GRADUATION PROJECT

Degree in dentistry

Fraud detection in dentistry using artificial intelligence

Madrid, Academic year: 2025/2026

Identification number: 64

Resumen

Introduccion: La odontologia ha cambiado mucho en los ultimos anos con la llegada de la inteligencia artificial, los escaneres intraorales y el software en la nube. Estas herramientas aportan ventajas claras, pero tambien facilitan nuevas formas de fraude, manipulacion de registros y engano a los pacientes. **Objetivos:** Este trabajo analiza las distintas formas de fraude en la odontologia digital, desde radiografias alteradas y marketing enganoso hasta algoritmos sesgados y brechas de seguridad, y explora como la IA puede servir tanto para detectar como para generar estos problemas. **Metodos:** Se realizo una revision sistematizada utilizando PubMed, Google Scholar y revistas dentales especializadas, incluyendo unicamente fuentes publicadas entre 2015 y 2026. Se seleccionaron 67 fuentes tras el cribado. **Resultados:** Los dentistas apenas distinguen las radiografias reales de las falsas (56% de acierto), las herramientas de IA marcan dientes sanos como enfermos en hasta el 85% de los casos en algunos estudios, y los principales ciberataques han expuesto los registros de casi 90 millones de pacientes dentales. **Conclusiones:** Las herramientas digitales han llegado para quedarse, pero sin normas adecuadas, marketing honesto y responsabilidad real, el riesgo de fraude solo crecera. La profesion necesita estandares de autenticacion de imagenes, pruebas independientes de productos de IA, mejor formacion para los graduados y una inversion seria en seguridad de datos.

Palabras clave:

Odontologia; Inteligencia artificial; Fraude digital; Etica; Radiologia.

ABSTRACT

Introduction: Dentistry has undergone considerable change in recent years. Artificial intelligence, intraoral scanners, and cloud-based software have become part of routine clinical work in practices worldwide. These tools offer clear advantages in terms of speed and precision. However, they also introduce vulnerabilities that remain largely under-discussed, particularly regarding fraud, record manipulation, and patient deception. **Objectives:** This work looks at the different ways fraud shows up in digital dentistry, from altered X-rays and exaggerated marketing to biased algorithms and data breaches. It also explores how AI can be used to catch these problems, not just cause them. **Methods:** A systematised review was carried out using PubMed, Google Scholar, and specialised dental journals. Only sources published between 2015 and 2026 were included, covering topics like digital image manipulation, AI ethics, automation bias, cybersecurity, orthodontic AI accuracy, and regulatory frameworks. A total of 67 sources were selected after screening. **Results:** The findings show that dentists can barely tell the difference between real and fake X-rays, that AI tools flag healthy teeth as diseased in up to 85% of cases in some studies, and that major cybersecurity attacks have exposed the records of nearly 90 million dental patients. The SmileDirectClub collapse left over two million people midtreatment with no support. Commercial AI platforms used in thousands of clinics have almost no independent studies proving they work in real life. **Conclusions:** Digital tools are here to stay, and that is a good thing. But without proper rules, honest marketing, and real accountability, the risk of fraud will only grow. Dentistry needs authentication standards for digital images, independent testing of AI products, better training for new graduates, and serious investment in data security. The profession that gets this right will earn and keep its patients' trust. The one that ignores it will not.

Keywords:

Dentistry; Artificial intelligence; Digital fraud; Ethics; Radiology.

TABLE OF CONTENTS

1. INTRODUCTION 1

1.1 Definition and background 1

1.2 Scientific background and current context 2

1.3 Problem statement 3

1.4 Justification and purpose of this review 4

1.5 Algorithmic bias: not everyone gets the same quality of care 5

2. OBJECTIVE 7

3. MATERIAL AND METHODS 8

4. RESULTS 10

5. DISCUSSION 10

6. CONCLUSION 24

7. SUSTAINABILITY 26

8. REFERENCES 27

9. ANNEXES 46

1. INTRODUCTION

1.1 Definition and background

1.1.1 Digital transformation in dentistry

If you walked into a dental clinic ten years ago, the experience was already starting to change. Intraoral scanners were replacing the messy impression trays, digital X-rays were taking over from film, and CAD/CAM machines could mill a crown while the patient waited. Fast forward to today, and the transformation is nearly complete. AI-powered software now reads radiographs, plans implant placements, designs clear aligners, and even predicts which patients are most likely to develop cavities (1,2).

Companies like Invisalign, Angel Aligner, and Spark depend heavily on prediction software that simulates how teeth will move over the course of treatment. In radiology, platforms like Pearl, Overjet, and VideaHealth promise to spot caries, bone loss, and periapical pathology with accuracy that matches or beats the average clinician (3,4). These are not small-scale experiments. Pearl alone reports that its AI processes millions of dental images across thousands of clinics (5). Overjet has partnered with major insurance carriers to review claims using computer vision. This technology is already shaping clinical decisions for millions of patients.

On the positive side, digital workflows save time, cut down on errors, and make it easier to share information between providers. A digital impression takes minutes instead of the ten or fifteen minutes needed for a traditional alginate mould. A panoramic X-ray is available instantly instead of waiting for film to develop. A CAD/CAM crown can be milled chairside in a single appointment instead of requiring a temporary and a second visit. For patients, this means less time in the chair, fewer appointments, and a generally smoother experience.

1.1.2 Risks associated with digitalisation

Benefits always come with trade-offs, and in this case, the trade-off is vulnerability. Those same tools that make diagnosis faster also make manipulation harder to detect. Cloud platforms that that let clinics access records from anywhere also give hackers a target worth attacking. And the same AI algorithms that flag disease can also flag healthy tissue, leading to treatments patients did not need. In an era where a dental X-ray is just a digital file that can be opened and edited in any photo programme, the question is no

longer whether fraud can happen digitally. The question is how much of it is already happening and nobody knows about it.

1.2 Scientific background and current context

1.2.1 Digital radiography and Ai limitations

Digital radiography is now standard in almost every dental practice. The advantages are clear: less radiation, instant images, easy storage, simple sharing (6).But these same features also make tampering easy. A dental X-ray saved as a JPEG or TIFF file can be opened in any photo editing programme and modified in ways that are nearly impossible to detect. A 2017 study asked 66 dentists to look at 20 dental radiographs and tell which ones had been digitally altered. They got it right only 56% of the time, barely better than flipping a coin (7).

AI adds new layers of concern. Many products are marketed as intelligent without disclosing what algorithm they use, what data they were trained on, or how they reach their conclusions (3,8). This is what researchers call the black-box problem: the AI gives you an answer, but nobody can explain why (9,10).If a dentist tells a patient they need a crown because the computer flagged something on their X-ray, but neither the dentist nor the patient can understand how the computer made that decision, the whole idea of informed consent falls apart.

1.2.2 Clinical and regulatory implications of Ai

Consider how this plays out in a typical clinic. A patient comes in for a routine check-up. The dentist takes a set of bitewing X-rays and the AI immediately highlights three areas in yellow, suggesting possible early caries. The patient sees the highlights on the screen and asks what those yellow spots mean. The dentist says the computer detected some early decay. The patient, trusting both the dentist and the impressive-looking technology, agrees to treatment. But what if two of those three spots were false positives, areas of normal anatomy that the AI mistook for disease? The patient just consented to unnecessary treatment based on a recommendation that nobody in the room fully understands. This is not a rare scenario. As we will see in the results, false positive rates for AI caries detection can reach disturbingly high levels in certain clinical contexts.

Orthodontics presents an even more complex situation. When a patient is shown a ClinCheck simulation of how their teeth will look after Invisalign treatment, they naturally assume that the result on screen is what they will actually get. But the clinical data shows that these predictions achieve as little as 18% accuracy for certain types of tooth movement (11).That gap between what AI promises and what it delivers is not small. It is enormous. And patients are making financial and health decisions based on these predictions without knowing how unreliable they can be.

Regulation has not caught up. While the European Union passed the AI Act in 2024, which classifies medical AI as high-risk and requires transparency and human oversight (12,13),most dental AI products still reach the market through the FDA's 510(k) pathway, which only asks manufacturers to show their product is similar to something already approved, not that it actually works better than a human clinician in a real clinical setting (4,14).A 2026 analysis of 52 FDA-cleared dental AI devices found that 60% did not even disclose their algorithm type, and half had no public documentation showing they had ever been tested in a real clinic (14).

1.3 Problem statement

1.3.1 Nature and scope of dental fraude

Dental fraud is not a minor issue affecting a handful of bad actors. It takes many forms: billing for procedures that were never done, performing treatments that patients did not need, altering clinical records to justify unnecessary work, and misrepresenting the capabilities of digital tools. What ties all of these together is that patients end up receiving care they did not need, or not receiving care they did, while trust in the profession slowly erodes (5,15).In the United States alone, federal enforcement actions against dental fraud have increased in recent years, with multiple large-scale schemes uncovered involving Medicaid billing, fake practices, and falsified
records (16,17).

1.3.2 Evidence from real-world cases and digital implifications

Federal prosecutions in the United States give us real examples. The company behind the Kool Smiles paediatric chain settled after performing unnecessary pulpotomies, extractions, and stainless steel crowns on children across more than 130 clinics in 17

states. Whistleblowers described a corporate culture that punished dentists who did not produce enough procedures and rewarded those who did the most (17). In these cases, the harm was not just financial. Children underwent invasive treatments they did not need, under general anaesthesia in some cases, with all the risks that come with it.

Perhaps the most disturbing recent case is that of Dr. Scott Charmoli from Wisconsin. Charmoli deliberately damaged his patients' teeth with his drill, then took photos and X-rays of the damage and submitted them to insurance as pre-existing conditions that needed crowns. He did over 1,000 crowns per year, far above the average for his state. He was convicted on five counts of healthcare fraud and sentenced to prison (18,19).

