ORIGINAL RESEARCH **Open access**

Generative Adversarial Network with Privacy Guarantees for Creating Synthetic Histopathology Images of Rare Pediatric Tumors for Training Deep Learning Models

Rania Hassan[1*], Dina Fathy[1]

Abstract

Rare pediatric tumors like sarcomas, neuroblastoma, medulloblastoma, and retinoblastoma pose a challenge for developing deep learning models due to the limited availability of histopathology images, which are distributed across multiple institutions. This scarcity is compounded by privacy concerns, as whole-slide images often contain sensitive clinical and genomic data, and generative adversarial networks (GANs) risk memorizing and leaking training samples. To address this, a differentially private GAN framework is proposed for synthesizing high-resolution histopathology patches of rare pediatric cancers. The framework incorporates a generator for image synthesis, a discriminator for realism assessment, per-sample gradient clipping, Gaussian noise injection, and a privacy accountant, ensuring provable privacy guarantees during the training process. The synthetic images generated can aid in data augmentation, model pre-training, and benchmarking without exposing identifiable pathology data, offering a privacy-preserving solution for dataset augmentation while emphasizing the importance of clinical validation.

Keywords Differential privacy, Membership inference, Generative adversarial network, Synthetic histopathology, Pediatric tumors, Whole-slide imaging

*Correspondence:
Rania Hassan
rania.hassan@gmail.com
[1] Department of Healthcare Analytics and AI Systems, Alexandria University, Alexandria, Egypt

Introduction

Rare pediatric cancers impose a disproportionate diagnostic and computational burden because each disease subtype may appear infrequently within a single hospital, yet accurate histopathological classification requires exposure to diverse morphologic patterns. Deep learning has shown promise for extracting diagnostic and molecular information from histopathology images, including tumor classification and mutation-associated morphology in adult cancers [1, 2]. Pediatric tumor applications are beginning to demonstrate similar potential, particularly in neuroblastoma and sarcoma histopathology, but the available datasets remain much smaller and less diverse than those used in common adult malignancies [3-6]. This creates a structural mismatch between the data requirements of high-capacity models and the limited availability of rare pediatric tumor whole-slide images.

Direct data sharing is not a sufficient remedy because pediatric pathology data are governed by heightened privacy expectations, institutional review constraints, and cross-border regulatory limits. Whole-slide images may contain tissue patterns associated with diagnosis, treatment history, or rare disease identity, making them difficult to treat as low-risk research images. Reviews of medical synthetic data emphasize that utility must be balanced against privacy, especially when released data could be

reused outside the original governance environment [7-9]. In this context, synthetic histopathology is attractive only if its release mechanism is explicitly privacy-preserving rather than merely visually realistic.

Conventional GANs can synthesize plausible histopathology images, but realism alone does not guarantee safety. Generative models may memorize rare examples, especially when training data are small, high-dimensional, and morphologically distinctive, making rare pediatric tumor images particularly vulnerable to unintended disclosure. Membership inference attacks have shown that adversaries may determine whether a particular record contributed to model training, and later work extended this concern to generative models and synthetic health data [10-13]. Therefore, a synthetic pathology framework for rare pediatric cancer must address privacy leakage as a first-order design constraint, not as a post hoc audit.

This article develops a conceptual framework for a differentially private GAN that generates synthetic histopathology patches of rare pediatric tumors while providing a formal privacy guarantee. The framework draws on digital pathology GAN research, high-resolution histopathology synthesis, and medical data privacy methods to define an architecture suitable for scarce pediatric cancer slides [14-21]. It does not report experiments, fabricated performance gains, or simulated results; instead, it specifies the design logic, privacy mechanism, evaluation strategy, and clinical boundaries of a privacy-preserving generative system. The roadmap proceeds from histopathology and generative modeling background to architecture, differential privacy integration, privacy auditing, downstream utility, and rare pediatric tumor adaptation.

