MINISTRY OF EDUCATION AND RESEARCH
NATIONAL UNIVERSITY OF SCIENCE AND TECHNOLOGY
BUCHAREST POLYTECHNICA

Doctoral School of Electronics, Telecommunications and Information Technology

Ph.D. THESIS
SUMMARY

Alexandra-Georgiana ANDREI

SISTEME DE INTELIGENȚĂ ARTIFICIALĂ PENTRU DIAGNOSTIC ȘI MONITORIZARE

ARTIFICIAL INTELLIGENCE SYSTEMS FOR DIAGNOSIS AND MONITORING OF MEDICAL CONDITIONS

Prof. Dr. Ing. Bogdan IONESCU
POLITEHNICA Bucharest

PhD Supervisor

BUCHAREST 2026

Table of contents

1 Introduction **1**

1.1 Presentation of the field of the doctoral thesis 1

1.2 Scope of the doctoral thesis 2

1.3 Content of the doctoral thesis 2

2 Theoretical Concepts **3**

2.1 The need for synthetic medical images 3

2.2 Data privacy in biomedical imaging 3

2.3 Medical imaging modalities 4

2.4 Generative models 4

3 Identifying Trends in Generative Adversarial Networks **5**

3.1 Scope and motivation 5

3.2 Methods 5

3.3 Results 5

3.4 Conclusions 6

4 Machine Learning Methods for Neurological Disorder Analysis **7**

4.1 Scope and motivation 7

4.2 EEG epilepsy systematic review 7

4.3 Freezing of gait detection 8

4.4 Conclusions 8

5 Privacy-Preserving Histopathological Image Augmentation **9**

5.1 Scope and motivation 9

5.2 Methods 9

5.3 Results 11

5.4 Conclusions 11

6 Investigating Generative Models' Impact on Biomedical Synthetic Images **12**

6.1 Introduction and motivation 12

6.2 Methods 12

6.3 Results 13

6.4 Cross-edition conclusions 14

7 Utility-Privacy Trade-offs in Synthetic Medical Image Generation 15
7.1 Introduction and motivation 15
7.2 Methods and evaluation framework 15
7.3 Results 16
7.4 Conclusions 16

8 Conclusions 17
8.1 Obtained results 17
8.2 Original contributions 18
8.3 List of original publications 20
8.4 Perspectives for further developments 23

References 27

Chapter 1

Introduction

Artificial Intelligence (AI) and Machine Learning (ML) have become central to biomedical research and clinical practice, enabling data-driven approaches to disease diagnosis, patient monitoring, and treatment planning across a broad spectrum of medical specialties. The ability to extract meaningful patterns from complex, high-dimensional data, such as images, biosignals, and wearable sensor recordings has opened new avenues for automated tools that support clinical decision-making. Despite this promise, deploying reliable AI systems in healthcare remains challenging. Medical data are scarce, sensitive, costly to annotate, and subject to strict regulatory constraints such as the General Data Protection Regulation (GDPR), making it difficult to build models that are simultaneously effective and trustworthy. This thesis addresses several of these challenges through contributions spanning multiple medical domains and data modalities. The work is united by a concern for developing AI methods that are not only performant but also responsible: producing clinically meaningful outputs, operating reliably under data constraints, and handling patient information with appropriate care for privacy.

1.1 Presentation of the field of the doctoral thesis

Medical imaging is one of the primary tools supporting diagnosis and treatment planning. The application of AI to medical images has demonstrated considerable promise, yet progress is consistently delayed by the limited availability of large, diverse, and wellannotated datasets. Generative models, particularly Generative Adversarial Networks (GANs) [7] and diffusion-based models [11] have emerged as a compelling approach to addressing data scarcity by producing realistic synthetic images that reproduce the statistical properties of real clinical data. However, their application raises important questions about data utility, clinical relevance, and patient privacy, since models trained on sensitive data may inadvertently retain identifiable traces of their training samples. Beyond medical imaging, similar challenges arise in the analysis of physiological and wearable sensor data. Neurological disorders such as epilepsy (affecting over 50 mil-

lion people worldwide [29]) and Parkinson's Disease (PD) (affecting over 10 million globally [23]) present particularly difficult data environments: EEG recordings for epilepsy are subject to strong inter-subject variability and scarce expert annotations, while wearable-based monitoring of motor symptoms in PD requires robust classifiers operating under real-world conditions with limited labelled data.