What makes the Charmoli case relevant to this review is that digital records played a central role. Evidence against him came partly from comparing before-and-after images that showed damage appearing only after his appointments. In an era of digital radiography, a more careful fraudster might have manipulated the before images to show pre-existing damage, making the scheme much harder to detect. Tools to do this are freely available and, as the research shows, the results would likely fool both clinicians and insurers (7).

These are not isolated stories. They point to a broader problem: when clinical records go digital, fraud becomes harder to catch, easier to scale, and potentially more harmful to patients. And as AI tools become more common in dental practice, they add new dimensions to this risk, from over-diagnosis driven by false positives to deceptive marketing that exaggerates what the technology can actually do (3,8).

1.4 Justification and purpose of this review

Given everything described above, it is fair to say that dentistry is at a crossroads. The technology is powerful and getting more powerful every year. But the safeguards, the rules, the training, and the ethical frameworks have not kept pace. Many dentists do not know how easy it is to alter a digital X-ray, or how unreliable some AI tools really are, or how exposed their patient data might be to cyberattack (8,20).

This review aims to pull together the evidence on all these fronts. Not to argue against technology, because that ship has sailed, but to map out where the risks are, what we know about them, and what the profession needs to do about them. The goal is to give

dental professionals, students, and regulators a clear picture of the current state of AI-related fraud in dentistry, grounded in real cases, real data, and real clinical consequences.

1.5 Algorithmic bias: not everyone gets the same quality of care

1.5.1 Origin and mechanisms of algorithm bias

There is another dimension to AI fraud that does not always get the attention it deserves: bias. AI systems learn from data, and if that data mostly comes from one type of population, the system will work well for that population and poorly for everyone else. In dentistry, this is not a distant concern. Allareddy et al. showed that orthodontic and craniofacial AI training datasets grossly underrepresent ethnic and racial minorities, leading to models that perform poorly on non-white populations (21).If an AI system was trained primarily on X-rays from affluent urban clinics, it may misread images from patients with different dental anatomy, different restoration materials, or different disease patterns.

1.5.2 Clinical transparency and lack of transparency

This kind of bias does not look like fraud in the traditional sense, but its effects can be just as harmful. A patient from an underserved community who receives an incorrect AI-assisted diagnosis, whether a false positive that leads to unnecessary treatment or a false negative that misses actual disease, is experiencing the consequences of a system that was never designed to serve them well. As AI tools become more widespread, this disparity could widen rather than narrow the gaps in dental care quality that already exist between different populations (1,21).

The problem is made worse by the fact that most commercial AI platforms do not publish demographic breakdowns of their training data or performance across different populations. Without this information, clinicians have no way to know whether the tool they are using works equally well for all their patients. Transparency about training data composition and performance across demographic groups should be a basic requirement for any AI system used in healthcare, but it is not yet standard practice in dental AI (3,8) (figure 1).

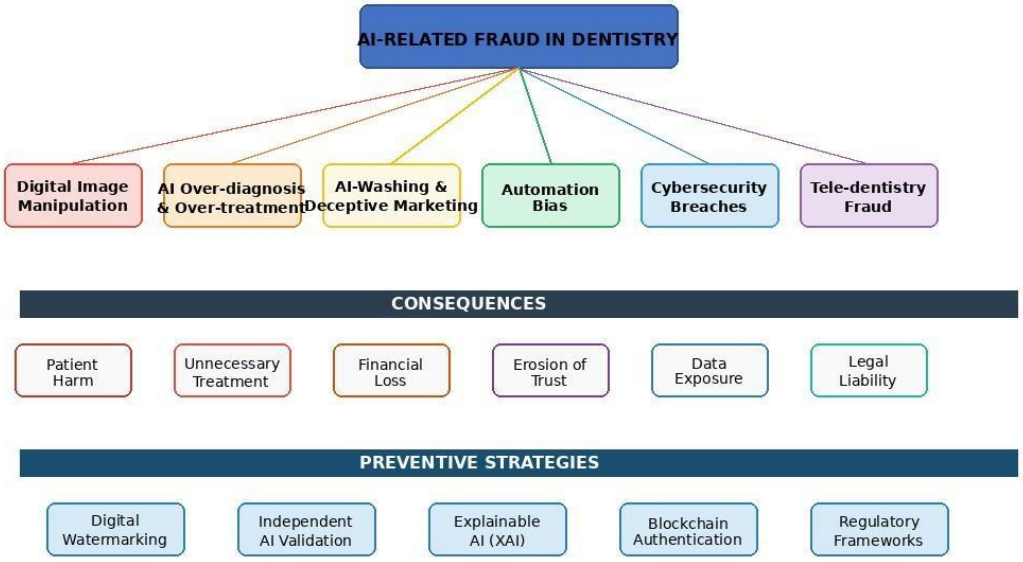


Figure 1. Conceptual map of AI-related fraud types, consequences, and preventive strategies in dentistry (3,8,22,23,24)

2. OBJECTIVES

To analyze and organize the different types of fraud and ethical problems that come from using AI and digital tools in dentistry, including image manipulation, misleading marketing, data misuse, algorithmic bias, and cybersecurity failures, and to suggest practical strategies to prevent them.

PICO: In the context of digital dental practice (P), do AI-based fraud detection tools and ethical frameworks (I), compared to traditional oversight mechanisms (C), improve the identification and prevention of digital fraud, algorithmic bias, and clinical misconduct (O)?

3. MATERIAL AND METHODS

This systematic review was guided by a PICO (population, Intervention, Comparaison, Outcomes) framework, in order to ensure a structured approach to data collection, the PICO question was as follow: In the context of digital dental practice (P), do AI-based fraud detection tools and ethical frameworks (I), compared to traditional oversight mechanisms (C), improve the identification and prevention of digital fraud, algorithmic bias, and clinical misconduct (O)?

A systematised literature review was carried out using PubMed and Google Scholar as the main databases, alongside manual searches in specialised journals such as the Journal of Dental Research, Journal of Dentistry, British Dental Journal, Dentomaxillofacial Radiology, the Journal of the American Dental Association, and the American Journal of Orthodontics and Dentofacial Orthopedics. Regulatory documents from the FDA, the European Commission, and professional bodies including the FDI World Dental Federation and the American Dental Association were also consulted. Legal and investigative sources from the U.S. Department of Justice, state attorney general offices, and the FDA MAUDE adverse event database were included for realworld case documentation.

The following search terms were used in various combinations, formatted according to standard PubMed and Google Scholar syntax: ("artificial intelligence"OR"AI" OR "deep learning" OR "machine learning") AND ("dentistry" OR "dental") AND ("fraud" OR "manipulation" OR "ethics" OR "bias" OR "cybersecurity" OR "over-diagnosis"); ("digital radiograph" OR "dental imaging") AND ("manipulation" OR "tampering" OR "forgery"); ("clear aligner" OR "Invisalign" OR "SmileDirectClub") AND ("accuracy" OR "prediction" OR "adverse events"); ("automation bias") AND ("clinical decision" OR "dental"); ("explainable AI" OR "XAI") AND ("dental diagnosis"); ("blockchain" OR "data security") AND ("dental records").

The initial search across all databases returned 847 results from academic databases and 42 additional records from regulatory and legal sources, for a total of 889 records. Before screening, 223 records were removed: 68 duplicates, 112 published outside the date range of 2015 to 2026, and 43 non-English publications. This left 666 records for screening.

Inclusion criteria required that sources be published between 2015 and 2026, written in English, and available in full text. Accepted study types included systematic reviews, meta-analyses, narrative reviews, prospective and retrospective clinical studies, scoping reviews, regulatory analyses, position papers from professional bodies, case reports, and documented legal proceedings. Exclusion criteria removed conference abstracts without full-text data, studies not relevant to at least one of the seven defined thematic areas, and sources with clear methodological weaknesses. A two-step screening process followed. First, title and abstract review was conducted to select papers relevant to at least one of the seven thematic areas: digital image manipulation and fraud detection, AI diagnostic accuracy and over-treatment, orthodontic AI prediction failures, cybersecurity in dental practice, automation bias and clinical decision-making, ethical and legal frameworks, and tele-dentistry fraud risks. This excluded 421 records, leaving 245 reports sought for retrieval. Of these, 38 could not be retrieved in full text, leaving 207 reports assessed for eligibility. Second, detailed full-text reading and quality assessment, during which 142 reports were excluded: 64 were not sufficiently relevant to AI fraud in dentistry, 38 had insufficient methodology, 22 presented duplicate findings already covered by other sources, and 18 had unavailable full text upon closer inspection. This process resulted in the final selection of 67 sources that form the basis of this review.

Table 1. Eligibility Criteria

Criteria	Inclusion	Exclusion
Publication Date	Studies published between 2015–2026	Articles published before 2015
Language	Articles written in English	Articles in other languages
Focus Area	Studies on AI in dentistry, including fraud, imaging, diagnosis, bias, cybersecurity, ethics, or tele-dentistry	Studies outside these topics
Study Type	Reviews, clinical studies, scoping papers, and official or legal reports	Abstracts without full text or low-quality studies
Target Population	Research involving human dental data or patients	Studies on non-human subjects
Relevance	Directly related to the research question	Not clearly related after review
Availability	Full-text available articles	Full text not accessible

4. RESULTS

The study selection process followed PRISMA 2020 guidelines. A total of 889 records were initially identified through database searches (n = 847) and additional sources (n = 42). After removing duplicates (n = 68), records outside the predefined date range (2015–2026) (n = 112), and non-English publications (n = 43), 666 records remained for screening. Following title and abstract review, 421 records were excluded, leaving 245 reports sought for retrieval. Of these, 38 could not be retrieved, resulting in 207 full-text articles assessed for eligibility. After detailed evaluation, 172 reports were excluded due to lack of relevance (n = 94), insufficient methodological quality (n = 38), duplicate findings (n = 22), or unavailable full text (n = 18). Ultimately, 35 studies were included in the final results synthesis.

