



Generative artificial intelligence for synthetic dental data: opportunities, privacy, and validation

Inteligencia artificial generativa para datos dentales sintéticos: oportunidades, privacidad y validación

Nairobi Hernández Bridon¹  

¹ Instituto de Investigaciones en Microbiología y Parasitología Médica. Consejo Nacional de Investigaciones Científicas y Técnicas (CONICET). Universidad de Buenos Aires. Ciudad Autónoma de Buenos Aires, Argentina

Submitted: 02-08-2025 | Revised: 30-09-2025 | Accepted: 20-11-2025 | Published: 26-11-2025

Corresponding Author: Nairobi Hernández Bridon 

ABSTRACT

Generative artificial intelligence has moved from producing plausible but imperfect dental images to a serious instrument for augmenting scarce datasets, and its trajectory in dentistry can be read in three stages. In the first stage, generative adversarial networks entered dental imaging as a tool for artifact correction and denoising, where the images were a means to a cleaner input rather than an end in themselves. In the second stage, the models turned to synthesis, producing synthetic panoramic, intraoral, and cephalometric images that augmented small datasets and addressed class imbalance, and the dental literature began to evaluate whether synthetic images improve downstream models. In the third, present stage, diffusion models raised the fidelity of synthetic radiographs, and the discussion shifted from whether images can be generated to whether they are private and whether they are valid, two questions the earlier literature largely ignored. The opportunities, augmentation, imbalance correction, rare pathology, and privacy-preserving data sharing, are real, but the privacy of synthetic dental data is not automatic, because generative models can memorize training examples, and the validation of synthetic data, that it is faithful, useful, and safe, is still in its infancy. The review concludes that generative dental data is a promising and rapidly maturing tool whose clinical and regulatory value now depends less on generating better images than on demonstrating that the images are safe to share and trustworthy to train on.

Keywords: artificial intelligence; generative models; synthetic data; dental imaging; privacy; validation.

RESUMEN

La inteligencia artificial generativa ha pasado de producir imágenes dentales plausibles pero imperfectas a convertirse en un instrumento serio para aumentar conjuntos de datos escasos, y su trayectoria en odontología puede leerse en tres etapas. En la primera, las redes generativas antagónicas entraron en la imagen dental como herramienta de corrección de artefactos y reducción de ruido, donde las imágenes eran un medio para una entrada más limpia. En la segunda, los modelos se orientaron a la síntesis, produciendo imágenes panorámicas, intraorales y cefalométricas sintéticas que aumentaron conjuntos pequeños y corrigieron el desequilibrio de clases. En la tercera, actual, los modelos de difusión elevaron la fidelidad de las radiografías sintéticas y el debate pasó de si se pueden generar imágenes a si son privadas y si son válidas. Las oportunidades son reales, pero la privacidad de los datos dentales sintéticos no es automática, porque los modelos generativos pueden memorizar ejemplos de entrenamiento, y su validación, que sean fieles, útiles y seguros, está aún en su infancia. La revisión concluye que los datos dentales generativos son una herramienta prometedora cuyo valor clínico y regulatorio depende hoy menos de generar mejores imágenes que de demostrar que son seguras para compartir y fiables para entrenar.

Palabras clave: inteligencia artificial; modelos generativos; datos sintéticos; imagen dental; privacidad; validación.

INTRODUCTION

The shortage of labeled data is the quiet constraint on nearly every deep-learning advance in dental imaging. A model that detects caries, segments the mandibular canal, or classifies a lesion is only as good as the annotated images it was trained on, and dental datasets are small, private, and skewed toward common conditions, while the rare pathologies and the atypical presentations that a diagnostic model most needs to learn are exactly the ones least represented.(1)

The standard remedies, data augmentation, transfer learning, and multi-institution collaboration, help but do not remove the underlying scarcity, and the sharing of dental images is constrained by privacy, because radiographs and intraoral photographs are, or can be made to be, personally identifiable.(2)

The two constraints, scarcity and privacy, pull in opposite directions: the field needs more data to train better models, and it needs to share less data to protect patients, and this tension is what makes synthetic data so attractive, because it appears to dissolve the conflict by producing data that is both abundant and, in principle, non-personal. Generative artificial intelligence entered this problem with that specific promise: instead of collecting more data, synthesize it. A generative model trained on a small set of dental images can, in principle, produce new images that carry the statistical structure of the originals without copying any individual patient, and the resulting synthetic data can augment training, correct class imbalance, and be shared where the real data cannot.(3)

This promise has driven a decade of work whose trajectory, read chronologically, reveals as much about the field's changing priorities as about its technical progress, and it is that trajectory, rather than any single model or paper, that this review sets out to read, because a trajectory read as a whole shows the field what it has learned and what it has, so far, only assumed.

The promise is worth stating carefully, because its three components, augmentation, imbalance correction, and privacy-preserving sharing, are distinct claims that require distinct evidence, and the field has not always kept them separate. Augmentation requires showing that a downstream model trained with synthetic data performs better; imbalance correction requires showing that the rare class, in particular, is better learned; and privacy-preserving sharing requires showing that the synthetic data does not leak the real patients, and the third of these is, of the three, the one the imaging literature has least often attempted to demonstrate.