1.2 Scope of the doctoral thesis

The scope of this thesis is to investigate the application of AI and ML to medical data across multiple modalities and clinical contexts, with a focus on three interrelated objectives: advancing synthetic medical data generation, assessing the privacy implications of generative models, and developing machine learning methods for clinical applications where data availability is limited.

Within the domain of synthetic image generation, the thesis examines generative architectures applied to histopathology and CT imaging, proposes privacy-preserving data augmentation frameworks, and systematically benchmarks generative models with respect to the trade-offs between image quality, downstream utility, and privacy risk. A central question is whether synthetic images are both clinically useful and privacy-safe, or whether they retain identifiable traces of the real patient data used during training.

Beyond image generation, the thesis addresses data-scarce clinical settings through two contributions in neurological disorder analysis: a systematic review of unsupervised and self-supervised learning methods for EEG-based epilepsy research, and an experimental study on the automatic detection of Freezing of Gait in Parkinson's disease from accelerometer data.

1.3 Content of the doctoral thesis

The thesis is structured in two parts. Part I (Chapters 1-2) establishes the theoretical and conceptual background. Part II (Chapters 3-7) presents the original contributions: a systematic review of GAN-based medical image synthesis (Chapter 3), unsupervised learning for neurological disorders (Chapter 4), privacy-preserving histopathological image augmentation (Chapter 5), a multi-year benchmarking campaign on generative model privacy (Chapter 6), and and a comparative study of utility–privacy trade-offs in CT synthesis (Chapter 7). Chapter 8 summarises the conclusions, original contributions, publication list, and future directions.

Chapter 2

Theoretical Concepts

2.1 The need for synthetic medical images

Tasks such as disease detection, segmentation, and outcome prediction require access to large, diverse, and well-labelled datasets. In clinical practice, three barriers consistently limit such access. First, legal and institutional constraints prevent most hospitals from sharing raw patient data. Second, annotation demands domain expertise from radiologists or pathologists, making the process costly and subject to inter-observer variability. Third, rare diseases and atypical presentations are structurally underrepresented, leading to class imbalance in any collected dataset. Synthetic medical images address these barriers.

2.2 Data privacy in biomedical imaging

Despite their potential, generative models introduce important privacy and security concerns. Recent studies have demonstrated that both GANs and diffusion models can memorise, replicate, or encode identifiable characteristics of the real images on which they were trained. In biomedical contexts this is especially problematic: imaging data may contain sensitive biometric or diagnostic information about identifiable patients. Three attack paradigms have been documented in the literature: linkage attacks (re-identifying a patient by matching synthetic to real records), attribute inference (deducing protected attributes from synthetic outputs), and membership inference (determining whether a specific real image was used during training) [26, 19, 18, 28]. Ensuring that synthetic medical images are both clinically useful and genuinely privacy-safe requires rigorous, standardised evaluation frameworks, representing a gap that directly motivates the privacy-focused contributions of this thesis.

2.3 Medical imaging modalities

The thesis focuses on two primary imaging modalities. Histopathology involves microscopic examination of tissue sections stained with Haematoxylin and Eosin (H&E), providing detailed cellular morphology critical for cancer diagnosis and grading. Whole-slide images (WSIs) are typically divided into smaller patches to facilitate automated analysis [8]. Computed Tomography (CT) uses X-ray measurements to produce cross-sectional images of the body, enabling volumetric analysis of anatomical structures, with lung tuberculosis imaging serving as the test domain for Chapters 6 and 7.