PRISMA 2020 flow diagram for new systematic reviews which included searches of databases and registers only

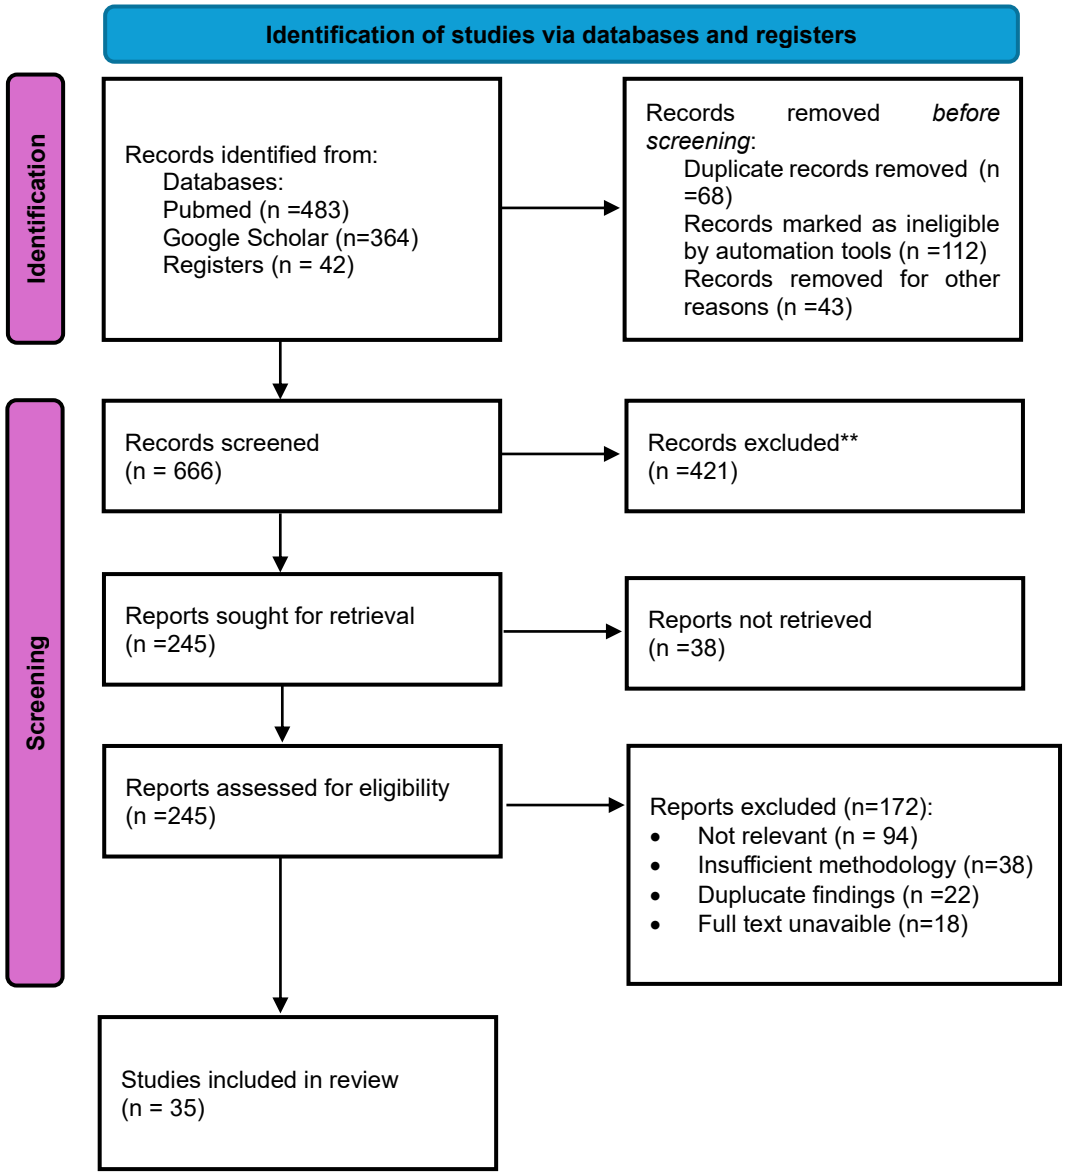


Figure 2. PRISMA 2020 flow diagram

Source: Page MJ, et al. BMJ 2021;372:n71. doi: 10.1136/bmj.n71.

This work is licensed under CC BY 4.0. To view a copy of this license, visit https://creativecommons.org/licenses/by/4.0/ 5.1 Digital Radiographic Manipulation and Detection

Table 2. Digital radiographic manipulation and detection (n=5)

Year / Author	Country	Study Type	Sample Size	Method	Control	Assessment Tool	Outcome
2017 DiazFlores-Garcia et al. (7)	Spain	Diagnostic accuracy study	66 dentists; 20 radiographs (10 original, 10 manipulated)	Adobe Photoshop manipulation of intraoral X-rays (added fillings, RCT)	Original unmodified radiographs	Blinded visual evaluation by practising dentists	Detection accuracy: 56% correct (6% above chance level); no difference by experience
2016 Bik et al. (25)	USA	Crosssectional analysis	20,621 papers from 40 biomedical journals	Forensic image screening for duplication / manipulation	None (observational)	Manual forensic image analysis + software tools	3.8% of papers had problematic figures; half appeared deliberate; rate increasing over time
2024 Amiri et al. (26)	Iran	Experimental study	Medical image datasets (exact n not specified)	AI-based copy-move forgery detection (SIFT + EOM + CHM)	Existing forgery detection algorithms	Accuracy, F1-score, precision, recall	Accuracy: 100%; F1score: 98.47%; superior to all comparison methods
2024 Abdalla et al. (27)	Egypt	Experimental study	JPEG image datasets	Error Level Analysis + CNN for manipulation detection	Existing detection frameworks	Training accuracy, validation metrics	Training accuracy: 99.05%; robust across image types

2026 Tordjman et al. (28)	France	Diagnostic accuracy study	17 radiologists from 12 centres in 6 countries	Detection of ChatGPT-4o generated radiographs vs real	Authentic radiographs	Blinded reader study; sensitivity and specificity	Unaware condition: 41% detected fakes; Warned condition: 70-75%; AIgenerated images fool trained radiologists

Table 3. AI Diagnostic Accuracy in Caries and Periapical Lesion Detection (n=6)

Year / Author	Country	Study Type	Sample Size	Method	Control	Assessment Tool	Outcome
2025 Arzani et al. (29)	Multiple	Umbrella review	14 systematic reviews; 137 primary studies	AI caries detection across multiple imaging modalities	Clinician diagnosis / histology	Pooled sensitivity, specificity, meta-analysis	Sensitivity: 0.85; Specificity: 0.90; ~1 in 5 positives are false
2025 Ammar et al. (30)	Multiple	Systematic review + meta-analysis	45 studies; 21 AI platforms	Commercial and research AI caries detection	Clinical / radiographic gold standard	PPV, sensitivity, specificity per platform	PPV range: 0.15-0.87; up to 85% false positives in lowprevalence groups
2024 Tural et al. (31)	Turkey	External validation study	991 carious lesions	AI caries detection on clinical photographs	Expert clinical diagnosis	Sensitivity, specificity, false negative rate	False negatives: 18.7%; Sensitivity for subtle lesions: 42.9%
2022 Mohammad-Rahimi et al. (32)	Multiple	Systematic review	Deep learning caries detection studies	Deep learning models (CNN architectures)	Histology / clinical reference	Accuracy; QUADAS-2 risk of bias	Accuracy: 7196%; most studies at high risk of bias

2024 Koenig et al. (33)	Multiple	Diagnostic accuracy study	Panoramic + CBCT images	AI periapical lesion detection	Expert radiologist consensus	Sensitivity, specificity per imaging type	Panoramic sensitivity: 30-33% (misses 2/3 of lesions); CBCT: significantly better
2025 Lal et al. (3)	Multiple	Critical review	44 studies; Pearl, Overjet, VideaHealth, others	Commercial dental AI platforms	Clinical performance claims vs evidence	Validation level (1-4 scale)	Most platforms at level 1-2; no independent real-world validation studies

Table 4. Orthodontic AI Prediction Accuracy (n=5)

Year / Author	Country	Study Type	Sample Size	Method	Control	Assessment Tool	Outcome
2020 Haouili et al. (11)	USA	Prospective clinical study	38 patients	ClinCheck prediction vs achieved tooth movement	ClinCheck predicted position (baseline)	3D model superimposition (ToothMeasure)	Overall accuracy: 50%; worst: maxillary incisor extrusion 18.3%; no improvement since 2009
2023 Ferlito et al. (34)	USA	Prospective study	50 patients	AI-based remote monitoring of clear aligner therapy	Repeated scans of same patients	GO/NO-GO consistency; tooth-level agreement	20% contradictory decisions; 0% agreement on which teeth had tracking problems
2023 Lombardo et al. (35)	Italy	Multicentre retrospective study	206 patients	ClinCheck prediction reliability	Pretreatment ClinCheck model	3D superimposition; arch length measurement	Arch length accuracy: 3038%; lower arch predictions least reliable
2015 Rossini et al. (36)	Multiple	Systematic review	Clear aligner clinical studies	Efficacy of controlled orthodontic movements	Predicted movement (ideal)	Measurement of achieved vs predicted movement	Extrusion accuracy: ~30%; rotation of rounded teeth: poorly controlled
2020 Wexler et al. (37)	USA	Crosssectional survey	470 DTC aligner users (SmileDirectClub)	Userreported experience questionnaire	None (descriptive survey)	Self-reported outcomes and side effects	6.6% had side effects requiring dentist visit; NDAs used to suppress negative reviews

Table 5. Cybersecurity Breaches Affecting Dental Practices (n=5)