The story begins with the generative adversarial network, introduced in 2014 as a way to train a generator and a discriminator against one another, and its rapid adoption in medical imaging, where it was first used not to create data but to clean it, because the model's ability to learn a data distribution was first turned, conservatively and understandably, toward restoring what noise and artifacts had degraded rather than toward inventing what had never existed.(4,5)

In dentistry, the early applications were artifact correction and denoising of computed tomography, tasks in which the generative model was a means to a better input rather than a source of new examples.(6,7)

From there the field moved, in a second stage, to synthesis proper: generating panoramic, intraoral, and cephalometric radiographs that could augment small datasets, and the dental literature began to ask the question that matters, whether synthetic images actually improve a downstream model.(8-10)

The present, third stage is defined by two shifts: the arrival of diffusion models, which produce more realistic images than their predecessors, and a turn toward the questions the earlier enthusiasm skipped, the privacy of synthetic data and the standards by which it should be validated.(11,12)

This review traces that trajectory and argues that the field's center of gravity has moved from generation to governance, a move that is less glamorous than the technical progress that preceded it but is the more important of the two, because it is the move that will determine whether the generated images are ever used outside the laboratory, in the clinic, the classroom, or the shared dataset, where their value and their risk both become real.

The review is organized chronologically because the technology's evolution is inseparable from its applications, and because the inflection points, the move from cleaning to synthesizing, and from synthesizing to validating, mark real changes in what the field is trying to accomplish, changes that a topical organization would flatten into a single undifferentiated body of work.

Three periods are delimited by explicit criteria: 2014 to 2018, the emergence of generative adversarial networks and their first medical and dental uses for correction; 2019 to 2022, the turn to synthesis and augmentation and its evaluation; and 2023 to 2025, the diffusion era and the turn to privacy and validation.

The periodization is not arbitrary: each boundary is marked by a concrete change in the lite-

ture, the first by the appearance of dental synthesis studies, the second by the arrival of diffusion models and of the privacy-and-validation framework, and the boundaries are therefore reproducible by a reader who examines the same sources. A periodization of this kind necessarily simplifies, and a study may straddle a boundary or anticipate a later development, but the simplification is useful precisely because it makes the trajectory legible, which is the first step toward judging it.

The purpose is to show not only what has been achieved but what has been overlooked, and to identify what the synthetic dental data of the future will need to demonstrate before it can be trusted in the clinic and before it can be shared where the real data cannot be.

DEVELOPMENT

2014-2018: the emergence of generative adversarial networks and their first dental uses

The generative adversarial network was introduced in 2014, and its defining feature, a generator and a discriminator trained adversarially, made it the first deep model capable of producing images that were visually plausible rather than merely statistically similar to their inputs, a capability that immediately attracted medical imaging.(4)

The early surveys of the field, and the first taxonomies of its variants, document a technology maturing rapidly in generality, and the broader medical-imaging surveys that followed show the adversarial network being taken up across ophthalmology, histopathology, and radiology within a few years of its introduction, which is the context in which the dental applications appeared.(4,13,14)

Its first dental appearances were conspicuously narrow, however: the generative model was used not to create new dental images but to correct existing ones. The clearest example is low-dose dental computed tomography, where a Wasserstein generative adversarial network was trained to remove the artifacts of reduced radiation, improving image quality without increasing dose, and where a transfer-learning generative network denoised dental computed tomography by learning the noise distribution rather than the signal.(6,7)

The half-scan artifact correction work that followed pushed the same idea further, using the generative network to reconstruct a clean image from a partial acquisition, a task in which the model is asked to fill in structure it has never seen, which is, in retrospect, the seed of the synthesis that the next period would embrace.(7)

In these applications the generative model was a sophisticated filter, and the evaluation was appropriately modest, image quality, dose, and diagnostic readability, rather than any claim about synthetic data as such. The period's narrowness is, in hindsight, its virtue: by asking the model only to correct an image rather than to invent one, the early studies could be evaluated against the original, uncorrupted image, a clean reference standard that the later, generative applications would lack.

The significance of this first period is that it established the generative model's credibility in dental imaging while leaving its data-generating potential almost entirely untapped, and it set the stage for the shift that followed by showing that the models could learn the structure of dental images well enough to manipulate it.

2019-2022: the turn to synthesis and augmentation

The second period is defined by a change of purpose: the generative model ceased to be a filter and became a factory. The dental literature in these years began to produce synthetic images for their own sake, panoramic, intraoral, and cephalometric radiographs generated to augment small datasets, and the studies cluster around two motivations, class imbalance and privacy.

