



Federated learning for patient data privacy in dental artificial intelligence

Aprendizaje federado para la privacidad de los datos de pacientes en la inteligencia artificial odontológica

Ana Rodríguez Yanes¹  

¹ Universidad de Ciencias Médicas de La Habana, Facultad de Ciencias Médicas “Miguel Enríquez”, Departamento de Estomatología General Integral. La Habana, Cuba.

Submitted: 02-03-2026 | Revised: 10-05-2026 | Accepted: 07-07-2026 | Published: 12-07-2026

Corresponding Author: Ana Rodríguez Yanes 

ABSTRACT

Dental artificial intelligence needs large, diverse datasets to generalize, but the records that would form those datasets are protected by privacy law and are held in scattered clinics. Federated learning trains a shared model without moving the data, and it is the most promising answer to that tension, but its evidence base in dentistry is young. This applied review synthesizes the evidence on federated learning for dental artificial intelligence, ordered by practical relevance, and states concrete implications for practitioners. The evidence shows that federated learning reaches accuracy close to centralized training for caries detection, panoramic pathology detection and oral cancer screening, and that it can be combined with differential privacy and interpretability at a modest cost. In practice, a dental network that wants to pool its data without pooling its data should start with a federated caries or pathology detector, because those tasks are the ones the evidence supports most strongly. For regulators, the evidence implies that federated learning is not a substitute for legal compliance but a technical complement. The field still lacks a dental benchmark and long-term deployment studies, and each gap is paired with a concrete recommendation.

Keywords: federated learning; privacy; dental informatics; artificial intelligence; data protection.

RESUMEN

La inteligencia artificial odontológica necesita conjuntos de datos amplios y diversos para generalizar, pero los registros que formarían esos conjuntos están protegidos por la ley de privacidad y se encuentran dispersos en clínicas. El aprendizaje federado entrena un modelo compartido sin mover los datos, y es la respuesta más prometedora a esa tensión, pero su base de evidencia en odontología es joven. Esta revisión aplicada sintetiza la evidencia sobre aprendizaje federado para la inteligencia artificial odontológica, ordenada por relevancia práctica, y enuncia implicaciones concretas para los profesionales. La evidencia muestra que el aprendizaje federado alcanza una exactitud cercana al entrenamiento centralizado para la detección de caries, la detección de patologías panorámicas y el cribado del cáncer oral, y que puede combinarse con privacidad diferencial e interpretabilidad a un coste moderado. En la práctica, una red odontológica que quiera aunar sus datos sin aunarlos debería empezar por un detector federado de caries o de patología, porque esas son las tareas que la evidencia respalda con más fuerza. Para los reguladores, la evidencia implica que el aprendizaje federado no sustituye al cumplimiento legal sino que lo complementa técnicamente. El campo aún carece de un conjunto de referencia odontológico y de estudios de despliegue a largo plazo, y cada brecha se acompaña de una recomendación concreta.

Palabras clave: aprendizaje federado; privacidad; informática odontológica; inteligencia artificial; protección de datos.

INTRODUCTION

A dental artificial intelligence model that generalizes to new patients must be trained on records from many clinics, but those records cannot be pooled, because they contain patient data protected by law and by professional duty. This is the concrete problem that federated learning addresses, and it is the reason the technique has attracted attention in dentistry. A clinic that wants to contribute to a shared caries detector cannot send its radiographs to a central server, and a vendor that wants to train a panoramic pathology model cannot collect the images it would need from a single site, and neither party can solve the problem by simply ignoring the law that protects the patient. The result is that dental artificial intelligence has been trained on smaller and less diverse data than it needs, and its accuracy on new populations suffers for it. The scale of the problem is easy to understate: a dental network of even a few dozen clinics holds millions of images that no single clinic and no vendor can lawfully assemble, and every one of those images is a training example that is currently going to waste. Federated learning is the technique that promises to put those examples to use without moving them, and the gap between that promise and its evidence is the subject of this review.

Current answers to this problem do not suffice, and each fails in a specific way that a practitioner should understand. Centralized data sharing requires legal agreements, ethics approvals and de-identification that are slow and incomplete, and Giouroukou et al.(1) showed that even de-identified medical images can leak identity through metadata and pixel-level features, so the privacy of centralized data is never fully assured. Local-only training, in which each clinic trains its own model, protects privacy but wastes the statistical power of many small datasets, and a model trained on one clinic does not generalize to another, which is the exact failure that a dental network with many small clinics cannot afford. Synthetic data, the third current answer, replaces real records with generated ones but cannot reproduce the rare findings that matter most for diagnosis, and a synthetic dataset that has never seen a rare oral lesion cannot teach a model to find it. Federated learning is the fourth answer, and it is the one this review evaluates: the model travels, the data does not, and the statistical power of many clinics is combined without combining their data. The claim to evaluate is therefore not whether federated learning is elegant, but whether it delivers that combination in practice, and the evidence that follows is organized around that test.