Year	Entity	Country	Attack Type	Sample Size (Records)	Attack Method	Impact on Practice	Outcome
2023 (38)	MCNA Dental	USA	Ransomware attack	8.9 million patient records	LockBit group; 700GB exfiltrated	Data published on dark web after ransom refused	SSNs, X-rays, Medicaid IDs exposed; largest dental breach in 2023
2023 (39)	Delta Dental of California	USA	Supply chain vulnerability	6.9 million patient records	MOVEit file transfer exploit (Clop gang)	Names, SSNs, financial data stolen	Second largest dental breach; affected multiple state programmes
2019 (40)	DDS Safe / PerCSoft	USA	Ransomware (supply chain)	~400 dental practices	Evil/Sodinokibi ia backup rovider	Complete loss of records; practices shut down	Average 9.6 days downtime; backup service itself was compromised
2023-2024 (41)	Henry Schein	USA	Ransomware (triple attack)	35 TB data claimed stolen	BlackCat/ALPHV (3 separate attacks)	Ecommerce offline for weeks	10-15% customer loss; forced restructuring
2024 (42)	Westend Dental (4	USA 2)	Ransomware + cover-up	Patient records compromised	MedusaLocker; practice lied to investigators	No HIPAA training; servers in break rooms	Penalty: $350,000; no password policies existed

Table 6. Automation Bias in Dental Clinical Decision-Making (n=5)

Year / Author	Country	Study Type	Sample Size	Method	Control	Assessment Tool	Outcome
2022 Glick et al. (43)	USA	Randomised experimental study	41 dental students	AI with incorrect furcation suggestions	Students without AI assistance	Diagnostic agreement rate; statistical comparison (p-value)	Students with AI agreed with incorrect diagnosis significantly more ($p<0.05$)
2023 Dratsch et al. (44)	Germany	Diagnostic reader study	27 radiologists; 50 mammograms	AI BI-RADS suggestions (correct and incorrect)	Radiologists without AI	Sensitivity, specificity, BI-RADS accuracy	All experience levels affected by wrong AI suggestions; specificity decreased
2024 Saragih et al. (45)	Multiple	Experimental study	210 participants	AI decision support with known errors	Participants without AI support	False agreement rate; effect of training	Training reduced but did not eliminate bias; experience was not protective
2017 Cabitza et al. (46)	Italy	Narrative review	Healthcare AI literature	Analysis of unintended consequences of clinical AI	Pre-AI clinical practice	Literature synthesis	Identified deskilling, overconfidence, and complacency as predictable side effects
2024 FDI World Dental Federation (22)	International	Policy statement + white paper	Global dental profession guidelines	Expert consensus on AI risks in dentistry	None (policy recommendation)	Expert panel consensus	Named 3 risks: deskilling, never-skilling, mis-skilling; recommended mandatory AI literacy training

Table 7. Ethical and Legal Frameworks for AI in Dentistry (n=5)

Year / Author	Country	Study Type	Sample Size	Method	Control	Assessment Tool	Outcome
2021 Morch et al. (20)	Multiple	Scoping review	178 AI dental studies	Analysis of ethical reporting in dental AI research	None (descriptive)	Proportion reporting ethical concerns	Only 12.4% of studies reported any ethical concerns; major gap identified
2023 Rokhshad et al. (47)	Multiple (ITU/WHO)	Expert framework development	11 ethical principles; expert panel	Delphimethod checklist for ethical dental AI	None (consensus building)	Expert validation; checklist reliability	11-principle checklist developed; covers transparency, bias, privacy, accountability
2024 FDI Policy Statement (22)	International	Policy statement	Global dental profession	Expert consensus on AI governance	None (policy)	Consensus agreement	AI must be useful and true; built on high-quality representative data
2025 Mathew et al. (48)	UK	Legal analysis	Case law and liability frameworks	Analysis of liability when AI and clinician disagree	Traditional malpractice standards	Legal precedent analysis	Both following and overriding AI carry liability risks; no settled case law exists
2025 D'Cruz L. (49)	UK	Legal commentary	Bolam test and Montgomery ruling	Application of existing law to AIassisted dentistry	Pre-AI legal standards	Case law review	Failure to use available AI may become breach of duty; current law inadequate
2026 B.K. et al. (13)	Multiple	Systematic review	16 records; 5 ethical domains	Mapping ethical and legal challenges of dental AI	None (systematic analysis)	Thematic synthesis across 5 domains	Key gaps in: patient autonomy, privacy, bias, transparency, regulation

Table 8. Tele-dentistry Fraud Risks and Remote Monitoring (n=4)

Year / Author	Country	Study Type	Sample Size	Method	Control	Assessment Tool	Outcome
2023 Ferlito et al. (34)	USA	Prospective study	50 patients; repeated scans	AI remote monitoring of aligner therapy (Dental Monitoring)	Same patients rescanned	GO/NO-GO consistency; inter-scan agreement	20% contradictory GO/NO-GO decisions; 0% agreement on which teeth had problems
2021 Squires et al. (50)	USA	Evidencebased review	Teleorthodontics literature	Analysis of clinical and ethical risks of remote orthodontics	In-person orthodontic care	Literature synthesis; risk assessment	Identified risks: inadequate supervision, malpractice, missed pathology
2024 Malpe et al. (51)	India	Narrative review	Tele-dentistry applications	Analysis of fraud risk in remote dental consultations	Traditional in-person consultations	Literature review; risk mapping	Remote format increases fraudulent billing risk; service verification difficult
2019 Bhargava et al. (52)	India	Literature review	Tele-dentistry evolution	Analysis of legal and ethical aspects of remote dentistry	Conventional care delivery	Narrative synthesis	Digital radiographs can be altered to falsify procedure completion; legal gaps identified
2025 McGrath et al. (53)	Hong Kong	Ethics analysis	AI monitoring in vulnerable populations	Analysis of AI remote monitoring ethical concerns	Standard care for vulnerable groups	Ethical framework analysis	Concerns: consent validity in elderly/children; privacy risks; algorithmic bias in underserved groups

5. DISCUSSION

5.1 Reliability and integrity of digital dental radiographs

The evidence on this point is both clear and troubling. Diaz-Flores-Garcia et al. asked 66 practising dentists to look at 20 radiographs and identify which ones had been digitally altered. They scored 56%, which is almost identical to what you would get from random guessing (7). This is not an old problem either. A 2016 study by Bik et al. screened over 20,000 biomedical papers and found that 3.8% contained images that looked like they had been deliberately tampered with, and the rate was going up, not down (25).

What makes this especially dangerous in daily practice is that altered images do not just fool dentists. They fool insurance companies too. Researchers have shown that when you digitally add pathology to a dental radiograph and submit a claim, insurers approve it without hesitation. No red flags, no second opinions, no questions asked (7). Pearl's fraud detection AI has already caught real-world examples: one provider used the same toothless panoramic image to get denture approvals for almost 40 different patients. Another submitted the same set of fullmouth X-rays 18 times (5).

In New Mexico, the Gardner case provides the clearest documented example of a dentist using digitally altered X-rays for fraud. He submitted the same manipulated radiograph for multiple patients, with fake caries added to justify crowns. Even after his licence was revoked, undercover investigators caught him still practising (54). Stories like this make you wonder how many similar cases go undetected.

Generative AI makes the outlook even more uncertain. A 2026 study published in Radiology found that radiologists could only identify AI-generated X-rays 41% of the time when they did not know fakes were in the mix. Even when they were told to look for them, accuracy only reached 70-75% (28). If trained radiologists cannot reliably tell real from fake, dentists have no chance. Digital watermarking and blockchain-based authentication have been proposed as solutions (23,24), but none has been widely adopted in dental practice.

Think about what this means for a daily clinical situation. A patient transfers from another practice and brings their digital X-rays on a USB drive. The new dentist opens the files, sees what appears to be extensive decay, and plans a comprehensive

treatment. But what if those images had been altered, either by the previous practitioner to justify billing or by someone else entirely? There is currently no way for the receiving dentist to verify whether those images are authentic. No digital signature, no timestamp, no chain of custody. The files could have been modified ten times and nobody would know.

Insurance companies face the same problem on a much larger scale. They receive hundreds of thousands of radiographic images with every batch of claims. Human reviewers cannot examine each one in detail, and until recently there was no technology capable of screening images for manipulation at scale. Pearl's AI-based fraud detection system represents the first serious attempt, and its early results suggest the problem is far bigger than anyone thought (5). We are still in the early days of this technology, and the tools for creating fakes are advancing at least as fast as the tools for detecting them.

5.2 The over-detection trap: when finding more disease is not always better

Commercial AI tools for dental imaging are marketed with numbers that sound impressive. Pearl says its system finds 37% more disease per radiograph than dentists find on their own. Overjet has documented a 100% increase in case acceptance at one dental group. VideaHealth reports a 43% drop in missed diagnoses (3,5). At first glance, this looks like progress. More detection should mean better care.

Yet the data tells a more complicated story. Arzani et al. reviewed 14 systematic reviews covering 137 individual studies and found that AI caries detection achieves a pooled sensitivity of 0.85 and specificity of 0.90 (29). That translates to roughly one in five AI-positive caries diagnoses being wrong, which means the patient gets flagged for disease that is not there. In some study populations, the picture is much worse. Ammar et al. found positive predictive values as low as 0.15 in certain contexts, meaning 85% of AI-flagged caries were false positives (30). An external validation study from Munich tested AI caries detection on clinical photographs of 991 lesions and found that the system missed 18.7% of actual cavities while also wrongly flagging healthy teeth (31).

This matters because dentistry already has a well-documented over-treatment problem. A 2024 Korean study sent people with perfectly healthy teeth to 196 dental clinics and found that two out of three visits resulted in a diagnosis of caries that did not exist (55).

Now imagine deploying AI that flags even more potential cavities on top of an environment that already tends to overdiagnose. The risk is obvious: more healthy teeth getting drilled, more crowns placed on teeth that needed nothing, and patients paying for procedures they never needed. Schwendicke et al. put it well: finding a cavity by itself does not help anyone; the value only comes from the treatment that follows, and if that treatment is based on a false positive, the patient is worse off, not better (56).