Background

Rare pediatric cancers from a histopathology perspective

Rare pediatric cancers are histologically heterogeneous, with sarcomas, neuroblastoma, medulloblastoma, retinoblastoma, Wilms tumor, and embryonal tumors each presenting distinct nuclear morphology, mitotic activity, stromal organization, and tissue architecture. Neuroblastoma histopathology, for example, has been approached using deep learning methods that encode morphologic features for classification, while sarcoma-focused studies highlight the need for multicenter integration because no single site captures sufficient subtype diversity [3-6]. These tumors are usually represented as whole-slide images after hematoxylin and eosin staining, from which diagnostically meaningful regions must be sampled at magnifications that preserve cellular and stromal detail. A privacy-preserving synthetic framework must therefore respect both subtype-specific morphology and the patch-level sampling structure through which whole-slide images are made computationally tractable.

Existing generative models for histopathology

Generative models in digital pathology have been used for image augmentation, stain transformation, high-resolution synthesis, segmentation support, and layout-conditioned tissue generation. GAN-based approaches have produced diagnostic-quality pathology images, selective synthetic augmentation, stitched histology regions, and high-resolution histopathology outputs, but reviews note recurring risks such as mode collapse, texture imitation without diagnostic preservation, and artifacts introduced by patch tiling [15, 16, 18-21]. More recent synthesis approaches explore bespoke cellular layouts, frequency-spatial hybrid generation, and ultra-resolution histopathology synthesis, indicating that visual fidelity is improving but remains technically constrained by tissue scale and morphology [22-24]. For rare pediatric tumors, these methods require privacy-aware redesign because small datasets increase the probability that a visually convincing output may be too close to an identifiable training image.

Privacy risks of generative models

Privacy risks in generative pathology arise because the model learns from real patient images while producing outputs that may be released beyond the original clinical institution. Membership inference attacks assess whether a target image was included in training, while model inversion, attribute inference, and leakage from gradients or generated samples may reveal sensitive information about rare cases [10-13]. These risks are amplified in pediatric oncology because rare tumor subtypes may be represented by very few slides, making memorization more likely and harder to detect through visual inspection alone. Consequently, a safe generative framework must combine

architectural choices with formal privacy mechanisms and adversarial auditing.

Differential privacy fundamentals

Differential privacy provides a mathematical definition of limited disclosure by bounding how much the output distribution of a mechanism can change when one training record is added or removed. In deep learning, the standard mechanism clips per-sample gradients to bound sensitivity and then adds calibrated Gaussian noise before parameter updates, with a privacy accountant tracking the accumulated privacy loss across training iterations [25-28]. The parameters ε and δ define the strength of the privacy guarantee, while Rényi or moments-based accounting enables more precise tracking over many optimization steps. The central trade-off is that stronger privacy generally increases noise, which may degrade convergence, image fidelity, and downstream diagnostic utility.

Framework Overview

High-level architecture

The proposed framework begins with rare pediatric tumor whole-slide images that are curated into a small real patch dataset, then trains a GAN under differential privacy using gradient clipping and Gaussian noise injection. The generator produces synthetic histopathology patches, while the discriminator or critic provides realism feedback during training but is never treated as a clinical validator. After training, generated patches pass through a privacy audit and a pathology-oriented fidelity review before being used for downstream classifier development or benchmarking. This architecture integrates the synthetic augmentation logic of histopathology GANs with the privacy-accounted training logic used in differentially private medical image learning [14, 18, 25-27].

Figure 1 presents the proposed differentially private GAN architecture for transforming scarce rare pediatric tumor WSI patches into privacy-governed synthetic histopathology images for downstream deep learning development.

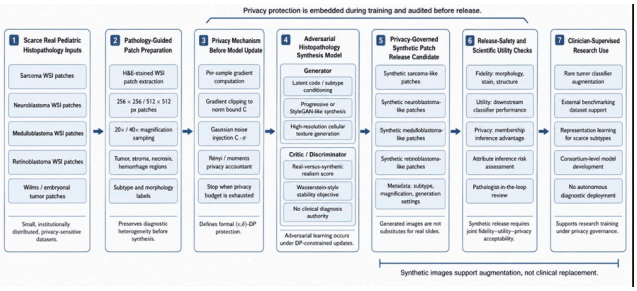

Figure 1. Differentially Private GAN Framework for Synthetic Rare Pediatric Tumor Histopathology Generation

Core assumptions

The framework assumes that a small but curated rare pediatric tumor dataset is available, such as tens to hundreds of diagnostically reviewed WSI-derived patches per tumor subtype rather than thousands of independent whole-slide images. It also assumes that the privacy mechanism is implemented at the training level rather than approximated through visual de-identification, because anonymization alone cannot remove all inference risk from high-dimensional medical images [7-9, 29]. Computational cost is expected to be higher than conventional GAN training because per-sample gradients, noise calibration, and privacy accounting increase optimization complexity. These assumptions position the framework as a governed research infrastructure for consortial pediatric pathology rather than a plug-and-play image augmentation method.