The legal frame makes the problem concrete, and it is worth stating because the technical answer is judged against it. Under the General Data Protection Regulation and comparable laws, patient data may not be pooled across clinics without a lawful basis, and Brinz et al.(2) set out the legal frameworks that govern data sharing for responsible dental artificial intelligence, while Uribe et al.(3) described the data foundations that a trustworthy dental artificial intelligence requires. Federated learning is attractive precisely because it sidesteps the pooling, and a regulator must decide whether training a model without moving the data is, for the purpose of the law, the same as not sharing the data at all. That question is not yet settled, and it is the reason this review addresses regulators as well as practitioners.

The purpose of this review is practical. It synthesizes the evidence on federated learning in dental and medical imaging, ordered by what a practitioner can use first, and it states, for each practical implication, a concrete action and its recipient. A dental network considering federated learning, a clinic deciding whether to participate, and a regulator deciding whether federated learning satisfies a privacy requirement should each find, in the discussion, a decision they can act on. The review is not a survey of every federated learning method, because a survey would bury the decision under the technique, and the decision is what the reader needs. It is also not a systematic review, because the question it answers, what should a practitioner do, is a question of judgment that a systematic review can inform but not settle, and the judgment is laid open in the discussion rather than hidden behind a protocol.

METHODS

The sources were selected so that a dental informatics professional could repeat the process. Web of Science, Pubmed, EMBASE, Scopus and Open Alex were searched for publications that combined federated learning, differential privacy or privacy-preserving machine learning with dental, medical imaging or electronic health record terms. Records were merged, de-duplicated by digital object identifier and screened so that a study was retained when it developed or evaluated a federated or privacy-preserving method for a dental or transferable medical imaging task, or when it reviewed the field. Reference metadata were verified against Crossref, and a study

whose record could not be verified was excluded rather than retained on a title alone. The dental terms were deliberately narrow, because the dental federated learning literature is small, and the medical terms were deliberately broad, because the transferable evidence is what fills the dental gap, and the distinction between the two is preserved throughout the synthesis so that a dental claim is never mistaken for a medical one.

The synthesis ordered the retained evidence by practical relevance rather than by method or by anatomy. A dental application was ranked ahead of a medical one, a task with a clear clinical use ahead of a methodological refinement, and a study that reported accuracy together with a privacy guarantee ahead of one that reported either alone. The ordering reflects a judgment about what a practitioner can act on first, and it is stated openly so that a reader who disagrees can reorder the evidence without losing any of it. Two tables and one decision figure organize the evidence: one table lists the federated dental and medical applications, the second lists the privacy techniques and their trade-offs, and the figure presents the decision a practitioner faces. The method is deliberately light, because the review's audience is a professional who needs a decision, not a systematic reviewer who needs an exhaustive protocol, and the trade-off between completeness and usefulness is resolved in favour of usefulness. This choice is defensible only if the evidence that supports each recommendation is named in the text, and it is, so that a reader who wants to check a recommendation can follow the citation to its source and judge the strength of the support for themselves.

DEVELOPMENT

The evidence is ordered from what a practitioner can adopt today to what remains research. Federated learning is ready for dental detection tasks, it works with privacy and interpretability added, and its remaining gaps are in deployment rather than in feasibility.

What works now

Federated learning has been demonstrated on the dental tasks that matter most, and it reaches accuracy close to centralized training. For caries detection, Gopinath et al.(4) built an automated system that trained across X-ray images without moving them, and Kumar et al.(5) combined a privacy-aware and interpretable framework for caries classification, which showed that the accuracy of a federated model need not come at the cost of an explanation a clinician can read. The caries task is the natural starting point because the data are abundant, the labels are well defined and the clinical value is immediate, and a dental network can therefore expect a federated caries detector to be its fastest win.

For panoramic pathology, Maniraj et al.(6) reported Fed-ADD, a federated detector of dental pathologies from panoramic radiographs, and Dhas et al.(7) added uncertainty estimation to a privacy-preserving federated model for multi-institutional dental diagnosis, so that the model reports when it is unsure rather than issuing a confident error. The uncertainty addition matters in practice, because a clinician will trust a model more when it says it is unsure than when it is silently wrong, and it is the single feature that most improves the clinical acceptability of a federated detector. For oral cancer, Wu et al.(8) applied privacy-preserving federated learning to multi-site oral squamous cell carcinoma detection, and this is the task for which the privacy benefit is largest, because oral cancer data are scattered and sensitive. In each of these cases, the required trade-off was modest: the federated model lost a small amount of accuracy relative to a centralized one, and it gained the ability to train on data that no central server could lawfully hold.

The readiness of the dental tasks is the key finding, because it reverses a common assumption. A practitioner might assume that federated learning is a research technique too immature for clinical dental work, but the dental demonstrations of caries, panoramic pathology and oral cancer detection show the opposite, that the technique is past feasibility and into the stage where a network can adopt it for the tasks that matter. The remaining uncertainty is not whether federated learning works, but whether it keeps working when the participants change and the data drift, and that uncertainty is the reason the discussion returns repeatedly to deployment.

The same readiness is visible in the wider medical imaging literature, which supplies the evidence that a dental team can borrow. Bakas et al.(9) reviewed federated learning in healthcare from research to deployment and concluded that the technique is mature for imaging, and Shukla et al.(10) reached the same conclusion for computational pathology, while Li et al.(11) reviewed cardiovascular prediction and found consistent, near-centralized performance across clinical applications.