To put this in practical terms, think about a dental clinic that sees 30 patients a day and takes bitewing radiographs on 20 of them. If the AI flags an average of three suspicious areas per set of images and the false positive rate is 20%, that clinic is generating 12 false alerts every single day. Over the course of a year, that adds up to roughly 3,000 false positive findings at a single practice. Even if only a fraction of those lead to actual treatment, the cumulative impact on patients and healthcare costs is substantial. Multiply this across the thousands of clinics now using AI diagnostic tools, and you start to understand the scale of the problem.

Financially, the incentives point in a troubling direction. AI companies earn money when clinics adopt their tools and when the tools find things. Clinics bill more when AI flags more pathology. Patients accept more treatment when they see an AI highlight on their X-ray. Everyone profits except possibly the patient whose healthy tooth just got a filling (22). No major commercial dental AI platform has published data on how often its system triggers unnecessary treatment. That silence is itself a problem.

There is also a psychological dimension worth considering. When a patient sees a colourful overlay on their X-ray with AI-generated annotations pointing to potential problems, the visual impact is powerful. It feels scientific, authoritative, and hard to argue with. Research shows that patients are far more likely to accept treatment when they see visual evidence on a screen compared to when a dentist simply tells them about a finding (5).AI annotations turn every Xray into a sales presentation, whether or not the findings are accurate (figure3).

5.3 Impact of artificial intelligence on clinical reason-making and decision-making

Automation bias is the tendency to trust a computer over your own eyes, even when the computer is wrong. It is well documented in aviation, nuclear power, and medicine, and now it is showing up in dentistry. Glick et al. ran the only randomised study testing this in dental students: when AI incorrectly flagged furcation involvement where none existed, students with AI assistance were significantly more likely to agree with the machine than those working without it (43).They did not catch the error. They went along with it.

The broader medical evidence backs this up. Dratsch et al. showed that even experienced radiologists were swayed by wrong AI suggestions when reading mammograms (44).Saragih et al. found in a 2024 study of 210 participants that automation bias persisted across all groups, though certified training helped reduce the rate of false agreement (45).Cabitza et al. warned back in 2017 that deskilling and overconfidence were predictable side effects of deploying machine learning in clinical settings (46).

What concerns me most is the next generation. The FDI World Dental Federation identified three risks in its 2026 guidelines: deskilling, which means losing abilities you once had; never-skilling, which means never developing them because the AI does the work for you; and mis-skilling, which means building a false sense of competence based on AI output rather than genuine clinical reasoning (22).If AI becomes the default first reader for every radiograph from day one of dental school, when will students learn to read X-rays on their own? And if they never develop that skill, what happens when the AI goes down, gives a wrong answer, or encounters a case outside its training data?

This is not just a theoretical worry. Talk to any experienced clinician and they will tell you that reading radiographs is a skill that takes years of practice. You learn to notice subtle density changes, to recognise the shadow of a crack, to distinguish a developing lesion from an overlapping restoration. This skill is built through thousands of repetitions, most of which happen during dental school and the first few years of practice. If those repetitions are shortcircuited by AI that reads the image first and tells the student what to look for, the student never develops the independent judgement needed to practise

safely. They become dependent on a tool that, as we have seen, gets it wrong a significant portion of the time.

Lal et al. described a case where AI software flagged a fissure sealant as caries (3). A fissure sealant. If the dentist had not overridden the system and double-checked, a perfectly protected tooth would have been drilled open for no reason. This kind of thing is probably happening in clinics every day, but we have no way to know how often because nobody is tracking it. In endodontics, AI systems miss up to two-thirds of periapical lesions on panoramic radiographs (33). In a clinical setting, that means a patient could walk out of an appointment with an undiagnosed infection because the AI did not flag it and the dentist relied on the AI instead of looking carefully themselves.

Banning AI from dental education is not the solution. That would be like banning calculators from mathematics classes unrealistic and counterproductive. Instead, dental schools should teach students how to use AI critically. They need to learn to read radiographs without AI first, then learn to work with AI while understanding its limitations. They need to know what the false positive and false negative rates are for the tools they use. They need practice in overriding AI when their clinical judgement tells them something different. This kind of training barely exists in dental curricula today, and that needs to change fast (22).

5.4 The black box and the consent problem

Most dental AI systems are proprietary. That means the companies keep their code, their training data, and their decision logic confidential. From a business perspective, that makes sense. From a clinical and ethical perspective, it creates a serious problem (9,10).

Informed consent requires that the patient understands what is being done, why, and based on what evidence. If a dentist recommends a procedure partly because an AI system flagged something, but neither the dentist nor the patient can explain what the AI actually detected or how it reached that conclusion, the consent is hollow. Rokhshad et al. developed a framework with 11 ethical principles for dental AI, including transparency and explainability (47).Roganovic created the first informed consent checklist specifically for AI-assisted dental treatment (57).Yet these tools are not being used in most clinics.

Explainability technology exists. Ma et al. demonstrated heat maps that show which parts of an image the AI focused on when making a decision (9).Oztekin et al. used Grad-CAM visualisations to make caries detection more transparent (58).These approaches let clinicians verify whether the AI is looking at the right thing. Yet they remain almost entirely in the research world. None of the major commercial platforms currently offers real-time visual explanations to clinicians during regular use (3).

Imagine the difference it would make if, every time the AI flagged a suspicious area on an X-ray, the dentist could see a heat map showing exactly which pixels drove the decision. If the highlight sits squarely on a shadow at the dentin-enamel junction, that is useful confirmation. If it is centred on a restoration margin or a fissure sealant, the dentist immediately knows to override the system. Without that kind of visual feedback, the clinician is forced to either trust the AI blindly or ignore it entirely. Neither option is good medicine.

Transparency concerns extend beyond individual diagnoses. When a dental practice adopts an AI platform, it is typically making a long-term commitment. Its software integrates with the practice management system, the billing workflow, and the patient communication process. Switching to a different platform is expensive and disruptive. This gives AI vendors significant market power with little accountability. If a platform's algorithm changes with an update, the practice may not even be informed, let alone given the chance to evaluate whether the new version performs better or worse. Clinicians deserve to know when the tool they rely on has changed how it works, and they deserve access to performance data that lets them make informed decisions about continuing to use it.

5.5 AI-washing: when marketing writes cheques the technology cannot cash

AI-washing is what happens when companies label basic software features as artificial intelligence to charge more or impress consumers, without the technology actually delivering what is promised. It is the tech equivalent of greenwashing, and it is everywhere in dental AI (3,8).

Ruiz et al. reviewed 13 commercial clear aligner programmes and found that not a single one had been independently validated in a peer-reviewed study, despite all of them

making AI claims in their marketing (2).Naved et al. analysed 52 FDA-cleared dental AI devices and discovered that 60% did not disclose their algorithm type and 50% had no public documentation of clinical deployment (14).These products are being used on real patients in real clinics every day, and there is almost no independent evidence that they do what they claim to do.

SmileDirectClub is the most dramatic example of what happens when AI-driven treatment is sold at scale without proper oversight. SDC used algorithm-based treatment planning with remote review by dentists who typically never examined patients in person. No X-rays, no periodontal check, no hands-on evaluation. The company treated over two million people before filing for bankruptcy in September 2023. By December 2023 it had shut down completely, stranding thousands of patients mid-treatment with locked accounts, inaccessible dental records, and a voided lifetime guarantee (59,60).

Reports in the FDA MAUDE database describe of patients who developed severe open bites, TMJ pain, tooth mobility, and nerve damage after SDC treatment (61).New York's Attorney General recovered $4.8 million for over 28,000 affected consumers (62).In Washington D.C., the Attorney General found that SDC had forced more than 17,000 customers to sign agreements prohibiting negative reviews, regulatory complaints, or even acknowledging the agreement existed, under penalty of $10,000 per violation (63). Meanwhile, the American Association of Orthodontists filed complaints with 36 state dental boards (64).As the British Dental Association put it, what many were thinking: it should not have taken a bankruptcy to protect patients from harm (65).

The clinical data on clear aligner prediction accuracy adds another dimension. Haouili et al. found that ClinCheck predictions achieved only 50% overall accuracy, with the weaknesses remaining essentially unchanged over a decade (11). Rossini et al. confirmed that extrusion movements achieve only about 30% accuracy and rotation of rounded teeth is poorly controlled (36). Lombardo et al. reported that arch length predictions, which are critical for deciding whether teeth need to be extracted, were accurate only 30-38% of the time (35).

What this means in practice is that patients are shown a digital picture of their future smile that looks precise and scientific, but the actual treatment delivers barely half of

what was promised on screen for overall movements, and as little as 18% for specific types of movement. A patient who consents to treatment based on that simulation has been given incomplete information, no matter how good the intentions of the treating clinician. When the result falls short, the patient feels misled, the clinician faces a difficult conversation, and in some cases, additional treatment with traditional braces becomes necessary at additional cost and time (figure 4).

Remote monitoring raises just as many concerns. Ferlito et al. tested AI-based monitoring of clear aligner therapy and found that one in five repeated scans gave contradictory instructions. It could not even agree with itself on which teeth had tracking problems, showing 0% agreement between repeated assessments (34). Patients and orthodontists relying on this technology to make real-time treatment decisions are, in effect, following advice that changes randomly depending on when the scan is taken. That is not guidance. That is noise.

5.6 Legal responsability and liability in AI-assisted dental practice

This is a question that nobody has fully answered yet. Mathew et al. laid out the central dilemma in their 2025 analysis: when AI and the clinician disagree, both choices carry risk. If you follow the AI and it turns out to be wrong, you may face liability for not exercising independent judgement. If you override the AI and it turns out to be right, you may face liability for ignoring available technology (48).D'Cruz went a step further, suggesting that as AI tools become more established, not using them could itself become a breach of the standard of care (49).

Current law was not designed for this. Legal standards like the Bolam test and Montgomery ruling, which govern medical liability in much of the English-speaking world, assume that a human makes the decision. They do not account for a situation where a machine recommends a course of action that the human then either follows or rejects (48,49).Under the EU AI Act, medical AI devices are classified as high-risk and imposes transparency requirements, but how this will play out in actual malpractice cases remains to be seen (12,13).