The intraoral image generation of Kokomoto and colleagues, using progressive growing of a generative adversarial network and evaluated by dentists, made the privacy motivation explicit, arguing that synthetic images could substitute for real patient photographs in teaching and training.(8)

The cephalometric generation work of Kim and colleagues extended the same idea to a harder, higher-resolution target and evaluated the realism of the output, while the panoramic refinement work of Kim and colleagues used the model in a hybrid role, cleaning a real image rather than inventing a new one.(9,15)

The restoration and prosthesis literature took synthesis in a functional direction, with conditional generative networks designing crowns, inlays, and occlusal surfaces directly, where the generated output is not an image for a human to read but a geometry for a machine to mill, and

the comparative studies of this line, which measured time efficiency, occlusal morphology, and internal fit, established that the generated restorations were not merely plausible but clinically usable.(16-19)

This functional line is worth distinguishing from the image-synthesis line, because it has a different and in some ways easier validation problem: a generated restoration either fits or it does not, and the fit can be measured, whereas a generated radiograph is judged by a reader or by a downstream model, and its value is harder to pin down. The two lines therefore matured at different speeds, with the restoration line reaching clinical usability earlier, and the distinction foreshadows the third period's recognition that validation is task-specific and cannot be reduced to a single realism score, a recognition that is, in effect, the intellectual content of the field's current turn.

The scope of synthesis widened further in the same years: cephalogram synthesis from cone-beam computed tomography produced the two-dimensional views that three-dimensional imaging lacks, tooth segmentation of three-dimensional scans used a generative model to overcome the scarcity of annotated meshes, and dental-based human identification used a generative network to address the small-sample problem in a forensic setting.(20-22)

The orthopantomogram augmentation pipeline of Waqas and colleagues, and the narrative review of generative design for restorations, round out the period by showing that synthesis had become a general-purpose tool across nearly every dental image and geometry type.(23,24) The unsupervised tooth-segmentation work of Arian and colleagues, which used synthetic data for domain adaptation rather than augmentation, foreshadows the privacy-motivated uses that the third period would make central.(25)

The augmentation question, whether synthetic images improve a downstream diagnostic model, was tested most directly in the periapical classification study of Yang and colleagues, which generated images to train a classifier for C-shaped canals and evaluated the gain, and in the periodontal inpainting work of Kearney and colleagues, which used a generative inpainting network to fill missing structure and improve clinical-attachment-level prediction.(10,26)

These two studies are the period's methodological high point, because they answer the question that the rest of the period's literature merely posed, and their modest, mixed results are a more durable contribution than any demonstration of realism. The inflection point of this period is the appearance of the first systematic review of generative adversarial networks in dental imaging, which consolidated the scattered applications and made the field legible, and the appearance of a general survey of generative adversarial networks in medical imaging that placed dentistry within the broader medical trend.(27,28)

By the end of the period the technology had been applied to nearly every dental image type, but the evaluation had lagged the generation: most studies demonstrated that images could be made, and only a minority demonstrated that the images helped.

The distinction between these two demonstrations is the period's lasting contribution, because it is the first appearance in the dental literature of the question, now central, of whether a synthetic image is useful rather than merely realistic, and the studies that answered it, by training a downstream classifier on synthetic data and measuring the change, established the template that the validation literature of the third period would formalize.(10,29)

2023-2025: diffusion models, privacy, and validation

The third period is defined by two changes that are related but distinct. The first is technical: diffusion models, which generate images by iteratively denoising from random noise, arrived in dental imaging and produced synthetic radiographs of a fidelity that the adversarial networks had not reached, including work showing that a diffusion model trained on a limited dataset of panoramic radiographs can generate realistic images for augmentation, and coupled diffusion models for metal-artifact reduction in cone-beam computed tomography.(11,30)

The diffusion approach has a specific advantage that explains its rapid adoption: it is more stable to train than an adversarial network and less prone to the mode collapse that limited the earlier models, and these properties matter more in dental imaging, where the datasets are small and the cost of a failed training run is high, than they do in the large-scale domains where diffusion models were first demonstrated.

The tooth-generation and panoramic-generation systems of this period, with disentangled control over anatomy, mark a further step toward controllable, high-fidelity synthesis, and the Pano-GAN work showed that a generative model can produce entire panoramic radiographs of sufficient realism to be evaluated against clinical readers.(31,32)

The bitewing-segmentation work of Tolstaya and colleagues, which used synthetic data to

improve a segmentation model, and the synthetic-orthopantomogram augmentation of Waqas and colleagues, both belong to this period and both measure the downstream benefit rather than the visual quality, which is the sign that the utility question had taken hold.(23,29)

The second change is thematic: the field began to take its own premise seriously, and to ask whether synthetic dental data is actually private and actually valid, questions that the earlier literature had answered by assumption.(1,33)

The synthetic-panoramic study of Schoenhof and colleagues is the clearest marker of the shift, because it claims explicitly that generative networks can produce images “free from personal data” and then, in effect, invites the question of whether that claim has been demonstrated, and the more careful treatments that followed show that it has not been, at least not in general.(33)

The privacy question turns on a well-documented property of generative models, that they can memorize and reproduce their training examples, which means a synthetic image is not automatically non-personal, and the dental literature’s early claim that synthetic images are “free from personal data” is, at best, an overstatement that the privacy literature corrects.(33,34)

The replica-detection and membership-inference methods that the neuroimaging and general medical literatures have developed make the privacy question operational rather than rhetorical, by testing directly whether a generated image is a near-copy of a training example, and these methods are transferable to dental images without modification.(35)