The cardiovascular case is instructive because it is a tabular, multi-site problem rather than

an imaging one, and its success shows that federated learning is not tied to a single data type, which matters for a dental network that also holds structured clinical fields alongside its images. What this means for dentistry is that the burden of proof has been met elsewhere, and a dental team does not need to re-derive the method, only to apply it to dental images. The risk of applying it is therefore a risk of integration rather than of invention, and an integration risk is one that a competent engineering team can retire, which is the most useful thing the medical evidence contributes to the dental decision.

The size of the accuracy trade-off is the first thing a practitioner asks about, and the evidence answers it consistently across studies. Across the dental and medical studies, federated training reached accuracy within a few percentage points of centralized training, and the gap narrowed as the number of participating sites grew, which is the expected behavior for a method that combines many small datasets into a single statistical pool.

The differential-privacy step added a further small loss, and the interpretability step added none, so the total cost of the recommended configuration, federated training with differential privacy and an explanation, is a few percentage points of accuracy against a large gain in the lawful reach of the training data. In practice, a clinic should weigh that few percentage points against the alternative, which is a model trained on a single clinic's data, and the single-clinic model is the less accurate of the two, so the trade-off is a net gain rather than a cost.

The breadth of the medical applications is itself evidence that the technique is general, and it is worth listing because a dental team can borrow the method from any of them. Federated learning has been demonstrated for stroke detection by Samuel et al.,⁽¹²⁾ for tuberculosis diagnosis by Bhattacharjee et al.,⁽¹³⁾ for skin lesion classification by Alfalahi et al.⁽¹⁴⁾ and Awais et al.,⁽¹⁵⁾ for breast cancer by Ahmed et al.,⁽¹⁶⁾ for liver cancer by Lou et al.,⁽¹⁷⁾ for thoracic disease by Zulqarnain et al.,⁽¹⁸⁾ for eye disease by Niro et al.,⁽¹⁹⁾ and for brain aging and dementia by Roca et al.⁽²⁰⁾ and Moniruzzaman et al.,⁽²¹⁾ and for lymph node metastasis from CT foundation models by Bhalla et al.⁽²²⁾ Each of these reported accuracy close to centralized training, and each required the same trade-off that a dental task would require, a small loss of accuracy in exchange for the ability to train on data that could not be pooled. The dental reader does not need to study each of these, but their number is the point: federated learning is not a niche technique, and the risk of applying it to dental images is low because the risk has already been paid elsewhere.

What can be added

Privacy is not all-or-nothing, and federated learning can be hardened with two techniques that the dental studies have begun to adopt. The reason this matters is that federated learning alone, without any added noise, is not a complete privacy guarantee, because the model updates that a clinic transmits can be inspected and, in some settings, reconstructed into the data that produced them, and a practitioner who believes otherwise is accepting a privacy claim that the security literature does not support. Differential privacy adds noise to the shared model updates so that no individual patient can be reconstructed from them, and Mohammadi et al.⁽²³⁾ set out the methods and trade-offs for medical deep learning, while Pavate et al.⁽²⁴⁾ described the convergence between differential privacy and federated learning in healthcare. The cost of differential privacy is a small accuracy loss, and the dental studies that added it, such as the caries framework of Kumar et al.,⁽⁵⁾ found the loss acceptable for the privacy gained. The noise parameter is itself a decision, because more noise means more privacy but less accuracy, and a clinic should choose the parameter with its data-protection officer rather than accept a vendor's default, because the default is chosen to protect the vendor, not the patient. Interpretability can also be preserved, and Dhas et al.⁽⁷⁾ and Kumar et al.⁽⁵⁾ both paired federated training with an explanation, which matters because a clinician will not adopt a privacy-preserving model that is also a black box. For the purpose of adoption, the combination of federated training, differential privacy and an explanation is therefore the configuration that a dental network should target, and it is achievable with the methods already published.

The workflow that a clinic follows is simple enough to state in a single paragraph, and stating it removes much of the mystery that surrounds the technique. The clinic keeps its data on its own machines and trains a local copy of the model, and it sends only the model's updated parameters to the aggregator, never the data that produced them, so the patient's record never leaves the building where it was made.

The aggregator averages the parameters it receives from many clinics and sends the averaged model back, and the cycle repeats, round after round, until the model stops improving. In practice, the clinic's burden is the computation of the local training step and the network cost of transmit-

ting the parameters, and both are within the reach of a modern dental practice that already runs imaging software on a daily basis. This simplicity is the practical advantage of federated learning, and it is the reason the technique can be adopted without a data scientist on every site, which is the very constraint that most small dental clinics face.

Beyond differential privacy and interpretability, a set of hardening techniques has been demonstrated that a dental network can adopt selectively. Homomorphic encryption allows the aggregator to combine encrypted model updates without ever seeing them, and Josephine Usha et al.(25) and Zuo et al.(26) applied it to multi-institutional imaging, at the cost of additional computation. Blockchain integration, described by Shahsavari et al.,(27) adds an audit trail to the aggregation, and swarm intelligence, demonstrated by SayedElahl et al.,(28) replaces the central aggregator with a peer-to-peer one, which removes the aggregator as a single point of trust.