The proposed solutions lean towards shared liability models, where responsibility is distributed between the clinician, the AI developer, and potentially the healthcare

institution (13). Until case law develops around actual AI-related dental malpractice claims, practitioners are operating in a grey area.

Think about a concrete scenario. A general dentist uses an AI system to evaluate a panoramic radiograph. Its algorithm flags an area near the apex of a lower molar as suspicious. Trusting the system, the dentist, refers the patient for endodontic treatment. A root canal is performed. Six months later, a different radiograph shows the original finding was a normal anatomical variation, not pathology. An unnecessary procedure on a healthy tooth. Who is responsible? The general dentist who followed the AI? The AI company whose product gave a false positive? The endodontist who proceeded without questioning the referral? Under current law, the answer is unclear. Most likely, the general dentist would bear primary liability, because the law still treats AI as a tool rather than a decision-maker. But that framework was built for drills and impression materials, not for software that reads X-rays and tells you what to do.

Insurance companies face their own legal exposure. If an insurer uses AI to approve or deny claims and the AI makes systematic errors, whether approving fraudulent claims or denying legitimate ones, the legal consequences could be substantial. Pearl's fraud detection AI has already exposed patterns that human reviewers missed for years (5), which raises an uncomfortable question: were insurers negligent for not using AI sooner? Or are they creating new risks by deploying technology that is itself imperfect? These questions do not have easy answers, and the legal system will likely need years to work them out.

5.7 Cybersecurity: the problem nobody wants to talk about

Dental practices hold exactly the kind of data that criminals find most valuable: Social Security numbers, insurance details, medical histories, and high-resolution images. And yet most practices are shockingly unprepared for cyberattack (8,20,66).

In 2023, the MCNA Dental breach exposed 8.9 million patient records after the LockBit ransomware group exfiltrated 700 GB of data and published everything on the dark web when a $10 million ransom was refused (38).Just months later, the MOVEit vulnerability hit Delta Dental of California, compromising 6.9 million records through the Russian-linked Clop gang (39).Back in 2019, the DDS Safe attack had disabled around 400 dental

practices simultaneously through a supply chain hack, locking dentists out of their own systems for an average of nearly ten days (40).

Henry Schein, the world's largest dental supply company, was hit three separate times between 2023 and 2024 by the BlackCat ransomware group. One attack took its e-commerce platform offline for weeks, drove 10-15% of its customers to competitors, and forced a restructuring targeting $75-100 million in annual cost savings (41).The Westend Dental case in Indiana was even worse in some ways: the practice suffered a ransomware attack and then tried to cover it up, telling investigators the data loss came from a hard drive that was accidentally formatted. They had no HIPAA training, no password policies, no risk assessment, and stored servers in break rooms and bathrooms. Staff had posted patient X-rays on social media. Indiana imposed a penalty of $350,000 (42).

Blockchain and federated learning have been proposed as solutions. Blockchain can create tamper-proof audit trails for dental records (23,24),while federated learning allows AI models to train on data from multiple clinics without the data ever leaving the individual practice (67).Both technologies are promising but adoption in dentistry is still minimal.

For the average dental practice, the practical impact of a cyberattack goes far beyond the technical. When systems go down, appointments have to be cancelled, patient histories become unavailable, and treatment planning grinds to a halt. One Wisconsin dentist affected by the DDS Safe attack described the situation bluntly: they had no access to patient charts, schedules, Xrays, or payment systems, and could not do proper treatment without chart history and radiographs (40). For solo practitioners or small group practices, a ransomware attack can be existentially threatening. The costs of recovery, the loss of patient trust, and the regulatory penalties for data exposure can easily exceed what a small practice can absorb (figure 5).

Most dental offices, frankly, treat cybersecurity as an afterthought. A typical small practice might use shared passwords, have no two-factor authentication, skip regular software updates, and store backups on the same network as the main system. Staff might open phishing emails without thinking twice. Patient data might be sent through unencrypted email. These are not edge cases. Multiple surveys and incident reports suggest they describe the norm rather than the exception (66). Until the profession takes

data security as seriously as it takes infection control, dental patients will remain at risk of having their most sensitive information exposed.

5.8 Tele-dentistry: convenience with hidden costs

Tele-dentistry expanded rapidly during the pandemic and has stayed. For follow-up consultations, patient education, and triaging urgent cases, it makes sense. For active treatment monitoring, though, especially with AI-driven tools, the evidence raises red flags.

Ferlito et al.'s 2023 study tested AI-based remote monitoring of clear aligner therapy on 50 patients and found results that should worry anyone relying on these tools. One in five repeated scans gave a different answer about whether treatment was on track. The system identified different problematic teeth on each scan, with literally zero agreement between attempts (34). For a technology that patients and clinicians are trusting to guide active orthodontic treatment, that level of inconsistency is unacceptable.

Beyond technical issues, tele-dentistry creates practical opportunities for fraud. It is difficult to verify whether services were actually delivered during a remote consultation (50,51). Digital images can be altered to show procedure completion that never happened (52). Insurance claims can be submitted for consultations that were minimal or never occurred. None of these are hypothetical problems. They are logical consequences of a system that removes the physical verification that in-person care provides.

For vulnerable populations, the risks are even higher. McGrath et al. highlighted that AI-powered remote monitoring raises particular ethical concerns when applied to elderly patients, children, and underserved communities who may have limited digital literacy and reduced ability to evaluate whether the care they receive remotely is adequate (53). These are often the very populations that tele-dentistry is supposed to help reach.

The balance here is tricky. Tele-dentistry genuinely improves access for patients in rural areas, for those with mobility limitations, and for routine follow-ups that do not require hands-on examination. Throwing it out entirely would be wrong. But deploying it with AI monitoring tools that contradict themselves 20% of the time, in a regulatory environment that has not figured out how to verify remote services, is a recipe for

problems. The profession needs clear standards for what can and cannot be managed remotely, combined with audit mechanisms that ensure billed services actually correspond to care delivered.

5.9 Future directions and recommendations for safe AI integration in dentistry

Looking at all the evidence gathered in this review, several directions stand out as particularly important for the near future. The first is the development and adoption of authentication standards for digital dental images. Whether through digital watermarking, blockchain-based timestamping, or DICOM-specific security protocols, there needs to be a way for anyone receiving a dental radiograph to verify that the image has not been modified since it was captured. This technology exists and has been proposed repeatedly in the literature since at least 2008, but the profession has not acted on it. Every year of inaction is another year in which manipulated images circulate through insurance systems and clinical records unchallenged.

Independent post-market surveillance is the second priority of commercial AI diagnostic tools. Right now, the regulatory model, where products can reach the market through pathways that do not require independent clinical validation, leaves a dangerous gap between marketing claims and clinical reality. Dental professional bodies, academic institutions, and regulatory agencies need to collaborate on creating systematic evaluation programs that test AI tools in real clinical conditions, across diverse patient populations, with transparent reporting of performance metrics including false positive and false negative rates.

Third, dental education must integrate AI literacy as a core competency. Students need to understand not just how to use AI tools but how to evaluate them critically. They need handson training in identifying automation bias in their own decision-making. They need to be able to read a study evaluating an AI system and determine whether the evidence supports the marketing claims. And they need to develop strong foundational clinical skills before they become dependent on algorithmic assistance.

Fourth, the profession needs to develop clear informed consent protocols for AI-assisted treatment. Patients have a right to know when AI influences their diagnosis or treatment plan, what the limitations of that AI are, and how the clinician has verified the AI's

recommendations. The consent checklists developed by Roganovic represent a good starting point, but they need to be adopted and adapted by professional bodies worldwide (57).

Finally, cybersecurity needs to be treated as a patient safety issue, not just a technical one. The data shows that dental practices are being attacked regularly and that the consequences for patients can be severe. Professional bodies should set minimum cybersecurity standards for dental practices, just as they set standards for sterilisation and radiation safety. Insurance carriers should require evidence of cybersecurity compliance. And dental schools should teach data security alongside clinical skills.

Beyond these specific recommendations, the profession needs a broader change in mindset. Right now, dental AI is largely treated as a consumer product: companies develop it, market it, and sell it to clinics, and the clinics use it on patients with minimal oversight. The assumption is that if a product has FDA clearance, it must be safe and effective. But as we have seen, FDA clearance through the 510(k) pathway does not actually prove clinical effectiveness. It proves similarity to an existing device. That is not the same thing, and the profession should stop treating it as if it were.

What dental AI needs is the kind of rigorous, independent, ongoing evaluation that we expect for pharmaceuticals. Nobody would accept a new drug being marketed to millions of patients based solely on the manufacturer's internal data. Yet that is essentially what is happening with dental AI. The companies run their own studies, publish their own numbers, and clinicians take them at face value. Independent academic centres should be testing these tools in real clinical conditions, with diverse patient populations, and publishing the results, good or bad, in peerreviewed journals. Until that happens, clinicians are essentially beta-testing commercial software on their patients.

Professional associations also have a role to play. While the FDI has taken a lead with its 2024 policy statement, but implementation at the national level is patchy at best (22).National dental associations should be developing practical guidelines that tell their members how to evaluate AI tools before adopting them, what to disclose to patients when AI contributes to a diagnosis, and what steps to take when AI and clinical

judgement disagree. These guidelines should not be aspirational. They should be specific, actionable, and backed by continuing education requirements.

Overall, while AI offers many benefits in dentistry, it also raises important risks, especially related to fraud, bias, and cybersecurity, mainly because control, detection systems, and regulation are still limited, as shown in Table 9.

6. CONCLUSION

Based on the findings of this systematised review, the following conclusions can be drawn:

Digital and AI-related fraud in dentistry takes multiple forms that can be grouped into seven categories: digital radiograph manipulation, AI-driven over-diagnosis, deceptive marketing of AI products (AI-washing), automation bias in clinical decision-making, cybersecurity breaches, teledentistry fraud risks, and gaps in ethical and legal accountability. These categories are interconnected, and the rapid pace of digitalisation means that all of them are growing in scope and complexity.