The validation question is the mirror image: a synthetic image that is realistic but not faithful, or that is faithful but not useful, is of limited value, and the evaluation frameworks developed in the broader synthetic-data literature, which measure fidelity, utility, and privacy together, are only beginning to reach dental imaging.(12,36)

The systematic comparison of synthetic and real data across observational studies, which found that analyses on synthetic data can reproduce real-data findings but can also diverge, is the cautionary result that grounds the validation question in evidence rather than in principle.(36)

The systematic reviews of synthetic data across medical modalities that appeared in this period converge on the same diagnosis, that generation has outpaced validation, and that the field’s next step is to close that gap.(37,38) The longitudinal and time-series synthetic-data reviews extend the same conclusion beyond imaging, and the fact that the diagnosis is shared across imaging and non-imaging modalities suggests it is a property of the field rather than of any single method.(38)

The inflection point of the present period is therefore not a new architecture but a new question, and the remainder of the review addresses it directly. The three periods and their inflection points are summarized in (Table 1).

Table 1. Chronology of generative approaches for dental synthetic data.

Period	Dominant method	Dental application	Inflection point
2014-2018	GAN	CT artifact correction and denoising	Generative model as filter
2019-2022	Conditional and progressive GAN	Radiograph synthesis, augmentation, restoration design	Generative model as factory
2023-2025	Diffusion models	High-fidelity radiographs, privacy and validation	From generation to governance

Note: GAN = generative adversarial network; the period boundaries mark the shift from correction to synthesis, and from synthesis to validation. Source: Summary of the reviewed literature.

DISCUSSION

The trajectory from cleaning to synthesizing to validating is the clearest finding of the review, and it is best read not as a sequence of technical upgrades but as a sequence of questions the field has learned to ask. In the first period the question was whether a generative model could clean a dental image, and the answer, demonstrated in the computed-tomography work, was yes. In the second period the question was whether it could create a dental image, and the answer, demonstrated across panoramic, intraoral, and cephalometric radiographs, was again yes. In the third period the question is whether the created image is private and valid, and here the answer is still, in the main, not yet, which is what makes this period the decisive one.(1,12)

The pattern is familiar from other fields: the technology arrives first, the applications follow, and the governance, privacy, and validation, arrives last, and the dental literature is not unusual in

this, but it is late, and the cost of the delay is that the early claims, that synthetic data is privacy-preserving and that it improves models, were made before the evidence existed to support them.

That the claims were made early is understandable, because the promise was seductive and the alternative, waiting years for validation, was unattractive, but the record now shows that the field could have reached its current questions sooner if it had asked them sooner, and the lesson is general: a generative method's output is only as trustworthy as the evaluation that follows it.

The three-stage reading also explains why the field's current difficulty is not technical: the models now generate images that are hard to distinguish from real ones, and the difficulty is that this very success makes the privacy and fidelity questions harder, not easier, because a model that reproduces its training distribution faithfully is a model that is, by that very faithfulness, at risk of reproducing its training examples. The best generator is, in other words, the one that most needs to be checked, which is a reversal of the intuition, common in the first period, that a better model would resolve the field's doubts.

What changed across the three periods is less the models than the attitude toward them. The technical difference between an adversarial network and a diffusion model is real, and the diffusion models do produce more realistic radiographs, but the more consequential change is the recognition, now appearing in the dental and the medical literature alike, that realism is not the same as fidelity, and that a realistic synthetic image can be a misleading one.(11,12)

A synthetic radiograph that looks like a patient but contains no real patient is the best case; a synthetic radiograph that looks like a patient because it is, in substance, a memorized patient, is the worst, and distinguishing the two requires a privacy analysis that most dental studies still omit.(34,36)

The same distinction applies to utility: the handful of dental studies that tested whether synthetic images improve a downstream classifier found modest and mixed gains, which is a different and more honest result than the unqualified claim that synthetic data solves the data shortage.(10,29)

What persisted across the periods is the gap between generation and evaluation, and it is this gap, rather than any technical limitation, that now bounds the field's progress. The persistence of that gap despite a decade of technical advance is itself worth reflecting on, because it suggests the gap is not a temporary immaturity but a structural feature of a field whose incentives reward the demonstration of a new capability more than the validation of an existing one, and the commentary that frames generative artificial intelligence as raising questions of human oversight, of keeping the machine in the loop rather than trusting its output, speaks to the same concern from the governance side: the more convincing the generated image, the more necessary the human who checks it.(39)

The opportunities, for their part, are real and worth stating without the caution that the privacy discussion requires. The most immediate is augmentation: a synthetic image that is faithful to the real distribution can enlarge a small training set and, in the dental studies that tested it, can improve a downstream classifier, if modestly.(23,29)

The second is imbalance correction: generative models can oversample the rare class, the atypical lesion, the uncommon anatomy, which is precisely where real datasets are thinnest and where a diagnostic model's errors are most dangerous, and this is the use that the bias-mitigation work in other domains has begun to demonstrate.(40)

The third is privacy-preserving sharing: if synthetic images can be shown not to leak their sources, they can be shared across institutions and into teaching where real images cannot, which would address the collaboration bottleneck that currently fragments dental research.(2,33) A fourth opportunity, less discussed but no less real, is the generation of teaching and simulation material, where the privacy bar is lower and the realism bar is what matters, and where the intraoral-generation work of the second period already showed the value.(8)

Each of these opportunities is conditional on the same thing, that the synthetic data be faithful, useful, and private, and it is the conditionality, rather than the opportunity, that the next two sections take up.