Explainability can be preserved in a federated model, as Penzel et al.(29) and Mohanraj et al.(30) both showed, and pseudo-labeling, demonstrated by Montaha et al.,(31) reduces the need for labeled data across sites, while the federated nnU-Net of Skorupko et al.(32) shows that even segmentation, the most data-hungry task, can be federated. For the purpose of this review, these techniques are options rather than requirements, and the practitioner should choose the lightest combination that satisfies the privacy need, because each added technique has a cost in engineering and in accuracy.

What remains research

The gaps are not in feasibility but in evidence of deployment. No dental benchmark yet exists against which a federated model can be judged, and the field relies on medical benchmarks that do not capture dental anatomy, so a federated caries model is being compared against medical imaging tasks that do not share its tooth-level structure.

No long-term study has followed a federated dental model from training through clinical use, so the claims of near-centralized accuracy rest on research datasets rather than on deployed performance. Security is a further open area, because Faraoun et al.(33)demonstrated a backdoor attack on federated medical imaging, and Giouroukou et al.(1)catalogued the metadata and pixel-level re-identification risks that federated learning alone does not remove. The research base is also uneven in a way that a practitioner should know before investing.

The field is growing fast, as the bibliometric analysis of Ghaeli et al.(34)documents, and the review literature, from Popuri et al.(35)through Daram et al.(36)and Kothari Lakshmi Kasturi et al.,(37)has moved from describing the promise to cataloguing the deployments, but the growth is concentrated in a few imaging tasks and has not reached dentistry at scale. The privacy risks are real but bounded, and Dolezel et al.(38)reviewed them systematically, while Raggio et al.(39) showed that federated training can be combined with a privacy guarantee for a single task, and Shrimali et al.(40)examined how the choice of aggregation function changes what is learned and what leaks. The pretraining-to-privacy path that Jiang et al.(41)described for ultrasound and the privacy-preserving translation of Li et al.(42)show the direction the field is moving, and Soltanieh et al.(43)reviewed the same path in neurology. These studies define the frontier, and a dental team that adopts federated learning is adopting a moving frontier rather than a finished tool, which is itself a reason to start with the simplest task that the evidence supports.

The research agenda that would close these gaps is short and specific, and it is worth stating because a dental network that is deciding whether to invest wants to know what it is investing in. A dental benchmark is the first need, because it would let the field measure its progress and would let a buyer compare federated models on equal terms, and building it is a task for a dental research consortium rather than for any single clinic.

A deployment study is the second need, because it would replace the inference that federated learning will work in practice with the demonstration that it did, and it is a task for a vendor or a network that is already committed to a federated product. A security evaluation is the third need, because the backdoor and leakage risks are real but their dental-specific severity is unknown, and it is a task for the security research community in partnership with a dental network. Each of these is a concrete action with a named recipient, and they are the same recommendations that the discussion states from the practitioner's side, stated here from the research side.

These gaps are concrete, and each is paired with a recommendation in the discussion, because the review's purpose is to tell a practitioner what to do, not only what is known. The applications are listed in (Table 1), and the privacy techniques with their trade-offs in (Table 2).

Table 1. Federated learning applications in dental and medical imaging ordered by practical relevance.

Application	Dental or medical	Privacy feature	Readiness
Caries detection	Dental	Differential privacy, interpretability	Demonstrated
Panoramic pathology detection	Dental	Uncertainty, differential privacy	Demonstrated
Oral squamous cell carcinoma detection	Dental	Privacy-preserving multi-site	Demonstrated
Cardiovascular prediction	Medical	FL across sites	Transferable
Computational pathology	Medical	FL across sites	Transferable

Note: FL = federated learning.

Table 2. Privacy-preserving techniques and their trade-offs.

Technique	What it protects	Cost
Federated learning alone	Data remain local	Model updates can leak information
Federated learning + DP	Adds noise to updates	Small accuracy loss
Federated learning + DP + interpretability	Adds explanation	Additional engineering
Synthetic data	Replaces real records	Loses rare findings

Note: DP = differential privacy.

DISCUSSION

The practical implication comes first, because it is the one a dental network can act on this year. A dental network that wants to train a shared model without pooling patient data should start with a federated caries or panoramic pathology detector, because those are the tasks for which the dental evidence is strongest and the accuracy loss relative to centralized training is smallest. The recipient of this implication is the dental network or the vendor that serves it, and the action is specific: select a federated framework, apply differential privacy at the update step, and add an explanation so the model can be audited. This implies a further decision, that the network should not wait for a perfect dental benchmark before starting, because the medical benchmarks are sufficient to begin and the dental benchmark can be built in parallel.

The comparison with the wider literature is oriented to that decision. Bakas et al.(9) reviewed healthcare federated learning and found the method mature for imaging, and Mir et al.(44) reviewed the ethics and concluded that federated learning is a privacy-preserving and equitable route but not a legal compliance mechanism, which is the point a regulator needs to hear. The dental reviews of Rischke et al.(45) and the data-sharing analyses of Brinz et al.(2) and Uribe et al.(3) all converge on the same conclusion: federated learning solves the technical problem of pooling data, but it does not remove the legal and governance problem of sharing it, and the two must be handled together.