Design principles

The first design principle is provable privacy: synthetic outputs should inherit protection from a differentially private training mechanism rather than relying only on the absence of visible identifiers. The second principle is utility preservation, meaning that synthetic images should retain diagnostically relevant morphology such as nuclear density, mitotic pattern, stromal organization, and subtype-associated architecture rather than merely reproducing plausible texture [19, 21-23]. The third principle is adversarial conservatism, because privacy must be evaluated against strong membership inference and synthetic health data attack scenarios rather than benign visual review alone [10-13]. The final principle is clinical non-substitution: synthetic patches can support model training and benchmarking, but they cannot replace real pediatric pathology review or prospective validation.

Table 1 consolidates the privacy, fidelity, utility, and governance design choices that determine whether DP-

GAN-generated pediatric tumor histopathology patches can be responsibly used for downstream model development.

Table 1. Privacy–Utility–Fidelity Design Matrix for DP-GAN-Based Synthetic Rare Pediatric Tumor Histopathology

Design Dimension	Framework Role	Core Mechanism	Risk
Scarce pediatric tumor input curation	Defines the real data boundary from which the model learns	Subtype-stratified WSI patch extraction from sarcoma, neuroblastoma, medulloblastoma, retinoblastoma, Wilms tumor, and embryonal tumor slides	Overre of morph underre of ra
Patch-level synthesis	Makes WSI-scale learning computationally feasible	256 × 256 or 512 × 512 high-magnification H&E patch generation	Loss contex co
Generator design	Produces synthetic histopathology images	Progressive or StyleGAN-like architecture with subtype-aware conditioning	Tex replica diagn pre
Critic or discriminator	Provides adversarial realism feedback	Wasserstein-style critic or stabilized discriminator	Unsta mod unrea p
Per-sample gradient clipping	Bounds each image's influence on training	Clip each real-image gradient to norm C before aggregation	Excess of identifia
Gaussian noise injection	Obscures individual data contributions	Add calibrated Gaussian noise proportional to C · σ	Memo tra me le
Privacy accounting	Tracks cumulative privacy loss	Rényi or moments accountant over training steps	Unc priva ex
Membership inference audit	Tests practical leakage resistance	Compare attack advantage for DP-GAN versus non-private GAN	Ad dis wheth image t
Downstream classifier evaluation	Tests whether synthetic images add scientific value	Real-only versus real plus DP-synthetic versus real plus non-DP synthetic comparison	Synth falsel a per
Clinical boundary setting	Prevents misuse of synthetic pathology images	Explicit restriction to research augmentation and benchmarking	Misinte diagn clinic

Gan Architecture for WSI Synthesis

Patch-based synthesis strategy

Patch-based synthesis is necessary because pediatric tumor whole-slide images are extremely large, heterogeneous, and computationally expensive to synthesize at full resolution. The framework extracts diagnostically annotated patches, such as 256 × 256 or 512 × 512 pixels, at 20× or 40× magnification while preserving representative regions of tumor, necrosis, stroma, hemorrhage, and normal-adjacent tissue. Prior histopathology GAN studies show that patch-level synthesis can support classification, augmentation, and high-resolution image generation, but they also demonstrate that tissue continuity and sampling bias remain important concerns [18-21, 30]. For rare pediatric tumors, patch selection should therefore be stratified by subtype and morphology so that the generator does not overrepresent common visual patterns while ignoring diagnostically rare features.