The privacy implications deserve to be stated precisely, because they are the place where the field's early enthusiasm is most likely to mislead. Synthetic dental data is not automatically private, and the property that matters is not whether an image was generated but whether it can be traced to a real patient, a question answered by the memorization and membership-inference literature rather than by the generative model's own claims.(34,35)

The legal and regulatory framing is maturing in parallel, with arguments that synthetic data can substitute for real data in regulatory decisions only if its fidelity and its privacy are both demonstrated, and with evaluation frameworks that make those two properties, and utility, explicit

and measurable.(12,41)

The distinction that the legal literature draws, between synthetic data as a privacy-enhancing technology and synthetic data as a privacy risk, depends on how the data was generated and whether it was audited, and a dental institution that treats all synthetic data as equally safe is making exactly the category error the literature warns against.(34)

The dental field has not yet adopted such a framework, and its synthetic-data papers generally report visual quality and a downstream accuracy, leaving the privacy question unexamined, which is the single most important gap the review identifies. The gap is all the more consequential because the privacy claim is often the stated justification for the work: a study that generates synthetic radiographs to enable sharing has, by that very aim, taken on the burden of showing that the sharing would be safe, and a study that omits that demonstration has left its own motivation unsupported.

The neighboring medical literatures have grappled with the same questions longer, and their answers are instructive: the ophthalmology surveys of generative adversarial networks, the cancer-imaging reviews of data synthesis, and the histopathology assessments of generative models all converge on the same caution, that realism is not fidelity and that generated images can leak their sources, and the more recent synthetic-data reviews across medical modalities make the evaluation framework explicit.(14,37,42,43)

The breast-imaging discussion of synthetic data and the in-silico trial literature add the regulatory dimension, arguing that synthetic data enters decision-making only when its provenance and its privacy are documented, and the bias-mitigation work in prostate grading shows the positive case, that synthetic data can correct a skew in the training distribution if it is validated to do so.(40,44,45)

A practical consequence is that a dental researcher or regulator considering synthetic data should treat “synthetic” as a property to be proven, not a guarantee to be assumed, and should require, at minimum, a statement of how fidelity, utility, and privacy were each measured.

The validation requirements follow from the same reasoning, and they can be stated as a short list that the current literature only occasionally satisfies. Fidelity, that the synthetic images resemble the real distribution, is the property most often and most superficially reported, because visual inspection and a small quality metric are easier than a distributional comparison.

Utility, that training on synthetic images improves a downstream model, is the property that matters clinically and is reported least, because it requires the expensive step of actually training a diagnostic model on the synthetic data, and the dental studies that did take that step found gains that were real but modest, which is a more useful result than an unreported one.(10,29)

Privacy, that the synthetic images do not leak their training examples, is the property that is almost never reported, because it requires a privacy analysis that the imaging literature does not habitually perform.(12,36)

The evaluation framework that the broader synthetic-data literature has converged on, measuring fidelity, utility, and privacy together and reporting all three, is the standard the dental field should adopt, and the longitudinal and tabular synthetic-data reviews show that the framework is applicable well beyond imaging.(12,38)

The review’s position is that all three are necessary and none is sufficient, and that a synthetic dental dataset that fails any one of them should not be described as ready for clinical or regulatory use, however realistic its images appear.

The three properties are also not independent, and a dataset that maximizes one may sacrifice another, because a model trained to reproduce the real distribution faithfully is, by that fidelity, closer to memorizing it, and a model that is aggressively privacy-protected may generate images too degraded to be useful, so the validation problem is not three separate tests but one trade-off that must be navigated and reported. A synthetic-data paper that reports only fidelity is therefore not merely incomplete; it is uninterpretable, because the reader cannot know what was traded away to achieve the reported quality.

The research agenda that follows from these conclusions is short and specific. The first priority is to make the privacy question routine: every dental synthetic-data study should report how it tested for memorization or leakage, and the replica-detection methods already used in other imaging domains should be adopted unchanged.(35)

The second is to make the utility question routine: synthetic images should be evaluated by training a downstream diagnostic model on them and reporting the change, not by a visual-quality metric alone, and the dental studies that have done this are the template.(10,29) The third is to make the fidelity question measurable rather than aesthetic: a distributional comparison, or the three-part fidelity-utility-privacy framework, should replace the reader’s subjective judgment.(12)

None of these requires new technology, only new reporting norms, and a field that adopts them would, within a few years, be able to say with evidence what it currently says by assumption, that its synthetic data is safe to share and worth training on.