This contrasts with the earlier enthusiasm of Schwendicke et al.,(46) who framed the chances and challenges of dental artificial intelligence before federated learning was widely demonstrated, and it is a contrast between a promise and its partial fulfilment. The differential-privacy literature sharpens the point, because Mohammadi et al.(23) and Pavate et al.(24) both showed that the privacy guarantee of federated learning is only as strong as the noise added to it, and Dolezel et al.(38) reviewed the residual risks systematically, while Giouroukou et al.(1) and Raggio et al.(39) demonstrated that the privacy of the pipeline, not just the model, is what a regulator must audit. The general reviews of Daram et al.,(36) Kothari Lakshmi Kasturi et al.(37) and Popuri et al.(35) repeat the same sequence from promise to deployment, and Ghaeli et al.(34) documents that the field's growth has outpaced its dental penetration.

The comparison between the four alternatives can be reduced to a single decision, and it is the decision the whole review serves. Centralized training offers the highest accuracy but is unlawful or impractical for pooled patient data, local training is lawful but underpowered, synthetic data is lawful but loses rare findings, and federated learning is lawful and retains the rare findings at a small accuracy cost.


In practice, the choice is therefore not between four equal options but between federated learning and a lesser alternative, and the lesser alternative is local training, which federated learning dominates on every axis except simplicity. For the purpose of this review, the practitioner should choose federated learning when the privacy of the data is a hard constraint, which in dentistry it

always is, and should choose a simpler method only when the task is so narrow that a single clinic's data already suffices, which is rare.

For practitioners, the implications are concrete and each names its recipient. A clinic asked to join a federated network should require, before it agrees, that the aggregator uses differential privacy and that the shared model includes an explanation, because those are the two features that turn a privacy claim into an auditable one, and the clinic's data-protection officer is the person who should verify them before the clinic commits its patients to the network. A vendor selling dental artificial intelligence should state which of its models were trained federated, because the buyer has a right to know whether the accuracy figure was earned on pooled or on distributed data, and the buyer's clinical director is the recipient of that statement, which should be made in the product documentation rather than only on request. A regulator evaluating a federated system should treat federated learning as a technical complement to, not a substitute for, a lawful basis for processing, and should ask whether the system's metadata, not just its pixels, are protected, because Giouroukou et al.(1) showed that metadata leaks identity even when the images are de-identified, and the audit should therefore follow the metadata as closely as it follows the pixels.

For the purpose of prioritization, the evidence ranks the tasks. Caries and panoramic pathology are the tasks to deploy first, oral cancer the task to invest in next because its privacy benefit is highest, and rare dental findings the tasks where synthetic data should not be used as a substitute, because synthetic data loses the rare findings that federated learning preserves. The decision figure presents this ranking as a flow (Figure 1), so that a practitioner who must choose a starting point can follow it to a decision without reading the whole review. The ranking is not a matter of preference but of evidence density, and a dental network that departs from it should know that it is departing from the strongest part of the literature into a part that is still being written.

Figure 1. Decision flow for choosing a privacy-preserving training approach in dental artificial intelligence.



The evidence density itself is a fact about the field, not a value judgment, and it is the reason the review can be confident about caries while being cautious about rare lesions, and a practitioner who reads only the confident parts of this review will still have a workable plan.

A further practical implication concerns the recurring costs, and it names the people who must carry them. Federated learning shifts the cost of data from movement to computation, because every participating clinic must train a local model and transmit its updates, and the clinic's information-technology staff should budget for that computation before the clinical staff is promised a working model. The aggregator, whether it is a vendor or a lead hospital, must maintain the security of the aggregation step, because Faraoun et al.(33) showed that a compromised aggregator can be used to plant a backdoor, and the aggregator's security officer is the person who owns that risk. The clinical lead who signs off on a federated project should require, in the project plan, a named person for each of these responsibilities, because a privacy-preserving technique that no one is responsible for is not preserved in practice.

The vendor perspective is distinct from the clinic perspective, and it is worth stating because the two are often conflated. A vendor that sells a centralized dental model has an incentive to keep training data under its own control, because the data is the asset, whereas a vendor that sells a federated service earns from the service rather than from the data, and the two business models produce different incentives for privacy. For practitioners, this implies that a clinic should ask a vendor whether its model was trained federated and, if not, why not, because the answer reveals whether the vendor's interest is in the clinic's data or in the clinic's problem, and the clinic's procurement lead is the person who should ask.

For regulators, the same distinction implies that a privacy audit should follow the data rather than the marketing, because a product described as privacy-preserving may still collect data centrally under a different name, and the audit is the only way to tell. This is a concrete action with a recipient, and it is the implication that ties the technical evidence to the commercial reality in which a dental network actually operates.

The equity dimension deserves its own statement, because it is the reason federated learning matters beyond privacy. A centralized dental model is trained on whichever clinics can share data, and those clinics are disproportionately large, well-resourced and urban, so the model inherits their patients and underperforms on everyone else. Federated learning lets a small rural clinic contribute to a shared model without surrendering its data, and Mir et al.(44) identified this as one of the principal equitable benefits of the technique in their ethics review.