2.4 Generative models

Generative Adversarial Networks (GANs) consist of a generator G and a discriminator D trained adversarially. G maps a latent vector $z \sim p_z(z)$ to a synthetic sample, while D classifies each image as real or fake [7]. Training is governed by the minimax objective:

$$\min_G \max_D V(D,G) = \mathbb{E}_{x\sim p_{\text{data}}}[\log D(x)] + \mathbb{E}_{z\sim p_z}[\log(1-D(G(z)))]. \tag{2.1}$$

Variants of GANs used in this thesis include DCGAN [24], cGAN [20], BigGAN [6], WGAN [4], Pix2pix [13], CycleGAN [30], StyleGAN [16], and PGGAN [15].

Denoising Diffusion Probabilistic Models (DDPMs) generate images by learning to reverse gradual noise-addition procedure over T steps [11]. The forward process $q(x_t\,|\,x_{t-1}) = \mathcal{N}(x_t; \sqrt{1-\beta_t}\,x_{t-1}, \beta_t I)$ corrupts data over T steps into isotropic Gaussian noise. The data distribution is progressively recovered using a a neural network then learns the reverse process $p_\theta(x_{t-1}\,|\,x_t) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t,t), \Sigma_\theta(x_t,t))$ Both unconditional and conditional (cDiffusion) variants are evaluated in this thesis.

Evaluation metrics. Image quality and distributional fidelity are assessed using the Fréchet Inception Distance (FID) [10], Inception Score (IS) [25], and precision/recall-based metrics (α- precision, β-recall, density, coverage). Privacy risk is quantified through membership inference attack success rates.

Chapter 3

Identifying Trends in Generative Adversarial Networks

3.1 Scope and motivation

GANs have become one of the most widely used models for realistic medical image synthesis, with applications spanning radiology, histopathology, ophthalmology, and dermatology. The rapid expansion of GAN-based methods has produced a large but fragmented literature. This chapter presents a systematic review of GAN-based medical image synthesis conducted in accordance with PRISMA guidelines [22], providing a unified methodological foundation for the experimental contributions of subsequent chapters.

3.2 Methods

A structured search string combining terms for Generative Adversarial Networks with medical imaging modalities was applied to PubMed, IEEE Xplore, and Scopus. Inclusion criteria required studies to: (i) use GAN-based methods for medical image generation or translation; (ii) target at least one medical imaging modality; (iii) report quantitative evaluation. A total of 11,592 records were identified. After removing duplicates, title/abstract screening, and full-text assessment, 286 studies were retained for in-depth analysis.

3.3 Results

Publication trends and datasets. Of the 286 included studies, 227 (79.2%) were published within the last five years. Regarding dataset accessibility, 137 articles used exclusively public datasets, 127 relied on private data, and 22 combined both. MRI is the most frequently used modality, followed by X-ray and retinal fundus imaging. CT and

histopathology are also well represented modalities. The dominance of neuroimaging and thoracic imaging reflects the availability of large public benchmarks in these areas.

Pre-processing and architectures. Pre-processing steps vary by modality: CT data undergo HU windowing and intensity normalisation; MRI data require skull stripping and bias field correction; histopathology slides require stain normalisation (Reinhard, Macenko) and patch tiling. Regarding generation paradigm, image-to-image translation is the most common, leveraging Pix2Pix (59 studies, 20.3%), CycleGAN (59 studies, 20.6%), and cGAN (58 studies, 20.3%). Unconditional synthesis is dominated by DCGAN and PGGAN variants; conditional synthesis increasingly employs StyleGAN families for high-resolution outputs.

Evaluation gaps FID and IS are the most commonly reported metrics, used in over 80% of studies. Clinical utility metrics (downstream task performance) appear in fewer than 40% of studies. Privacy assessment, verifying that synthetic images do not leak training data, is reported in only 5 studies, despite widespread claims of privacy-preserving benefits.