Generator architecture

The generator should use a progressive or StyleGAN-like design capable of producing high-resolution cellular texture while controlling artifacts introduced by upsampling. Histopathology synthesis research has shown that diagnostic-quality outputs require more than global realism; generated patches must preserve meaningful nuclear contours, chromatin distribution, stromal relationships, and tumor-specific morphology [19, 22-24]. Adaptive or instance normalization may help stabilize stain and texture variation, but excessive normalization could suppress subtype-specific histological cues. Because pediatric tumor datasets are small, the generator must be explicitly constrained by differential privacy and evaluated for memorization rather than optimized only for visual fidelity.

Discriminator or critic

The discriminator or critic estimates whether a patch resembles the real pediatric tumor distribution, but its role in this framework is limited to training feedback rather than diagnostic judgment. Wasserstein-style critics are conceptually attractive because they can provide smoother optimization signals and may be more stable under noisy updates than standard adversarial losses, particularly when training data are scarce and privacy noise is present. Existing histopathology GAN studies demonstrate that discriminator design influences synthetic image realism, segmentation compatibility, and downstream augmentation utility [18, 20, 21, 30]. In the proposed DP-GAN, the critic must be trained through the same privacy-aware update mechanism when it directly accesses real images, ensuring that privacy protection is embedded in adversarial learning rather than added only to the released generator.

Differential Privacy Integration

Per-sample gradient clipping

Per-sample gradient clipping is the first privacy mechanism because it limits how much any single real histopathology patch can influence a model update. For each minibatch, the gradient contribution from every real image is computed separately and clipped to a fixed norm bound C before aggregation, thereby controlling sensitivity even when a rare pediatric tumor patch has unusually distinctive morphology. Differentially private medical image learning studies emphasize that this step is essential for making subsequent noise calibration meaningful, because Gaussian noise cannot provide a valid privacy guarantee when individual gradients are unbounded [25-28]. In the proposed framework, clipping should be applied to all adversarial updates that depend on real pediatric pathology patches, including critic or discriminator updates and any generator updates whose gradients are mediated through real-image comparisons.

Gaussian noise injection

After clipping, Gaussian noise scaled to the clipping bound and a noise multiplier σ is added to the aggregated gradient before the optimizer updates model parameters. This noise weakens the influence of any individual pediatric tumor patch on the learned generator distribution, reducing the likelihood that a released synthetic image reveals whether a specific patient slide was included in training. The privacy-utility literature on synthetic medical data shows that stronger noise can improve protection but may reduce fidelity and downstream usefulness, making calibration a core design decision rather than a technical afterthought [7, 9, 29, 31]. For rare pediatric histopathology, the selected σ should be justified by the intended release setting, the sensitivity of the disease subtype, and the expected adversary's ability to compare generated patches with candidate patient images.

Privacy budget accounting

A privacy accountant tracks the accumulated ε for a chosen δ across all training steps, batch sampling events, and optimization epochs. Rényi differential privacy or moments-based accounting is appropriate because adversarial training requires repeated access to small real datasets, and privacy loss grows with continued optimization [25-28]. Training should stop when the predefined privacy budget is exhausted, even if visual fidelity could improve with additional epochs, because the framework prioritizes formal release safety over unconstrained image realism. Once training is complete, the released generator and its synthetic outputs are treated as products of the differentially private mechanism, subject to subsequent privacy audit through membership inference and synthetic health data risk assessment [10-13].

Privacy Guarantees and Attack Resilience

Formal (ε,δ)-DP guarantee for synthetic outputs

The formal privacy claim of the framework is that the trained generator is the output of a differentially private learning mechanism, so any synthetic image sampled from it is protected by the same bounded disclosure guarantee. Under (ε,δ)-differential privacy, the presence or absence of a single pediatric tumor patch in the training set should not substantially change the probability distribution of released generator outputs, which is why the guarantee is stronger than conventional anonymization or visual de-identification [25-28]. This is especially important for rare pediatric cancers because individual slides may contain unusual morphologic signatures that could otherwise be memorized by a high-capacity generator. The guarantee does not mean that synthetic images are clinically perfect or risk-free; rather, it defines a mathematically bounded relationship between the training dataset and the released synthetic distribution.