Current detection methods are limited but promising. AI-based forgery detection tools achieve near-perfect accuracy in controlled settings, but none has been adopted in routine dental practice. Digital watermarking, blockchain authentication, and metadata analysis have been proposed as safeguards for dental images, yet implementation remains minimal across the profession. There is currently no standard authentication protocol for dental radiographs.

Ethical and legal frameworks are lagging behind the technology. Only 12.4% of published dental AI studies report any ethical concerns. Informed consent is compromised when AI contributes to clinical decisions that neither the dentist nor the patient can fully explain. Legal liability when AI and clinician disagree remains unresolved, with both following and overriding the AI carrying potential risk for the practitioner (20,47,48,49).

The clinical impact of AI-related fraud on patients is real and documented. False positive rates for AI caries detection reach up to 85% in some populations, directly increasing the risk of unnecessary treatment. Clear aligner prediction software achieves as little as 18% accuracy for certain tooth movements, yet patients make treatment decisions based on these simulations. Over two million patients were left without support when SmileDirectClub collapsed, and nearly 90 million dental patient records have been exposed in recent cyberattacks.

A prevention framework for digital fraud in dentistry should combine technical safeguards (image authentication, independent AI validation, explainable AI), educational measures (AI literacy training for dental students, awareness of automation bias), regulatory reforms (mandatory post-market surveillance of dental AI, harmonised

international standards), and institutional protocols (cybersecurity standards, informed consent checklists for AI-assisted treatment). None of these measures is sufficient alone; they need to work together.

7. SUSTAINABILITY

This project encourages a more responsible and sustainable model of dentistry by combining technological progress with ethical awareness. Preventing and identifying digital fraud goes beyond protecting individual patients; it helps maintain confidence in the healthcare system over time. From a social standpoint, the careful and transparent use of AI can improve the relationship between dentist and patient. When practitioners question and interpret AI results instead of accepting them automatically, innovation becomes beneficial for a wider population. Reducing algorithmic bias is also essential to ensure fair and consistent care for all patients.

On an economic level, errors and fraud lead to unnecessary spending that could be better invested in meaningful treatments. Although implementing secure AI systems and cybersecurity measures requires upfront investment, it significantly lowers the risk of higher costs related to incorrect procedures, legal issues, or data breaches. In healthcare, prevention is often more cost-effective than correction.

Digital practices also contribute to environmental sustainability by limiting paper use, reducing chemical waste, and avoiding repeated imaging. Finally, training future professionals to use AI critically while preserving their clinical judgment strengthens the profession’s long-term stability and adaptability.

8. REFERENCES

1. Schwendicke F, Samek W, Krois J. Artificial intelligence in dentistry: chances and challenges. J Dent Res.2020;99(7):769-74.
2. Ruiz DC, Muresanu S, Du X, Elgarba BM, Fontenele RC, Jacobs R. Unveiling the role of artificial intelligenceapplied to clear aligner therapy: a scoping review. J Dent. 2025;154:105564.
3. Lal A, Nooruddin A, Umer F. Concerns regarding deployment of AI-based applications in dentistry - a review.BDJ Open. 2025;11(1):27.
4. Shujaat S, Aljadaan H, Alrashid H, Aboalela AA, Riaz M. FDA-approved AI solutions in dental imaging: anarrative review of applications, evidence, and outlook. Int Dent J. 2026;76(1):109315.
5. Pearl AI. Computer vision breakthrough allows dental insurers to review 100% of claims-based oral imagerywith superhuman precision [press release]. 2023.
6. Asuquo NE, Abbas SG, Bendiab G, et al. Vulnerability to cyberattacks and sociotechnical solutions for healthcare systems: systematic review. J Med Internet Res. 2024;26:e46904.
7. Diaz-Flores-Garcia V, Labajo-Gonzalez E, Santiago-Saez A, Perea-Perez B. Detecting the manipulation ofdigital clinical records in dental practice. Radiography. 2017;23(4):e103-7.
8. Naved N, Adnan N, Umer F. Development trend of artificial intelligence (AI) in dentistry: exploring FDAcleareddental devices. BMC Oral Health. 2026;26:11.
9. Ma J, Schneider L, Lapuschkin S, et al. Towards trustworthy AI in dentistry. J Dent Res. 2022;101(11):1263-8.
10. Maart RD, Mulder R. Ethical considerations for artificial intelligence in dentistry. S Afr Dent J.2024;79(3):142-8.
11. Haouili N, Kravitz ND, Vaid NR, Ferguson DJ, Tamburrino RK. Has Invisalign improved? A prospectivefollow-up study on the efficacy of tooth movement with Invisalign. Am J Orthod Dentofacial Orthop. 2020;158(3):420-5.
12. Minssen T, Vayena E, Kesselheim AS, et al. Navigating the EU AI Act: implications for regulated digitalmedical products. NPJ Digit Med. 2024;7:261.
13. B.K. et al. Artificial intelligence-driven dentistry: a systematic review of ethical and legal challenges. Int JDent. 2026;2026:1870800.
14. Schwab K, Uribe SE, Goncalves TM, et al. Concerns regarding deployment of AI in dentistry: a focus groupstudy. J Dent. 2025;152:105458.
15. Association of Certified Fraud Examiners. Dental fraud: a multi-faceted international crime problem. FraudExaminer. 2022.

16. United States Department of Justice. Pennsylvania brothers convicted of decades long racketeeringconspiracy [press release]. 2023.
17. United States Department of Justice. Houston dental clinic operator heads to prison for role in $6M pediatricfraud scheme [press release]. 2022.
18. United States Department of Justice. Federal jury finds Grafton dentist guilty of healthcare fraud and falsestatements [press release]. 2022.
19. United States Department of Justice. Grafton dentist sentenced to 54 months imprisonment and ordered topay over $1 million in forfeiture [press release]. 2022.
20. Morch CM, Atsu S, Cai W, et al. Artificial intelligence and ethics in dentistry: a scoping review. J Dent Res.2021;100(13):1452-60.
21. Allareddy V, Oubaidin M, Rampa S, et al. Call for algorithmic fairness to mitigate amplification of racial biasesin AI models used in orthodontics and craniofacial health. Orthod Craniofac Res. 2023;26(Suppl 1):124-30.
22. FDI World Dental Federation. Artificial intelligence in dentistry: FDI policy statement and white paper. Int DentJ. 2024;74(6):1138-45.
23. Mokhamed T, Abu Talib M, Moufti MA, et al. The potential of blockchain technology in dental healthcare: aliterature review. Sensors. 2023;23(6):3277.
24. Cerda Mardini D, et al. Blockchain for trustworthy artificial intelligence in dentistry. JDR Clin Trans Res.2025;10(4):469-71.
25. Bik EM, Casadevall A, Fang FC. The prevalence of inappropriate image duplication in biomedical researchpublications. mBio. 2016;7(3):e00809-16.
26. Amiri E, Mosallanejad A, Sheikhahmadi A. CFDMI-SEC: an optimal model for copy-move forgery detection ofmedical image using SIFT, EOM and CHM. PLoS One. 2024;19(7):e0303332. 33
27. Abdalla MH, Ali HA, Omer OA. Detecting image manipulation with ELA-CNN integration: a powerfulframework for authenticity verification. PeerJ Comput Sci. 2024;10:e2205.
28. Tordjman M, et al. The rise of deepfake medical imaging: radiologists' diagnostic accuracy in detectingChatGPT-generated radiographs. Radiology. 2026;318:e252094.
29. Arzani E, et al. Examining the diagnostic accuracy of artificial intelligence for detecting dental caries across arange of imaging modalities: an umbrella review with meta-analysis. PLoS One. 2025;20(7):e0329986.
30. Ammar N, et al. Artificial intelligence platforms in dental caries detection: a systematic review andmeta-analysis. J Oral Rehabil. 2025;52(3):245-58.

31. Tural M, et al. Caries detection and classification in photographs using an artificial intelligence-based model: an external validation study. Diagnostics. 2024;14(20):2281.
32. Mohammad-Rahimi H, Motamedian SR, Rohban MH, et al. Deep learning for caries detection: a systematicreview. J Dent. 2022;122:104115.
33. Koenig LJ, et al. Periapical lesions in panoramic radiography and CBCT imaging: assessment of AI'sdiagnostic accuracy. Diagnostics. 2024;14(10):1041.
34. Ferlito T, Schneider D, Hsiou D, et al. Assessment of artificial intelligence-based remote monitoring of clearaligner therapy: a prospective study. Am J Orthod Dentofacial Orthop. 2023;164(2):194-200.
35. Lombardo L, et al. The reliability of ClinCheck accuracy before and after Invisalign treatment: a multicenterretrospective study. Appl Sci. 2023;13(8):4670.
36. Rossini G, Parrini S, Castroflorio T, Deregibus A, Debernardi CL. Efficacy of clear aligners in controllingorthodontic tooth movement: a systematic review. Angle Orthod. 2015;85(5):881-9.
37. Wexler A, Nagappan A, Beswerchij A, Choi R. Direct-to-consumer orthodontics: surveying the userexperience. J Am Dent Assoc. 2020;151(8):625-36.e4.
38. MCNA Dental data breach: LockBit ransomware attack affecting 8.9 million individuals. Security Affairs.
39. Delta Dental of California data breach: 7 million individuals affected. HIPAA Journal. 2024.
40. Krebs B. Ransomware bites dental data backup firm. Krebs on Security. 2019 Aug 29.
41. Henry Schein discloses data breach a year after ransomware attack. Bleeping Computer. 2024.
42. Westend Dental HIPAA violations: Indiana AG agrees to $350,000 penalty. HIPAA Journal. 2024.
43. Glick A, Clayton M, Angelov N, Chang J. Impact of explainable artificial intelligence assistance on clinicaldecision-making of novice dental clinicians. JAMIA Open. 2022;5(2):ooac031.
44. Dratsch T, Chen X, Rezazade Mehrizi M, et al. Automation bias in mammography: the impact of artificialintelligence BI-RADS suggestions on reader performance. Radiology. 2023;307(4):e222176.
45. Saragih ID, et al. Automation bias in AI-decision support: results from an empirical study. Stud HealthTechnol Inform. 2024;316:1592-6.
46. Cabitza F, Rasoini R, Gensini GF. Unintended consequences of machine learning in medicine. JAMA.2017;318(6):517-8.
47. Rokhshad R, Ducret M, Chaurasia A, et al. Ethical considerations on artificial intelligence in dentistry: aframework and checklist. J Dent. 2023;135:104593. 34