3.4 Conclusions

This review has highlighted the rapid adoption of GANs for medical image generation across various modalities and clinical contexts. GANs have shown strong potential in applications such as data augmentation, cross-modality translation, and pathology simulation, demonstrating their ability to produce anatomically realistic and clinically useful images. At the same time, several barriers remain. Many approaches still rely on paired datasets, evaluation practices are fragmented and often fail to capture clinical utility, and privacy considerations are rarely addressed. These challenges underline the need for more robust, clinically meaningful, and privacy-preserving frameworks. GANs have emerged as a promising tool for medical image synthesis, but significant work is still needed before they can be trusted in everyday clinical practice. Overcoming the challenges discussed in this review will be key to moving these methods beyond experimental studies and into safe, reliable, and widely used applications.

Chapter 4

Machine Learning Methods for Neurological Disorder Analysis

4.1 Scope and motivation

This chapter presents two complementary contributions addressing neurological disorders through label-efficient AI methods: (i) a systematic literature review of unsupervised and self-supervised learning applied to EEG-based epilepsy research; (ii) an experimental study on the automatic detection of Freezing of Gait (FoG) in Parkinson's Disease from tri-axial accelerometer data.

4.2 EEG epilepsy systematic review

Methods. The review was conducted following PRISMA guidelines. Five databases were searched: Embase, Cochrane, Web of Science, PubMed, and Google Scholar. Inclusion criteria required studies to apply unsupervised or self-supervised algorithms to EEG recordings from human subjects with epilepsy, published after 2013 in English. A total of 11,592 records were identified. 108 studies were retained, after screening.

Results. Of the 108 included studies, 82 (75.9%) used exclusively public datasets with CHB-MIT [27] and Bonn University datasets [3] being used most frequently. Seizure detection was the dominant task (57% of studies), followed by seizure prediction (25%), patient characterisation (2%), seizure localisation (4%), and type classification (22%). Two modelling paradigms were identified: patient-specific models (42 studies, showing higher peak accuracy but limited generalisability) and generalised multi-patient models (63 studies, showing more stable cross-patient performance). Methodologically, autoencoders and clustering-based approaches dominate, while contrastive and self-supervised methods are an emerging trend. Seventy out of 108 studies (73%) did not report artifact removal procedures, raising reproducibility concerns. The Short-Time

Fourier Transform was the most widely used time-frequency representation. The lack of standardised evaluation protocols is identified as the central limiting factor.

4.3 Freezing of gait detection

Methods. FoG it is a transient inability to initiate or sustain walking, one of the most disabling motor symptoms of PD. Automatic detection from wearable sensors is clinically valuable for monitoring disease progression and triggering assistive interventions [9]. Five deep learning frameworks were evaluated: two AlexNet variants (AlexNet-A with an additional hidden layer; AlexNet-B without), LeNet, VGG-16, and an autoencoder-based feature extractor followed by an SVM classifier. All models received as input a segment from 2, 3,4, and 5-second, 50%-overlapping windows of tri-axial accelerometer signals. The study systematically evaluated the effect of sensor placement (ankle, thigh, trunk), signal axis (X, Y, Z, magnitude), and dataset balancing (undersampling, no balancing) on detection performance. The Daphnet FoG dataset was used [5], comprising 8 h 20 min of recordings from 10 PD patients.

Results. All five frameworks achieved peak accuracies of up to 90.92%. The autoencoder and SVM and LeNet frameworks provided the most consistent results across sensor placements. Thigh and trunk sensors consistently produced the most discriminative signals with ankle sensors yielded more variable performance. Magnitude-based inputs outperformed individual-axis inputs across all architectures. The best configuration (autoencoder, thigh, magnitude) achieves 90.92% accuracy, outperforming comparable CNN-based approaches from the literature by 0.8-7%.

4.4 Conclusions

The EEG epilepsy review consolidates a large, heterogeneous literature and identifies standardised evaluation and dataset diversity as priority gaps. The FoG detection study provides actionable guidance for sensor placement and architecture selection, with results that surpass prior work. Both studies highlight data availability and evaluation standardisation as primary areas for future development.

Chapter 5

Privacy-Preserving Histopathological Image Augmentation

5.1 Scope and motivation

Histopathological examination of H&E stained tissue sections is essential for diagnosing colorectal cancer (CRC), the third most common cancer globally, and breast cancer [21]. Developing robust computational models is limited by the availability of large, well-annotated datasets. This chapter proposes a Conditional Deep Convolutional GAN (cDCGAN) for multi-class histopathological image synthesis and rigorously evaluates both the clinical quality and privacy properties of the generated images.

