

ETHICAL CONSIDERATIONS AND GOVERNANCE REGARDING THE USE OF ARTIFICIAL INTELLIGENCE IN DENTISTRY

*Edson Laranjeiras and ** (MAS) - 2026-09-02 at 15:19

Abstract

Artificial intelligence (AI) is increasingly used in dentistry for image interpretation, disease detection, risk prediction, treatment planning, prognosis, and administrative decision-making. Although these applications may improve diagnostic consistency and clinical efficiency, their implementation raises substantial ethical and governance concerns. This paper examines the principal requirements for responsible AI use in dentistry, focusing on patient autonomy, informed consent, privacy, fairness, transparency, accountability, human oversight, and equitable access. A structured narrative synthesis was conducted using selected peer-reviewed reviews and international policy documents addressing AI in dentistry and health care. The evidence indicates that dental AI systems demonstrate promising performance, particularly in radiographic tasks, but are frequently developed using retrospective, single-centre datasets that may not represent the diversity of clinical populations and practice environments. Ethical governance should therefore extend across the entire AI lifecycle, including justification of clinical need, representative data collection, privacy-preserving development, external validation, subgroup analyses, prospective evaluation, cybersecurity, version control, and post-market monitoring. AI should augment rather than replace professional judgement. Clear allocation of responsibility among dentists, health institutions, developers, manufacturers, and regulators is essential. Future research should prioritize patient-centred outcomes, fairness, calibration, cost-effectiveness, and the effects of AI on professional reasoning and dentist–patient relationships.

Keywords

artificial intelligence; dentistry; ethics; clinical governance; patient safety

Introduction

Artificial intelligence encompasses computational methods capable of identifying patterns, generating predictions, and supporting decisions from large datasets. In dentistry, machine-learning and deep-learning systems have been investigated for caries detection, periodontal and periapical lesion identification, assessment of bone loss, cephalometric analysis, orthodontic planning, implant planning, oral-cancer screening, forensic identification, and prediction of treatment outcomes. These applications may improve diagnostic support and workflow efficiency, but their clinical value cannot be determined by technical accuracy alone.^{1,2}

The use of AI in dentistry affects fundamental ethical principles. Dental images and records contain sensitive health information; algorithmic errors may delay diagnosis or result in inappropriate treatment; and unequal representation in training datasets may produce clinically important disparities. Moreover, clinicians may over-rely on algorithmic recommendations, potentially weakening independent reasoning and obscuring professional responsibility. Responsible

** Dentist and C.A.I.O.. **Multi-Agent Systems: MÉO-Research, Igor-Editor, Little Dentist-Reviewer and Coordinator.*

implementation therefore requires governance mechanisms that protect patients while permitting clinically justified innovation.

This paper examines ethical considerations and governance requirements for AI in dentistry. It emphasizes that AI should remain a decision-support technology operating under meaningful human supervision and within established standards of clinical accountability. The Figure 1. Artificial Intelligence in Dentistry: Applications and Responsible Use. The figure summarizes major AI applications in dentistry, including image interpretation, disease detection, risk prediction, treatment planning, prognosis, and administrative decision-making. It also highlights key principles for responsible implementation, including patient autonomy, informed consent, privacy, fairness, transparency, accountability, human oversight, and equitable access.



Figure 1. Artificial Intelligence in Dentistry: Applications and Responsible Use.

Methods

A structured narrative synthesis was undertaken using the six sources identified in the research materials. The evidence base comprised 2 peer-reviewed publications addressing AI applications and challenges in dentistry, 2 World Health Organization policy documents concerning ethics and regulation of AI for health, 1 methodological guideline for AI-related clinical trials, and 1 European Union legislative instrument. The sources were analyzed thematically according to clinical performance, autonomy, privacy, fairness, transparency, accountability, human oversight, justice, and lifecycle governance.

Because the available evidence was heterogeneous and included reviews, policy documents, reporting guidance, and legislation, no quantitative meta-analysis

was attempted. The synthesis focused on principles applicable to dental practice, research, health institutions, and developers.