Empirical membership inference resistance

Although the framework is conceptual, it requires any future implementation to evaluate membership inference resistance as a privacy audit rather than relying only on the formal accountant. Membership inference attacks against machine learning and generative models demonstrate that adversaries may distinguish training samples from non-training samples when models overfit or memorize rare examples [10-12]. In synthetic health data, this risk is clinically meaningful because a successful attacker could infer that a child's rare tumor slide contributed to a dataset, even if the released image is not an exact copy [13]. Therefore, the framework defines attack resilience as a combined requirement: the DP accountant must report the final privacy budget, and empirical attacks should show reduced membership advantage relative to a non-private GAN trained on the same source distribution.

Downstream Utility for Classifier Training

Augmentation strategy

The intended utility pathway is to augment a small real pediatric histopathology dataset with DP-generated synthetic patches, then train a downstream classifier for subtype recognition, triage support, or feature representation learning. Prior work on selective synthetic augmentation and histopathology GAN synthesis suggests that generated images can improve training diversity when synthetic samples are morphologically plausible and diagnostically aligned with the target task [18-21]. In this framework, synthetic patches should not be mixed indiscriminately with real images; instead, they should be labeled by tumor subtype, magnification, tissue compartment, and generation settings so that downstream models do not learn artifacts as disease signals. The downstream classifier may use convolutional or transformer-based architectures, but its final evaluation must occur on held-out real pediatric pathology slides rather than on synthetic validation data.

Utility under differential privacy

Utility under differential privacy is governed by a trade-off between privacy strength and synthetic image fidelity. Smaller ε values generally imply stronger privacy but require more noise during training, which may reduce nuclear sharpness, mitotic visibility, stromal continuity, and the subtype-specific texture needed for classifier learning [25-28]. Synthetic data utility studies show that high apparent fidelity does not automatically translate into useful downstream performance, and that privacy-preserving generation must be evaluated along privacy, fidelity, and task utility dimensions simultaneously [7-9, 29, 31]. For rare pediatric tumor augmentation, the acceptable privacy budget should be selected according to clinical sensitivity and release scope, with the goal of improving real-data-only model training without permitting unsafe memorization of scarce cases.

Evaluation Strategy

Table 2 specifies an evaluation logic that separates visual plausibility, diagnostic feature preservation, privacy protection, membership inference resistance, and downstream classifier utility.

Table 2. Evaluation Logic for Synthetic Rare Pediatric Tumor Histopathology Generated under Differential Privacy

Evaluation Domain	Primary Question	Recommended Comparison	Main M

Visual fidelity	Do synthetic patches resemble plausible pediatric tumor histopathology?	Real patches versus DP-synthetic patches	Patho review consis nucl morph stro organi
Distributional fidelity	Does the synthetic distribution cover the real patch distribution?	Real patch embeddings versus synthetic patch embeddings	FID ada patho embed preci recall, s cove
Diagnostic feature preservation	Are tumor-relevant structures retained?	Tumor-region synthetic patches versus subtype-matched real patches	Nuc density, count necr represe stro archite
Privacy guarantee	Does the training mechanism provide bounded disclosure?	DP-GAN training configuration versus predefined privacy budget	Final clipping C, n multip numb ste
Membership inference resistance	Can an adversary infer training-set participation?	DP-GAN versus non-private GAN under matched attack setting	Atta accura positiv false p rate, a advar
Attribute inference risk	Could generated images reveal sensitive tumor-associated features?	Generated patches with known subtype or molecular labels versus adversarial attribute prediction	Attril predi accu confic thresh rare-fe leak
Downstream augmentation value	Do synthetic images improve classifier training?	Real-only versus real plus DP-synthetic versus real plus non-DP synthetic	AURC sco balar accu calibr subt sensi
Bias and subtype equity	Are rare subtypes adequately represented?	Performance across tumor categories and morphology strata	Per-su sensi confu mat minority rec
Pathologist validation	Are synthetic images acceptable for research training use?	Blinded review of real and synthetic patches	Diagn plausi artif detec confic rati rejectio
Release governance	Is the synthetic dataset safe and appropriately bounded?	Proposed release package versus privacy and use restrictions	Docume comple intende stater audit r acc cont

Fidelity metrics

Fidelity evaluation should combine generic image realism metrics with pathology-specific morphology assessment. Measures such as Fréchet Inception Distance, precision, and recall can estimate whether synthetic patches approximate the real image distribution, but digital

pathology requires additional checks for nuclear morphology, mitotic figure preservation, stromal structure, necrosis patterns, and stain variability [14-17]. High-resolution histopathology synthesis studies also indicate that tiling artifacts, unrealistic cellular boundaries, and texture-only replication can remain hidden when only global image metrics are used [21-24]. Therefore, the framework treats fidelity as a layered construct: visual similarity is necessary, but diagnostic feature preservation and pathologist review are required before synthetic patches are considered useful for model development.