48. Mathew MG, Jose S, Yadav S, Cherian J. AI in dentistry: a legal minefield? Br Dent J. 2025;238:838.
49. D'Cruz L. Some dento-legal considerations about AI in dentistry. Br Dent J. 2025;239:408.
50. Squires T, Michelogiannakis D, Rossouw PE, Javed F. An evidence-based review of the scope and potentialethical concerns of teleorthodontics. J Dent Educ. 2021;85(1):92-100.
51. Malpe M, Choudhari SG, Nagtode N, Gharde PM. Beyond the chair: exploring the boundaries ofteledentistry. Cureus. 2024;16(6):e62286.
52. Bhargava A, Sabbarwal B, Jaggi A, et al. Teledentistry: a literature review of evolution and ethicolegalaspects. J Glob Oral Health. 2019;2(2):128-33.
53. McGrath C, Chau CWR, Molina GF. Monitoring oral health remotely: ethical considerations when using AIamong vulnerable populations. Front Oral Health. 2025;6:1587630.
54. Dentistry Today. Digital dental radiographs and their nefarious manipulation. Dentistry Today. 2021.
55. Kim HJ, et al. Overdiagnosis of dental caries in South Korea: a pseudo-patient study. BMC Oral Health.2024;24:1432.
56. Schwendicke F, Cejudo Grano de Oro J, Garcia Cantu A, et al. Artificial intelligence for caries detection: value of data and information. J Dent Res. 2022;101(11):1277-85.
57. Roganovic J. Developing a consent checklist for AI in dentistry: thematic analysis and pilot survey validation.Digit Health. 2025;11:20552076251393227.
58. Oztekin F, Katar O, Sadak F, et al. An explainable deep learning model to prediction dental caries usingpanoramic radiograph images. Diagnostics. 2023;13(2):226.
59. SmileDirectClub stops operating months after filing for bankruptcy. Fortune. 2023 Dec 11.
60. SmileDirectClub bankruptcy: shutdown leaves customers in limbo. Philadelphia Inquirer. 2023 Dec 12.
61. FDA MAUDE adverse event reports: SmileDirectClub aligners. FDA MAUDE Database. 2019-2023.
62. Attorney General James recovers $4.8 million for consumers wrongly charged by SmileDirectClub. New YorkState Attorney General. 2024.
63. AG Schwalb announces SmileDirectClub must release consumers nationwide from restrictive nondisclosureagreements. D.C. Office of Attorney General. 2023.
64. American Association of Orthodontists. State dental board complaints: SmileDirectClub. AAO. 2019-2023.

65. British Dental Association. SmileDirectClub collapse confirms the risks to patients of remote orthodontics. BrDent J. 2024;236:13.
66. Nalliah RP, Praveen S, Allareddy TV, et al. Cybersecurity threats and preparedness: implications for dentalschools. J Dent Educ. 2025;89(4):523-8.
67. Rischke R, Schneider L, Muller K, et al. Federated learning in dentistry: chances and challenges. J Dent Res.2022;101(11):1269-73.

9. Annexes

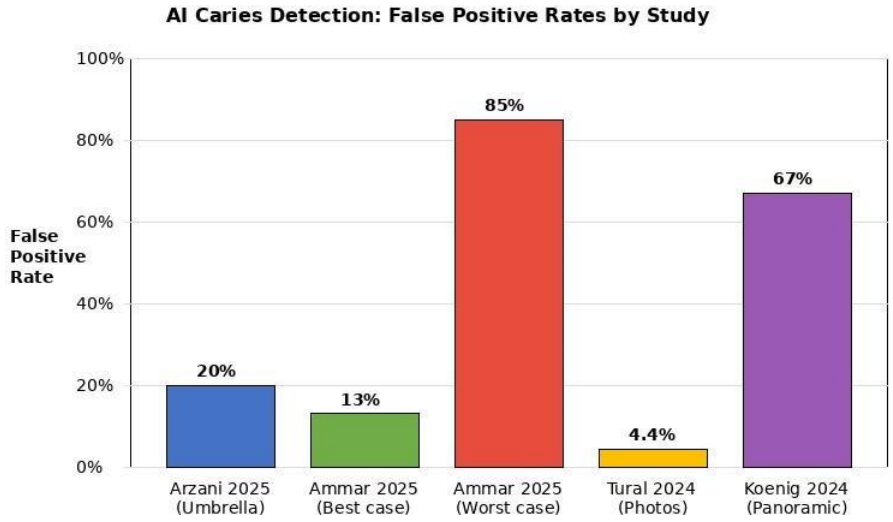


Figure 3: AI caries detection false positive rates across major studies (29,30,31,33)

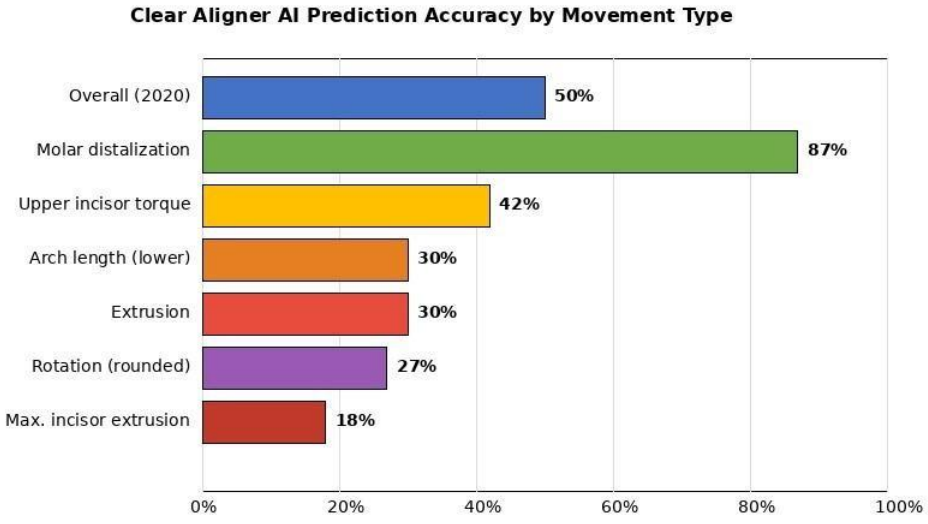


Figure 4: Clear aligner AI prediction accuracy by movement type. (11,35,36)

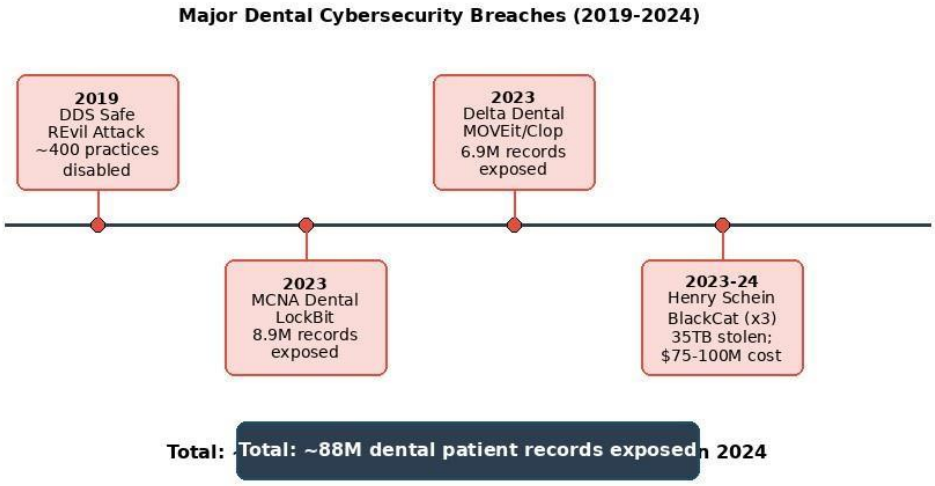


Figure 5. Timeline of major cybersecurity breaches affecting dental organisations (2019-2024) (38,39,40,41)

Table 9. Key risks and regulatory challenges of AI in dentistry

Fraud Domain	Key Finding	Risk Level	Detection Capability	Regulatory Status
Radiograph manipulation	Dentists detect at chance level; insurers approve 100% of fakes	Critical	AI fraud detection emerging; no standard adopted	No authentication standard for dental images
AI over-diagnosis	False positive rates up to 85%; no tracking of unnecessary treatment	High	No systematic monitoring exists	FDA 510(k); no post-market surveillance
AI-washing / marketing	0 of 13 aligner programmes validated; 60% FDA devices hide algorithms	High	Limited; no disclosure requirements	EU AI Act 2024; dental-specific gaps remain
Automation bias	Students agree with wrong AI ($p<0.05$); errors increase 26% with faulty AI	ModerateHigh	XAI tools exist but not used clinically	FDI 2024 policy; no enforcement
Cybersecurity	88M records exposed in 10 months; attacks up 755% since 2020	Critical	Blockchain proposed; adoption minimal	HIPAA/GDPR; dental-specific guidance lacking

Tele-dentistry fraud	20% contradictory AI decisions; service verification difficult	Moderate	No standard verification protocol	Inconsistent regulation
Legal accountability	Both following and overriding AI carry liability; no settled case law	High	Consent checklists proposed; rarely used	EU AI Act, FDA, FDI; not harmonised