A multi-agent system (MAS) was used to support the research, synthesis, editing, and review process. The collaborative AI agents included MÉO-Research, responsible for research support and evidence organization; Igor-Editor, responsible for editorial refinement and scientific consistency; Little Dentist-Reviewer, responsible for reviewing the content from a dental perspective; and Coordinator, responsible for coordinating the multi-agent workflow. AI-generated contributions were treated as supportive rather than autonomous scholarly outputs. The human author retained responsibility for the manuscript and performed the final review and approval before publication.

Results

Clinical promise and evidentiary limitations

AI has demonstrated promising performance in several dental diagnostic and planning tasks, particularly those involving radiographic image classification and landmark detection.^{1,2} However, reported performance is often derived from retrospective datasets obtained from a single institution or imaging system. Such studies may not adequately represent differences in patient age, ethnicity, socioeconomic circumstances, dentition, disease prevalence, image quality, or clinical workflow.

Consequently, high internal accuracy does not guarantee reliable performance in another population or practice environment. External validation, calibration assessment, and prospective evaluation are necessary to establish whether an AI system improves clinical outcomes rather than merely reproducing labels in a research dataset. Clinically meaningful evaluation should compare AI-assisted care with usual care and assess diagnostic errors, treatment decisions, patient outcomes, and resource use.

Autonomy, consent, and privacy

Patient autonomy requires understandable communication about the role of AI when it materially contributes to diagnosis, treatment planning, or risk assessment. Patients should be informed that an algorithm may assist the dentist, what type of information it uses, its known limitations, and that the treating professional remains responsible for interpreting the result. Consent is especially important when dental images or records are collected for model development or reused for secondary purposes.

Dental data governance should include lawful processing, data minimization, secure storage, role-based access, cybersecurity safeguards, retention policies, and procedures for data sharing. The World Health Organization identifies protection of autonomy, privacy, and confidentiality as central elements of ethical

AI governance in health care.³ De-identification reduces direct identifiers but does not eliminate the possibility of re-identification, particularly when datasets are combined with other information.

Bias, fairness, and justice

AI systems may produce unequal error rates when training data underrepresent particular demographic or clinical groups. Relevant sources of variation include age, ethnicity, socioeconomic status, disability, dentition, disease severity, imaging device, and image acquisition conditions. Overall accuracy can conceal poor performance in smaller or clinically vulnerable subgroups. Fairness assessment should therefore report subgroup-specific sensitivity, specificity, predictive values, calibration, and clinically consequential error patterns.

Justice also concerns access. AI systems requiring expensive software, high-quality imaging equipment, or specialized technical support may benefit well-resourced urban practices while widening disparities in rural or low-resource settings. Governance should consider affordability, infrastructure, workforce effects, environmental costs, and whether the technology addresses a genuine clinical need rather than merely increasing automation. The Figure 2. Key Challenges and Opportunities of Artificial Intelligence in Dentistry. The figure highlights four key dimensions for responsible AI implementation in dentistry: clinical promise and evidentiary limitations; bias, fairness, and justice; transparency, explainability, and accountability; and lifecycle governance.




Figure 2. Key Challenges and Opportunities of Artificial Intelligence in Dentistry.

Transparency, explainability, and accountability

Clinicians require clear information about an AI system's intended use, training population, input requirements, validation results, uncertainty, known failure modes, and update history. Explainability may help clinicians identify implausible outputs and communicate with patients, but an explanation should not be regarded as proof that a prediction is correct. A technically interpretable output can still be based on biased data or an inappropriate clinical question.

Accountability must be distributed explicitly among dentists, institutions, developers, manufacturers, and regulators. Dentists should verify AI outputs, integrate them with clinical findings, document the reasoning behind decisions, and retain the ability to override algorithmic recommendations. Institutions should establish procurement standards, training requirements, incident-reporting pathways, audit procedures, and rules for software updates. Developers and manufacturers should communicate limitations, maintain cybersecurity, investigate failures, and support post-market monitoring.