This review has limitations that follow from its method. It is narrative rather than systematic, and its selection of sources is guided by relevance and by the chronological structure rather than by an exhaustive protocol, so it may omit relevant work, particularly work published outside the indexed English-language literature and work on the commercial systems whose details are not public, and the true extent of commercial synthetic-data generation in dental software is likely understated here.

It does not appraise the risk of bias of the individual studies, and its synthesis of utility and privacy is qualitative, so its statements should be read as indicative rather than as pooled estimates, and a reader seeking a quantitative answer to the question of how much synthetic data improves a given model will not find it here.

The chronological periods, though delimited by explicit criteria, are a simplification of a continuous development, and the assignment of a study to one period or another is a judgment that a different reviewer might draw slightly differently. The review also foregrounds imaging and says less about synthetic tabular and longitudinal dental data, which is a growing adjacent literature with its own validation challenges, and a future synthesis that treated both would be more complete. These limitations bound the precision of the conclusions without changing their direction, which is that the field has matured technically and has only begun to mature methodologically.

CONCLUSIONS

Generative artificial intelligence for dental data has completed two of its three necessary steps. It has demonstrated, first in adversarial networks and then in diffusion models, that it can clean dental images and create new ones, and it has shown, in a smaller number of studies, that the created images can help downstream models.

What it has not yet demonstrated, at the standard the field itself is beginning to articulate, is that the synthetic images are private and that they are valid, and these two demonstrations are now the bottleneck.

The opportunities, augmentation, imbalance correction, rare pathology, and privacy-preserving sharing, are real and increasingly well-motivated, but they are claims that require proof, and the proof requires measuring fidelity, utility, and privacy together rather than reporting visual quality alone.

The field's next stage, and the one that will decide whether synthetic dental data becomes a clinical instrument or remains a research curiosity, is not a better generator but a better evaluation of the generators it already has.

The trajectory is nonetheless encouraging in its direction: the field has moved, in a decade, from using generative models as filters to using them as factories, and it is now learning to use them as instruments whose outputs must be audited, and that progression, from generation to governance, is the mark of a technology that is maturing rather than merely proliferating.

The next chapter of that progression is already visible in outline, and it will be written not by the generation of yet more convincing images but by the discipline of demonstrating, for each generated image, that it is faithful, that it is useful, and that it is safe.