Downstream utility metrics

Downstream utility should be measured by training classifiers under controlled data regimes and testing them only on held-out real pediatric histopathology images. The core comparison is between a real-only baseline, a real plus DP-synthetic training set, and a real plus non-private synthetic training set, with metrics such as AUROC, F1 score, balanced accuracy, calibration, and subtype-level sensitivity [3, 5, 6, 18, 30]. This comparison separates the value of synthetic augmentation from the cost of privacy noise, while also identifying whether non-private gains depend on unsafe memorization. Because adult cancer histopathology models have shown strong performance in mutation prediction and classification tasks, evaluation in pediatric tumors must be stricter rather than assumed transferable from common adult datasets [1, 2].

Privacy metrics

Privacy evaluation should include both formal and empirical metrics, because a privacy accountant alone may not capture practical attack behavior under realistic adversarial assumptions. Membership inference attack accuracy, true positive rate, false positive rate, and attack advantage should be reported for DP-GAN and non-private GAN variants, with special attention to rare morphologic subtypes that may be easiest to memorize [10-13]. Attribute inference should also be considered if generated images could reveal sensitive tumor features associated with molecular subtype, treatment response, or rare diagnostic category. A synthetic release should be considered acceptable only when fidelity and utility are adequate and privacy attacks remain constrained within the intended risk tolerance.

Rare Pediatric Tumor Focus

Handling extreme scarcity

Extreme scarcity can be addressed by pre-training a generative model on broader histopathology data and then differentially private fine-tuning on scarce pediatric tumor patches. This strategy may improve convergence and image quality because the model learns general tissue texture, staining variation, and cellular organization before accessing the sensitive rare pediatric dataset [15, 19, 21, 24]. However, pre-training must be carefully governed: adult or generic pathology images may help the generator learn histological priors, but they cannot substitute for pediatric tumor morphology. The DP fine-tuning phase is therefore the privacy-critical stage, because it is the point at which the model learns from rare pediatric sarcoma, neuroblastoma, medulloblastoma, retinoblastoma, or other scarce tumor patches.

Domain shift from adult to pediatric histology

Domain shift is a central limitation because pediatric tumors often show embryonal, small round blue cell, sarcomatous, or developmental morphologies that differ from adult epithelial cancers commonly used in digital pathology benchmarks. Studies of neuroblastoma and pediatric sarcoma show that pediatric histology requires attention to tumor-specific morphology and multicenter variation rather than simple transfer from adult cancer image models [3-6, 32]. Attention-based conditioning, subtype-aware latent codes, and style-transfer mechanisms may help adapt adult-pretrained generative priors to pediatric morphology, but these mechanisms must remain within the DP training boundary when they learn from sensitive pediatric images. The framework therefore treats pediatric domain adaptation as a constrained transfer problem: useful prior knowledge may be imported, but rare child-specific tissue information must be learned under formal privacy protection.

Limitations

Technical limitations

Differential privacy introduces noise that may degrade fine-grained histological details, including mitotic figures, apoptotic bodies, nuclear atypia, chromatin texture, and subtle stromal invasion patterns. High-resolution WSI synthesis also remains technically difficult because full slides contain multiscale structure, tissue folds, scanner variation, and spatial dependencies that patch-based GANs

cannot fully reproduce [20, 21, 23, 24, 29]. DP training may further increase instability, computational cost, and non-convergence risk, particularly when pediatric tumor datasets contain very few examples per subtype [25, 26, 28]. As a result, the proposed framework should be interpreted as a privacy-preserving augmentation architecture, not as a claim that synthetic slides can fully replicate rare pediatric tumor pathology.