For practitioners, the implication is that a dental network should recruit the small and the rural clinics first, because they are the ones whose data are most valuable and most underused, and the network's outreach lead is the person who should recruit them. This is a concrete action with a recipient, and it is the one implication of this review that is about justice rather than about efficiency, but it follows from the same evidence and the same logic.

The limitations are integrated with their recommendations, because a limitation without a remedy is not useful to a practitioner. The first limitation is that the dental evidence is small and recent, and most of the transferable evidence is medical, so a dental team should treat the medical results as a floor rather than a ceiling and should validate the federated model on its own images before trusting it on patients.

The second limitation is that no dental benchmark exists, so a dental network should contribute to building one, because a field without a benchmark cannot tell whether it is improving, and the network's research lead is the person who should coordinate that contribution. The third limitation is that the security of federated learning is not settled, as the backdoor attack of Faraoun et al.(33) shows, so a deployed system should include anomaly detection at the aggregator and should not assume that federated learning is inherently secure, and the aggregator's security officer should own that assumption explicitly.

The fourth limitation is that the long-term behavior of a federated dental model, as new clinics join and leave and as the data drift, has not been studied, so a deployed system should log its accuracy over time and revalidate when the participant set changes, and the clinical lead should review that log. Each of these is a concrete action with a named recipient, and together they define what the field should do next.

CONCLUSIONS

Federated learning is the right answer to the privacy problem in dental artificial intelligence, and it is ready for the tasks that matter most, caries detection and panoramic pathology detection, where it reaches accuracy close to centralized training while keeping patient data local and secure.

The combination that a dental network should adopt, namely federated training with differential privacy and an explanation, is achievable with published methods, and the medical imaging literature supplies the transferable evidence that the dental field can borrow. The workflow is simple enough for a small clinic to join, and the cost is a few percentage points of accuracy against a large gain in the lawful reach of the data.

What the field still lacks is a dental benchmark and long-term deployment evidence, and each gap has a concrete remedy: build the benchmark, run the deployment studies, and add anomaly detection to the aggregator.

Federated learning protects privacy but does not replace legal compliance, and the two must be managed together by the clinic, the vendor and the regulator, each of whom has a specific and named action to take.

A clinic should verify the privacy features before it joins a network, a vendor should disclose

how its model was trained, and a regulator should audit the data rather than the marketing, and if each of the three does its part faithfully, the promise of private, collaborative dental artificial intelligence becomes a working tool rather than a technical aspiration.