5.2 Methods

Proposed cGAN architecture. The proposed cDCGAN builds on DCGAN [24] and cGAN [20], with architectural modifications for multi-class histopathology synthesis. Class labels are used as conditioning signals for both generator and discriminator. The conditional generator comprises five transposed-convolution layers (kernel 4, 128 filters, stride 2), merging a latent vector of size 1500 with class one-hot labels via a concatenation layer. Each deconvolutional layer is followed by LeakyReLU activation with a slope of 0.2, and the output uses Tanh function. The discriminator contains five convolutional blocks with 4$times$4 convolutions (stride 2), ending in a sigmoid output. The model is trained with batch size 125, Adam optimiser with a learning rate of 0.0002, $beta_1$ of 0.5, and cross-entropy loss on 79,026 training images from the NCT-CRC-HE-100K dataset [17]. Seven tissue classes are included: DEB, LYM, MUC, MUS, STR, TUM, NORM.

Privacy assessment. The privacy investigation adopts the ImageCLEFmedical GANs protocol [2]: given a set of generated images G and a mix of real images containing

images used and not used during training, the task is to identify which real images influenced training. The proposed approach uses two stages: (i) deep feature extraction via pretrained CNNs (ResNet50, EfficientNet, DenseNet121); (ii) unsupervised clustering (k-means and hierarchical) of generated image features. The hypothesis is that training images lie closer to the cluster structure of generated images than unseen images.

Downstream utility evaluation. Utility is measured via downstream classification performance. Four architectures (DenseNet, ResNet-50, EfficientNet, MobileNetV2) are trained on mixed real and synthetic data and compared to real-only baselines. A Turing test with four histopathology specialists additionally assessed perceptual realism.