Clinical limitations

Synthetic histopathology images may fail to represent the complete biological and diagnostic variation of real pediatric tumors, especially when rare subtypes, treatment effects, or unusual morphologic presentations are absent from the source dataset. Pathologist-in-the-loop validation is essential because computational fidelity metrics may miss clinically meaningful artifacts or misleading patterns that could bias downstream classifiers [15-17, 31]. Clinical deployment remains distant, since synthetic images should support research augmentation and benchmarking rather than autonomous diagnosis or replacement of expert review. Regulatory frameworks for synthetic pathology data are also immature, so governance must define release scope, audit requirements, documentation standards, and restrictions on clinical use.

Conclusion

The proposed framework defines a privacy-preserving generative pathway for rare pediatric tumor histopathology. It combines GAN-based WSI patch synthesis with per-sample gradient clipping, Gaussian noise injection, and privacy budget accounting to create synthetic images under an explicit differential privacy guarantee. The framework is intended for conceptual design, dataset augmentation planning, and governance-oriented AI development rather than for reporting experimental performance.

Its central advantage is that privacy protection is embedded in the training mechanism rather than added as an afterthought. This design can reduce memorization risk, constrain membership inference attacks, and allow synthetic patches to support classifier training when real pediatric tumor datasets are too small for robust deep learning. At the same time, utility must be judged by real-slide evaluation, pathologist review, and privacy auditing rather than visual realism alone.

Future implementation should occur within pediatric rare-cancer curation consortia that can combine archival slides, expert annotation, privacy governance, and prospective validation. Consortia such as international neuroblastoma and pediatric oncology research networks could provide the multicenter structure needed to test whether DP-generated synthetic histopathology improves model development without exposing individual patients. The long-term goal is not to replace real pathology data, but to make rare pediatric cancer AI research more scalable, privacy-preserving, and clinically accountable.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 11 Jan 2026 Revised: 04 Mar 2026 Accepted: 02 Apr 2026
Published online: 20 July 2026

References

Kather JN, Pearson AT, Halama N, Jäger D, Krause J, Loosen SH, et al. Deep learning can predict microsatellite instability directly from histology in gastrointestinal cancer. Nat Med. 2019;25(7):1054-6.

Coudray N, Ocampo PS, Sakellaropoulos T, Narula N, Snuderl M, Fenyö D, et al. Classification and mutation prediction from non–small cell lung cancer histopathology images using deep learning. Nat Med. 2018;24(10):1559-67.

Gheisari S, Catchpoole DR, Charlton A, Kennedy PJ. Convolutional deep belief network with feature encoding for classification of neuroblastoma histological images. J Pathol Inform. 2018;9:17.

Frankel AO, Lathara M, Shaw CY, Wogmon O, Jackson JM, Clark MM, et al. Machine learning for rhabdomyosarcoma histopathology. Mod Pathol. 2022;35(9):1193-203.

Ramesh S, Dyer E, Pomaville M, Doytcheva K, Dolezal J, Kochanny S, et al. Artificial intelligence-based morphologic classification and molecular characterization of neuroblastic tumors from digital histopathology. NPJ Precis Oncol. 2024;8(1):255.

Thiesen AH, Domanskyi S, Foroughipour A, Zhang J, Sheridan TB, Neuhauser SB, et al. Multicenter histology image integration and multiscale deep learning support machine learning-enabled pediatric sarcoma classification. Cancer Res. 2026;(in press / incomplete citation).

Appenzeller A, Leitner M, Philipp P, Krempel E, Beyerer J. Privacy and utility of private synthetic data for medical data analyses. Appl Sci (Basel). 2022;12(23):12320.

Kaabachi B, Despraz J, Meurers T, Otte K, Halilovic M, Kulynych B, et al. A scoping review of privacy and utility metrics in medical synthetic data. NPJ Digit Med. 2025;8(1):60.

Adams T, Birkenbihl C, Otte K, Ng HG, Rieling JA, Näher AF, et al. On the fidelity versus privacy and utility trade-off of synthetic patient data. iScience. 2025;28(5):article number not provided.

Shokri R, Stronati M, Song C, Shmatikov V. Membership inference attacks against machine learning models. In: 2017 IEEE Symp Secur Priv. 2017. p. 3-18.

