflecting the prototype-stage training schedule. The principal contribution is the end-to-end integration and the explicit, traceable mapping from detected conditions to ranked treatments through patient-specific factors.

Future work will proceed in three directions. Technical hardening will expand the training schedule, evaluate larger backbones and introduce a strict train/validation/test partition on clinically validated data. Clinical evaluation will assess recommendation accuracy and appropriateness with practising dentists and prosthodontists, including top-k agreement studies and usability testing measuring perceived usefulness, trust and workflow integration. System evolution will calibrate the knowledge base against published clinical guidelines, compare the deterministic engine with learned ranking approaches, including supervised and reinforcement-learning models trained on expert-labelled treatment decisions, and integrate the platform with electronic health record (EHR) systems and longitudinal patient records.

Acknowledgements

The authors thank ESOF Metro Campus Galle for facilitating the research environment and the Roboflow community for the openly licensed Dental X-Ray Panoramic dataset used during model training and validation.

References

1. Schwendicke F. Artificial intelligence in dentistry: a comprehensive overview. *J Dent Res.* 2020;99(7):769-74.
2. Sivari E. Deep learning in diagnosis of dental anomalies and diseases. *Diagnostics (Basel).* 2022;13(15):2512.
3. Khan MA. A comprehensive artificial intelligence framework for dental diagnosis and charting. *J Clin Med.* 2022;14(14):1-23.
4. Arsiwala-Scheppach LT. Machine learning in dentistry: a scoping review. *J Clin Med.* 2023;12(3):937.
5. Chen H. A deep learning approach to automatic teeth detection and numbering based on object detection in dental periapical films. *Sci Rep.* 2019;9(1):3840.
6. Londono J. Evaluation of deep learning and convolutional neural network algorithms accuracy for detecting and predicting anatomical landmarks on 2D lateral cephalometric images. *Saudi Dent J.* 2023;35:487-97.
7. Fonseka H. Automatic dental X-ray detection using a hybrid multi-convolution neural network. *J Med Imaging.* 2020;12(3).
8. Shekarapp C. Artificial intelligence-powered dentistry: enhancing patient care and efficiency. *Dent Technol J.* 2023;16(1):15-9.
9. Overjet. AI for dental X-ray analysis: complete guide [Internet]. 2025 [cited 2025 Jul 20]. Available from: <https://www.overjet.com/>
10. Denti.AI. AI-powered dental scribe, perio charting and X-ray diagnostics [Internet]. 2024 [cited 2025 Jul 20]. Available from: <https://www.denti.ai/>
11. WeDiagnostiX. AI-powered dental X-ray report [Internet]. 2024 [cited 2025 Jul 20]. Available from: <https://wediagnostix.ai/en/automated-dental-x-ray-analysis-with-ai>
12. Pearl AI. The growth of AI in dental radiology [Internet]. 2025 [cited 2025 Jul 20]. Available from: <https://hellopearl.com/>
13. Aldahabani S. The future of dental prosthetics: artificial intelligence and machine learning applications. *Afr J Biomed Res.* 2024;27(3):1063-70.

14. Najeeds M. Artificial intelligence (AI) in restorative dentistry: current trends and future prospects. J Dent Res. 2025;25(1):392.
15. Joda T. Health data in dentistry: an attempt to master the digital challenge. Public Health Genomics. 2019;22(1-2):1-7.
16. Bernauer SA. The use and performance of artificial intelligence in prosthodontics: a systematic review. Sensors (Basel). 2021;21(19):6628.
17. Revilla-León M. Artificial intelligence applications in implant dentistry: a systematic review. J Prosthet Dent. 2023;129(2):293-300.