REFERENCES

1. Umer F, Adnan N. Generative artificial intelligence: synthetic datasets in dentistry. *BDJ Open*. 2024;10(1):13. doi: 10.1038/s41405-024-00198-4
2. Nakayama LF, de Matos JCRG, Stewart IU, Mitchell WG, Martinez-Martin N, Regatieri CVS, et al. Retinal Scans and Data Sharing: The Privacy and Scientific Development Equilibrium. *Mayo Clinic Proceedings: Digital Health*. 2023;1(2):67-74. doi: 10.1016/j.mcpdig.2023.02.003
3. Khosravi B, Purkayastha S, Erickson BJ, Trivedi HM, Gichoya JW. Exploring the potential of generative artificial intelligence in medical image synthesis: opportunities, challenges, and future directions. *The Lancet Digital Health*. 2025;7(9):100890. doi: 10.1016/j.landig.2025.100890
4. Pan Z, Yu W, Yi X, Khan A, Yuan F, Zheng Y. Recent Progress on Generative Adversarial Networks (GANs): A Survey. *IEEE Access*. 2019;7:36322-36333. doi: 10.1109/access.2019.2905015
5. Jabbar A, Li X, Omar B. A Survey on Generative Adversarial Networks: Variants, Applications, and Training. *ACM Comput Surv*. 2022;54(8):1-49. doi: 10.1145/3463475
6. Hu Z, Jiang C, Sun F, Zhang Q, Ge Y, Yang Y, et al. Artifact correction in low-dose dental CT imaging using Wasserstein generative adversarial networks. *Medical Physics*. 2019;46(4):1686-1696. doi: 10.1002/mp.13415
7. Hegazy MAA, Cho MH, Lee SY. Image denoising by transfer learning of generative adversarial network for dental CT. *Biomed Phys Eng Express*. 2020;6(5):055024. doi: 10.1088/2057-1976/abb068
8. Kokomoto K, Okawa R, Nakano K, Nozaki K. Intraoral image generation by progressive growing of generative adversarial network and evaluation of generated image quality by dentists. *Sci Rep*. 2021;11(1):18517. doi: 10.1038/s41598-021-98043-3
9. Kim M, Kim S, Kim M, Bae HJ, Park JW, Kim N. Realistic high-resolution lateral cephalometric radiography generated by progressive growing generative adversarial network and quality evaluations. *Sci Rep*. 2021;11(1):12563. doi: 10.1038/s41598-021-91965-y
10. Yang S, Kim KD, Arijj E, Takata N, Kise Y. Evaluating the performance of generative adversarial network-synthesized periapical images in classifying C-shaped root canals. *Sci Rep*. 2023;13(1):18038. doi: 10.1038/s41598-023-45290-1
11. Kirkwood B, Choi BY, Bynum J, Salinas J. AI Diffusion Models Generate Realistic Synthetic Dental Radiographs Using a Limited Dataset. *J Imaging*. 2025;11(10):356. doi: 10.3390/jimaging11100356
12. Hernandez M, Osorio-Marulanda PA, Catalina M, Loinaz L, Epelde G, Aginako N. Comprehensive evaluation framework for synthetic tabular data in health: fidelity, utility and privacy analysis of generative models with and without privacy guarantees. *Front Digit Health*. 2025;7:1576290. doi: 10.3389/fdgh.2025.1576290
13. Sharma P, Kumar M, Sharma HK, Biju SM. Generative adversarial networks (GANs): Introduction, Taxonomy, Variants, Limitations, and Applications. *Multimed Tools Appl*. 2024;83(41):88811-88858. doi: 10.1007/s11042-024-18767-y
14. You A, Kim JK, Ryu IH, Yoo TK. Application of generative adversarial networks (GAN) for ophthalmology image domains: a survey. *Eye and Vis*. 2022;9(1):6. doi: 10.1186/s40662-022-00277-3
15. Kim HS, Ha EG, Lee A, Choi YJ, Jeon KJ, Han SS, et al. Refinement of image quality in panoramic radiography using a generative adversarial network. *Dentomaxillofacial Radiology*. 2023;52(5):20230007. doi: 10.1259/dmfr.20230007
16. Tian S, Wang M, Dai N, Ma H, Li L, Fiorenza L, et al. DCPR-GAN: Dental Crown Prosthesis Restoration Using Two-Stage Generative Adversarial Networks. *IEEE J Biomed Health Inform*. 2022;26(1):151-160. doi: 10.1109/jbhi.2021.3119394
17. Tian S, Wang M, Yuan F, Dai N, Sun Y, Xie W, et al. Efficient Computer-Aided Design of Dental Inlay Restoration: A Deep Adversarial Framework. *IEEE Trans Med Imaging*. 2021;40(9):2415-2427. doi: 10.1109/tmi.2021.3077334
18. Yuan F, Dai N, Tian S, Zhang B, Sun Y, Yu Q, et al. Personalized design technique for the dental occlusal surface based on conditional generative adversarial networks. *Numer Methods Biomed Eng*. 2020;36(5):e3321. doi: 10.1002/cnm.3321
19. Cho JH, Yi Y, Choi J, Ahn J, Yoon HI, Yilmaz B. Time efficiency, occlusal morphology, and internal fit of anatomic contour crowns designed by dental software powered by generative adversarial network: A comparative study. *Journal of Dentistry*. 2023;138:104739. doi: 10.1016/j.jdent.2023.104739
20. Huang Y, Fan F, Syben C, Roser P, Mill L, Maier A. Cephalogram synthesis and landmark detection in dental cone-beam CT systems. *Medical Image Analysis*. 2021;70:102028. doi: 10.1016/j.media.2021.102028
21. Kim T, Cho Y, Kim D, Chang M, Kim YJ. Tooth Segmentation of 3D Scan Data Using Generative Adversarial Networks. *Applied Sciences*. 2020;10(2):490. doi: 10.3390/app10020490
22. Lin Y, Fan F, Zhang J, Zhou J, Liao P, Chen H, et al. DHI-GAN: Improving Dental-Based Human Identification Using Generative Adversarial Networks. *IEEE Trans Neural Netw Learning Syst*. 2023;34(12):9700-9712. doi: 10.1109/tnnls.2022.3159781
23. Waqas M, Hasan S, Chori AF, Alfaraj A, Faheemuddin M, Khurshid Z. Synthetic Orthopantomography Image Generation Using Generative Adversarial Networks for Data Augmentation. *International Dental Journal*. 2025;75(6):103878. doi: 10.1016/j.identj.2025.103878
24. Broll A, Goldhacker M, Hahnel S, Rosentritt M. Generative deep learning approaches for the design of dental restorations: A narrative review. *Journal of Dentistry*. 2024;145:104988. doi: 10.1016/j.jdent.2024.104988
25. Arian MSH, Sifat FA, Ahmed S, Mohammed N, Farook TH. Unsupervised tooth segmentation from three dimensional scans of the dental arch using domain adaptation of synthetic data. *International Journal of Medical Informatics*. 2025;195:105769. doi: 10.1016/j.ijmedinf.2024.105769