Lifecycle governance

International guidance recommends governance across the complete AI lifecycle rather than only at the point of deployment.^{3,4} Relevant stages include defining a justified clinical purpose, selecting representative data, protecting privacy, conducting independent external validation, evaluating subgroup performance, and performing prospective clinical studies. After implementation, systems require monitoring for performance drift, changes in disease prevalence, altered imaging technology, cybersecurity threats, and unintended effects on clinical practice.

Reporting frameworks can improve transparency in research involving AI. CONSORT-AI provides guidance for clinical trial reports involving AI interventions and emphasizes description of the algorithm, its role in the intervention, human–AI interaction, and deviations from intended use.⁵ Related frameworks, including SPIRIT-AI, DECIDE-AI, and TRIPOD+AI, support protocol development, early clinical evaluation, and prediction-model reporting.

Regulatory approaches are also developing. In the European Union, the Artificial Intelligence Act establishes obligations for AI systems according to risk, while medical-device regulation remains relevant to systems intended for clinical use.⁶ Regulation provides an important baseline but cannot replace professional judgement, institutional governance, or clinical accountability.

Discussion

The ethical implementation of AI in dentistry depends on recognizing that algorithmic performance is context-dependent. A system that performs well under controlled research conditions may fail when applied to different populations, imaging devices, disease prevalences, or workflows. Accordingly, validation

should be viewed as an ongoing process rather than a one-time technical certification.

Meaningful human oversight is particularly important because automation bias may cause clinicians to accept algorithmic recommendations without sufficient scrutiny. Conversely, excessive distrust may prevent clinicians from using a valid and beneficial tool. Appropriate education should therefore include AI literacy, interpretation of uncertainty, recognition of common failure modes, documentation standards, and procedures for overriding or reporting unsafe outputs. AI should support, not replace, the dentist–patient relationship and the professional obligation to consider individual circumstances.

Transparency also has practical limits. Complex models may not provide explanations that are fully understandable or clinically reliable. Governance should consequently require more than explainability. Independent validation, subgroup analysis, auditability, documented limitations, and evidence of patient benefit are equally important. Patients should not be presented with a misleading impression that an algorithm is objective merely because it is computational.

Several research gaps remain. Evidence is limited regarding prospective patient outcomes, long-term safety, cost-effectiveness, calibration drift, minority-population performance, patient preferences, and the effect of AI on clinical reasoning. Dental research would benefit from multicentre and internationally diverse datasets, standardized outcome definitions, independent replication, and comparisons between AI-assisted and usual care. Studies should report subgroup errors and patient-centred outcomes, including trust, understanding, satisfaction, treatment adherence, and perceived involvement in decision-making.

Conclusion

AI has substantial potential to support diagnosis, treatment planning, prevention, and workflow management in dentistry. Its ethical acceptability, however, depends on governance that protects autonomy, privacy, fairness, safety, and professional accountability. Dental AI systems should be developed from representative data, independently validated, prospectively evaluated, transparently reported, and continuously monitored after deployment. Patients should be informed when AI materially influences their care, and dentists should retain meaningful authority to interpret, question, and override algorithmic recommendations. Effective governance requires coordinated responsibility among clinicians, institutions, developers, manufacturers, researchers, patients, and regulators. AI should therefore be implemented as a clinically governed adjunct to dental expertise rather than as an autonomous substitute for professional judgement.

REFERENCES

1. Schwendicke F, Samek W, Krois J. Artificial intelligence in dentistry: chances and challenges. *J Dent Res*. 2020;99(7):769-774.
2. Khanagar SB, Al-Ehaideb A, Maganur PC, et al. Developments, application, and performance of artificial intelligence in dentistry—a systematic review. *J Dent Sci*. 2021;16(1):508-522.
3. World Health Organization. Ethics and Governance of Artificial Intelligence for Health. World Health Organization; 2021.
4. World Health Organization. Regulatory Considerations on Artificial Intelligence for Health. World Health Organization; 2023.
5. Liu X, Rivera SC, Moher D, Calvert MJ, Denniston AK; SPIRIT-AI and CONSORT-AI Working Group. Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: the CONSORT-AI extension. *Nat Med*. 2020;26:1364-1374.
6. European Union. Regulation (EU) 2024/1689: Artificial Intelligence Act. Official Journal of the European Union. Published 2024.