REFERENCES

1. Giouroukou K, Marias K, Tsiknakis M, Klontzas ME. Rethinking Privacy in Medical Imaging AI: From Metadata and Pixel-Level Identification Risks to Federated Learning and Synthetic Data Challenges. *Radiology: Artificial Intelligence*. 2026;8(1):e250273. doi: 10.1148/ryai.250273
2. Brinz J, Eslamiamirabadi N, Salamati A, Tresp V, Schwendicke F, Tichy A. Data sharing for responsible artificial intelligence in dentistry: a narrative review of legal frameworks and privacy-preserving techniques. *Journal of Dentistry*. 2025;163:106130. doi: 10.1016/j.jdent.2025.106130
3. Uribe SE, Schwendicke F. Data Foundations for Trustworthy AI in Dentistry. In: *Artificial Intelligence for Oral Health Care*. Springer Nature Switzerland; 2025. p. 151-163. doi: 10.1007/978-3-031-84047-0_10
4. Gopinath B, Priyadharsan JH, Ragasudhan R, Sanjay S. Development of an Automated Dental Caries Detection System using Federated Learning in X-Ray Images with Data Security. 2025 3rd International Conference on Advancements in Electrical, Electronics, Communication, Computing and Automation (ICAECA). 2025:1-5. doi: 10.1109/icaeca63854.2025.11012522
5. Kumar J, Almustafa KM, Madanat R, Sharma AK, Sutcu M, Katrib J. Privacy-aware and interpretable deep learning framework for dental caries classification. *Intelligence-Based Medicine*. 2025;12:100294. doi: 10.1016/j.ibmed.2025.100294
6. Maniraj K, Kumar MS. Fed-ADD: Federated Deep Learning for Automated Detection of Dental Pathologies from Panoramic Radiographs. 2025 First International Conference of Advances in Engineering and Computing Technologies for Sustainable Development (AECTSD). 2025:1-7. doi: 10.1109/aectsd65988.2025.11411534
7. Dhas JTM, Surendra D, Mayil VV. Uncertainty Aware Privacy Preserving Federated Deep Learning for Trustworthy Multi Institutional Dental Disease Diagnosis. 2026 5th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 50. 2026:1-6. doi: 10.1109/otcon69332.2026.11629829
8. Wu Y, Chang H. Automatic detection of multi-site oral squamous cell carcinoma based on privacy-preserving federated learning. *BMC Oral Health*. 2026;26(1):989. doi: 10.1186/s12903-026-08388-2
9. Bakas S, Li X, Shah P, Roth HR. Federated Learning in Healthcare: From Research to Real-World Deployment. *Annu Rev Biomed Eng*. 2026;28(1):163-186. doi: 10.1146/annurev-bioeng-080125-041414
10. Shukla S, Doyle S. Federated learning in computational pathology: a literature review. *J Med Imag*. 2025;12(06):1-26. doi: 10.1117/1.jmi.12.6.061412
11. Li J, Xiang W, Shang D, Li S, Li Q. Federated learning for cardiovascular disease prediction: a systematic review of clinical applications, validation, and translation readiness. *Front Cardiovasc Med*. 2026;13:1831342. doi: 10.3389/fcvm.2026.1831342
12. Samuel R, Pandi T. Federated deep learning for stroke detection: a privacy-preserving, explainable CNN-LSTM framework. *Front Neurosci*. 2026;20:1847277. doi: 10.3389/fnins.2026.1847277
13. Bhattacharjee S, Sapkal V, Sharma VK, Rewatkar A, Khidse SV, Ingole PK. Federated learning for privacy-preserving multi-center tuberculosis diagnosis using chest imaging data. *Indian Journal of Tuberculosis*. 2026;73(3):377-385. doi: 10.1016/j.ijtb.2025.11.011
14. Alfalahi MAM, Karan O, Kurnaz S, Türkben AK. Federated learning for privacy-preserving skin cancer classification using deep neural networks. *Front Oncol*. 2026;16:1720748. doi: 10.3389/fonc.2026.1720748
15. Awais M, Junejo RH. FedSwin-LHTP: Structure-Aware Hessian-Inspired Token Pruning for Efficient Federated Skin Lesion Classification. *Diagnostics*. 2026;16(16):2637. doi: 10.3390/diagnostics16162637
16. Ahmed F, Sánchez D, Haddi Z, Domingo-Ferrer J. FSCL-BC: Federated supervised contrastive learning for breast cancer diagnosis with high sensitivity. *Computer Methods and Programs in Biomedicine*. 2026;285:109500. doi: 10.1016/j.cmpb.2026.109500
17. Lou L, Govindarajan V, Shaikh ZA, Liu R, Ayadi M, Wang H, et al. FedLiverNet: a federated learning framework for privacy-preserving and efficient liver cancer detection. *Sci Rep*. 2026;16(1):16745. doi: 10.1038/s41598-026-46244-z
18. Zulqarnain M, Hussain SJ, Aslam MZ, Fiaz A, Islam M. Federated Learning for Thoracic Disease Classification Using Convolutional Neural Networks and Differential Privacy. *Healthcare Tech Letters*. 2026;13(1):e70080. doi: 10.1049/htl2.70080
19. Niro F, Renzo MD, Agnello P, Petyx M, Martinelli F, Maddalena M, et al. A Federated Machine Learning Approach for the Detection and Visualisation of Eye Diseases Using Activation Maps. *Sensors*. 2026;26(16):5288. doi: 10.3390/s26165288
20. Roca V, Tommasi M, Andrey P, Bellet A, Schirmer MD, Henon H, et al. Federated learning for MRI-based BrainAGE: A multicenter study on post-stroke functional outcome prediction. *NeuroImage*. 2026;339:122165. doi: 10.1016/j.neuroimage.2026.122165
21. Moniruzzaman M, Uddin MS, Ahmed A, Aktarujjaman M, Ahmed M, Rahman A. FEDI-CODE: A federated and causally-informed framework for dementia risk prediction using multi-site patient data. *PLoS One*. 2026;21(6):e0351957. doi: 10.1371/journal.pone.0351957
22. Bhalla P, Gaviria DD, Kupczyk P, Hosseini ASA, Conradi L, Fehrenbach U, et al. Federated CT foundation models for multi-center detection of lymph node metastasis in pancreatic cancer. *Sci Rep*. 2026;16(1):12051. doi: 10.1038/s41598-026-47631-2
23. Mohammadi M, Vejdanihemmat M, Lotfinia M, Rusu M, Truhn D, Maier A, et al. Differential privacy for medical deep learning: methods, tradeoffs, and deployment implications. *npj Digit Med*. 2026;9(1):93. doi: 10.1038/s41746-025-02280-z