Hayes J, Melis L, Danezis G, De Cristofaro E. Logan: membership inference attacks against generative models. arXiv:1705.07663. 2017.

Chen D, Yu N, Zhang Y, Fritz M. GAN-Leaks: a taxonomy of membership inference attacks against generative models. In: Proc ACM SIGSAC Conf Comput Commun Secur. 2020. p. 343-62.

Zhang Z, Yan C, Malin BA. Membership inference attacks against synthetic health data. J Biomed Inform. 2022;125:103977.

Chen Y, Yang XH, Wei Z, Heidari AA, Zheng N, Li Z, et al. Generative adversarial networks in medical image augmentation: a review. Comput Biol Med. 2022;144:105382.

Tschuchnig ME, Oostingh GJ, Gadermayr M. Generative adversarial networks in digital pathology: a survey on trends and future potential. Patterns. 2020;1(6):article number not provided.

Jose L, Liu S, Russo C, Nadort A, Di Ieva A. Generative adversarial networks in digital pathology and histopathological image processing: a review. J Pathol Inform. 2021;12:43.

Alajaji SA, Khoury ZH, Elgharib M, Saeed M, Ahmed AR, Khan MB, et al. Generative adversarial networks in digital histopathology: current applications, limitations, ethical considerations, and future directions. Mod Pathol. 2024;37(1):100369.

Xue Y, Ye J, Zhou Q, Long LR, Antani S, Xue Z, et al. Selective synthetic augmentation with HistoGAN for improved histopathology image classification. Med Image Anal. 2021;67:101816.

Levine AB, Peng J, Farnell D, Nursey M, Wang Y, Naso JR, et al. Synthesis of diagnostic quality cancer pathology images by generative adversarial networks. J Pathol. 2020;252(2):178-88.

Deshpande S, Minhas F, Graham S, Rajpoot N. SAFRON: stitching across the frontier network for generating colorectal cancer histology images. Med Image Anal. 2022;77:102337.

Li W, Li J, Polson J, Wang Z, Speier W, Arnold C. High-resolution histopathology image generation and segmentation through adversarial training. Med Image Anal. 2022;75:102251.

Deshpande S, Dawood M, Minhas F, Rajpoot N. SynCLay: interactive synthesis of histology images from bespoke cellular layouts. Med Image Anal. 2024;91:102995.

Liu Q, Zhou T, Cheng C, Ma J, Hoque TM. Hybrid generative adversarial network based on frequency and spatial domain for histopathological image synthesis. BMC Bioinformatics. 2025;26:29.

Cechnicka S, Ball J, Baugh M, Reynaud H, Simmonds N, Smith AP, et al. URCDM: ultra-resolution image synthesis in histopathology. In: Int Conf Med Image Comput Comput Assist Interv. 2024; p. 535-45.

Torfi A, Fox EA, Reddy CK. Differentially private synthetic medical data generation using convolutional GANs. Inf Sci. 2022;586:485-500.

Ziller A, Usynin D, Braren R, Makowski M, Rueckert D, Kaissis G. Medical imaging deep learning with differential privacy. Sci Rep. 2021;11:13524.

Adnan M, Kalra S, Cresswell JC, Taylor GW, Tizhoosh HR. Federated learning and differential privacy for medical image analysis. Sci Rep. 2022;12:1953.

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;(early online / incomplete citation).

Achterberg J, Haas M, van Dijk B, Spruit M. Fidelity-agnostic synthetic data generation improves utility while retaining privacy. Patterns. 2025;6(10):article number not provided.

Wu H, Gao R, Sheng YP, Chen B, Li S. SDAE-GAN: enable high-dimensional pathological images in liver cancer survival prediction with a policy gradient based data augmentation method. Med Image Anal. 2020;62:101640.

Pantanowitz J, Manko CD, Pantanowitz L, Rashidi HH. Synthetic data and its utility in pathology and laboratory medicine. Lab Invest. 2024;104(8):102095.

Zormpas-Petridis K, Noguera R, Ivankovic DK, Roxanis I, Jamin Y, Yuan Y. SuperHistopath: a deep learning pipeline for mapping tumor heterogeneity on low-resolution whole-slide digital histopathology images. Front Oncol. 2021;10:586292.