26. Kearney VP, Yansane AIM, Brandon RG, Vaderhobli R, Lin GH, Hekmatian H, et al. A generative adversarial inpainting network to enhance prediction of periodontal clinical attachment level. *Journal of Dentistry*. 2022;123:104211. doi: 10.1016/j.jdent.2022.104211
27. Yang S, Kim KD, Arij E, Kise Y. Generative adversarial networks in dental imaging: a systematic review. *Oral Radiol*. 2024;40(2):93-108. doi: 10.1007/s11282-023-00719-1
28. Islam S, Aziz MT, Nabil HR, Jim JR, Mridha MF, Kabir MM, et al. Generative Adversarial Networks (GANs) in Medical Imaging: Advancements, Applications, and Challenges. *IEEE Access*. 2024;12:35728-35753. doi: 10.1109/access.2024.3370848
29. Tolstaya E, Tichy A, Paris S, Schwendicke F. Improving machine learning-based bitewing segmentation with synthetic data. *Journal of Dentistry*. 2025;156:105679. doi: 10.1016/j.jdent.2025.105679
30. Zhang Z, Yan J, Shi Y, Cui Z, Xu J, Shen D. Coupled Diffusion Models for Metal Artifact Reduction of Clinical Dental CBCT Images. *IEEE Trans Med Imaging*. 2025;44(12):5103-5116. doi: 10.1109/tmi.2025.3587131
31. Yu W, Guo X, Li W, Liu X, Chen H, Yuan Y. ToothMaker: Realistic Panoramic Dental Radiograph Generation via Disentangled Control. *IEEE Trans Med Imaging*. 2025;44(12):5233-5244. doi: 10.1109/tmi.2025.3588466
32. Pedersen S, Jain S, Chavez M, Ladehoff V, de Freitas BN, Pauwels R. Pano-GAN: A Deep Generative Model for Panoramic Dental Radiographs. *J Imaging*. 2025;11(2):41. doi: 10.3390/jimaging11020041
33. Schoenhof R, Schoenhof R, Blumenstock G, Lethaus B, Hoefert S. Synthetic, non-person related panoramic radiographs created by generative adversarial networks in research, clinical, and teaching applications. *Journal of Dentistry*. 2024;146:105042. doi: 10.1016/j.jdent.2024.105042
34. Nisevic M, Milojevic D, Spajic D. Synthetic data in medicine: Legal and ethical considerations for patient profiling. *Comput Struct Biotechnol J*. 2025;28:190-198. doi: 10.1016/j.csbj.2025.05.026
35. Aydin OU, Koch A, Hilbert A, Rieger J, Lohrke F, Ishida F, et al. RELICT-NI: Replica Detection in Synthetic Neuroimaging—A Study on Noncontrast CT and Time-of-Flight MRA. *Neuroinform*. 2025;23(4):54. doi: 10.1007/s12021-025-09745-2
36. Reiner Benaim A, Almog R, Gorelik Y, Hochberg I, Nassar L, Mashiach T, et al. Analyzing Medical Research Results Based on Synthetic Data and Their Relation to Real Data Results: Systematic Comparison From Five Observational Studies. *JMIR Med Inform*. 2020;8(2):e16492. doi: 10.2196/16492
37. Ibrahim M, Khalil YA, Amirrajab S, Sun C, Breeuwer M, Pluim J, et al. Generative AI for synthetic data across multiple medical modalities: A systematic review of recent developments and challenges. *Computers in Biology and Medicine*. 2025;189:109834. doi: 10.1016/j.compbimed.2025.109834
38. Miletic M, Sariyar M. Synthetic data generation methods for longitudinal and time series health data: a systematic review. *BMC Med Inform Decis Mak*. 2025;26(1):30. doi: 10.1186/s12911-025-03326-8
39. Pesapane F, Cuocolo R, Sardanelli F. The Picasso's skepticism on computer science and the dawn of generative AI: questions after the answers to keep "machines-in-the-loop". *Eur Radiol Exp*. 2024;8(1):81. doi: 10.1186/s41747-024-00485-7
40. Van Booven DJ, Chen CB, Kryvenko ON, Punnen S, Sandoval V, Malpani S, et al. Mitigating bias in prostate cancer diagnosis using synthetic data for improved AI driven Gleason grading. *npj Precis Onc*. 2025;9(1):151. doi: 10.1038/s41698-025-00934-5
41. Alloza C, Knox B, Raad H, Aguilà M, Coakley C, Mohrova Z, et al. A Case for Synthetic Data in Regulatory Decision-Making in Europe. *Clin Pharma and Therapeutics*. 2023;114(4):795-801. doi: 10.1002/cpt.3001
42. Osuala R, Kushibar K, Garrucho L, Linardos A, Szafranowska Z, Klein S, et al. Data synthesis and adversarial networks: A review and meta-analysis in cancer imaging. *Medical Image Analysis*. 2023;84:102704. doi: 10.1016/j.media.2022.102704
43. Alajaji SA, Khoury ZH, Elgharib M, Saeed M, Ahmed ARH, Khan MB, et al. Generative Adversarial Networks in Digital Histopathology: Current Applications, Limitations, Ethical Considerations, and Future Directions. *Modern Pathology*. 2024;37(1):100369. doi: 10.1016/j.modpat.2023.100369
44. Pesapane F, D'Amelio L, Nicosia L, Mallardi C, Bozzini A, Meneghetti L, et al. The Role of Synthetic Data and Generative AI in Breast Imaging: Promise, Pitfalls, and Pathways Forward. *Diagnostics*. 2025;15(23):2996. doi: 10.3390/diagnostics15232996
45. Simalatsar A. Synthetic biomedical data generation in support of In Silico Clinical Trials. *Front Big Data*. 2023;6:1085571. doi: 10.3389/fdata.2023.1085571

FUNDING

No financing

CONFLICT OF INTEREST

None

AUTHORSHIP CONTRIBUTIONS

Drafting – original draft: Nairobi Hernández Bridon.
Writing–review and editing: Nairobi Hernández Bridon.