24. Pavate A. Privacy-preserving federated learning with differential privacy for healthcare AI: a convergence and utility analysis. *Journal of Artificial Intelligence Machine Learning and Neural Network*. 2025;69. doi: 10.55529/jaiml.52.69.78
25. Josephine Usha L, Saranya KR, Suganya Y, Mallikka R. Adaptive homomorphic federated learning framework for multi-institutional medical imaging with optimized diagnostic accuracy. *Sci Rep*. 2026;16(1):20599. doi: 10.1038/s41598-026-45821-6
26. Zuo A, Feng Z, Ping Y, Tao S, Sun H, Chen Y. FedGraphHE: A privacy-preserving federated graph neural network framework with dynamic homomorphic encryption and robust aggregation. *PLoS One*. 2026;21(1):e0339881. doi: 10.1371/journal.pone.0339881
27. Shahsavari Y, Baseri Y, Hafid A, Dambri OA, Makrakis D. Integration of Federated Learning and Blockchain in Health Care: Tutorial on Medical Data, Architectures, Privacy, Security, and Regulatory Compliance. *J Med Internet Res*. 2026;28:e80178-e80178. doi: 10.2196/80178
28. SayedElahl MA, Farouk RM, Ali AE, Ahmed E. Federated learning with swarm intelligence for efficient and secure medical image analysis. *Sci Rep*. 2026;16(1):14734. doi: 10.1038/s41598-026-50882-8
29. Penzel N, Scheliga D, Oppermann H, Mäder P, Hauelsen J, Denzler J, et al. Model utility and explainability in federated learning - A case study in healthcare using fundus oculi datasets. *Journal of Biomedical Informatics*. 2026;177:105010. doi: 10.1016/j.jbi.2026.105010
30. Mohanraj S, Radhakrishnan S. FuzzyFed-CNN: secure and explainable multimodal federated learning for early Alzheimer's diagnosis. *Front Artif Intell*. 2026;9:1852196. doi: 10.3389/frai.2026.1852196
31. Montaha S, Rahman R, Ghosh T, Sheikhi F, Maleki F. Federated Pseudo-Labeling: A Data-Centric, Privacy-Preserving Framework for Medical Image Segmentation. *IEEE J Biomed Health Inform*. 2025;29(12):8823-8830. doi: 10.1109/jbhi.2025.3622934
32. Skorupko G, Avgoustidis F, Martín-Isla C, Garrucho L, Kessler DA, Pujadas ER, et al. Federated nnU-Net for privacy-preserving medical image segmentation. *Sci Rep*. 2025;15(1):38312. doi: 10.1038/s41598-025-22239-0
33. Faraoun H, Bellafqira R, Coatrieux G, Kallas K. A Durable Backdoor Attack on Medical Imaging via Federated Learning. In: *Studies in Health Technology and Informatics*. IOS Press; 2026. p. 600-604. doi: 10.3233/shti260241
34. Ghaeli R. Federated learning for healthcare: A bibliometric analysis of privacy-preserving machine learning applications in medical imaging, electronic health records, and clinical diagnostics based on 200 highly cited publications (2018–2025). *105267/jhe*. 2026;2(4):217-242. doi: 10.5267/j.ijdns.2026.38
35. Popuri V. Securing Healthcare Data: Federated Learning for Privacy-Preserving AI in Medical Applications. *IJMT*. 2024;11(3):64-82. doi: 10.37745/ijmt.2013/vol11n36482
36. Daram S. Federated Learning in Medical AI: Advancing Privacy-Preserving Data Sharing for Collaborative Healthcare Research. *IJAIDSM*. 2025;6:65-72. doi: 10.63282/3050-9262.ijaidsm-v6i2p107
37. Kothari Lakshmi kasturi, Jajam Naresh. Federated Learning in Healthcare: A Privacy-Preserving Approach to Medical AI. *IJIS*. 2025;12-16. doi: 10.63856/fm5jw605
38. Dolezel D, Lalani K, Watzlaf V, Butler-Henderson K, Lambert E, Morton M, et al. AI Privacy and Security in Healthcare: A Systematic Literature Review. *Int J Telerehabil*. 2026;18(1):6747. doi: 10.63144/ijt.2065.6747
39. Raggio CB, Zaffino P, Spadea MF. A privacy-preserving federated learning framework for generalizable CBCT to synthetic CT translation in head and neck. *Front Digit Health*. 2026;8:1812254. doi: 10.3389/fdgth.2026.1812254
40. Shrimali B, Geddam R, Ghayvat H, Shah R, Ginglani M. Impact of Aggregation Functions and Learning Settings in Privacy-Preserving Federated Learning for Skin Cancer Detection. *Healthcare Tech Letters*. 2026;13(1):e70096. doi: 10.1049/htl2.70096
41. Jiang Y, Feng CM, Ren J, Wei J, Zhang Z, Hu Y, et al. From pretraining to privacy: federated ultrasound foundation model with self-supervised learning. *npj Digit Med*. 2025;8(1):714. doi: 10.1038/s41746-025-02085-0
42. Li W, Li Z, Shi Y, Lam S, Xiao L, Cheung AHY, et al. Privacy-preserving Virtual Contrast-enhanced MRI for Nasopharyngeal Carcinoma: A Multicenter Study. *International Journal of Radiation Oncology*Biophysics*. 2026. doi: 10.1016/j.ijrobp.2026.06.3057
43. Soltanieh S, Khalvati F, Yeh EA. Federated Learning in Neurology: Bridging Data Privacy and Artificial Intelligence for Brain Health. *Semin Neurol*. 2026;46(01):038-048. doi: 10.1055/a-2769-6752
44. Mir BA, Abbas SR, Lee SW. Federated Learning in Healthcare Ethics: A Systematic Review of Privacy-Preserving and Equitable Medical AI. *Healthcare*. 2026;14(3):306. doi: 10.3390/healthcare14030306
45. Rischke R, Schneider L, Müller K, Samek W, Schwendicke F, Krois J. Federated Learning in Dentistry: Chances and Challenges. *J Dent Res*. 2022;101(11):1269-1273. doi: 10.1177/00220345221108953
46. Schwendicke F, Samek W, Krois J. Artificial Intelligence in Dentistry: Chances and Challenges. *J Dent Res*. 2020;99(7):769-774. doi: 10.1177/0022034520915714

FUNDING

No financing

CONFLICT OF INTEREST

None

AUTHORSHIP CONTRIBUTIONS

Drafting – original draft: Ana Rodríguez Yanes.
Writing–review and editing: Ana Rodríguez Yanes.