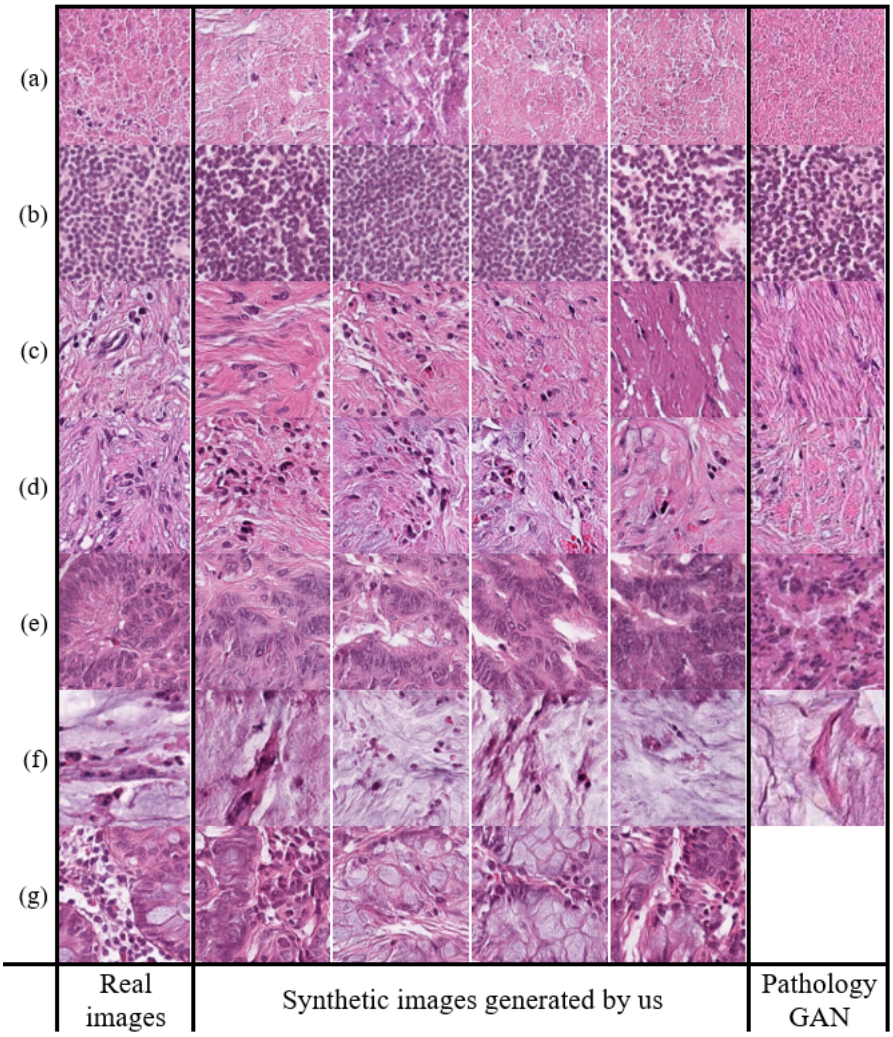


Fig. 5.3 Examples of real and synthetic images for NCT-CRC-HE-100K colorectal cancer dataset: (a) DEB, (b) LYM, (c) MUS, (d) STR, (e) TUM, (f) MUC, and (g) NORM.

Table 5.2 FID score comparison with state-of-the-art methods on the CRC dataset (R = real, S = synthetic images).

Method	**5k R vs 5k S**	**10k R vs 10k S**
PathologyGAN [1]	83.52	56.03
Conditional Diffusion [12]	292.86	286.25
StyleGAN [16]	202.42	194.01
Proposed	**50.14**	**41.79**

5.3 Results

Image quality. Figure 5.3 shows examples of real and generated images across all seven tissue classes. Generated images capture characteristic morphological properties in terms of texture, staining patterns, cellular organisation at a quality that is difficult to distinguish from real images in blind evaluation. Quantitative evaluation yields a mean FID of 64.09 across all classes (LYM achieves 21.2). As shown in Table 5.2, the proposed method substantially outperforms off-the-shelf baselines.

Perceptual evaluation and privacy. The Turing test shows that specialists were unable to reliably distinguish real from synthetic images, with average detection accuracy near chance across most tissue types.

Privacy assessment: fingerprint detection accuracy remains consistently around 50% regardless of the feature extractor, distance metric, or clustering algorithm employed, confirming that the generated images do not retain identifiable training data fingerprints.

Downstream utility. Augmenting training data with synthetic patches consistently improves classifier accuracy , with gains of 1.83–35.36% depending on architecture. Generalization on an independent breast cancer dataset confirms the framework's utility beyond the training domain.

5.4 Conclusions

This chapter demonstrates that conditional GAN-based synthesis can simultaneously achieve high image quality, downstream utility, and privacy preservation. The proposed cDCGAN outperforms PathologyGAN and off-the-shelf diffusion/StyleGAN baselines on FID, while generating images that are perceptually indistinguishable from real tissue and that do not expose training data fingerprints.

Chapter 6

Investigating Generative Models' Impact on Biomedical Synthetic Images

6.1 Introduction and motivation

Despite growing use of generative models in medical imaging, systematic evaluations of privacy leakage in synthetic images remain scarce. To address this gap, the ImageCLEFmedical GANs benchmarking series was launched as part of ImageCLEF, providing a structured framework for assessing how synthetic images encode, preserve, or reveal information about their training data[1]. Three central research questions are addressed: (1) are training-data fingerprints detectable in synthetic images? (2) can the generative architecture that produced a synthetic image be identified? (3) can synthetic images be attributed to specific training data subsets?

6.2 Methods

Each edition provided participants with curated datasets of real and synthetic biomedical CT images generated by different architectures, alongside a standardised evaluation protocol. The generative models covered include cGAN, BigGAN, DCGAN, and diffusion models. The AI Multimedia Lab's methods evolved across editions: a feature-extraction + clustering pipeline in 2023; deep CNN classification in 2024; and a Siamese network architecture in 2025. For Subtask 2 (model fingerprint detection), features are extracted using MobileNetV2, EfficientNet, and DenseNet-121, PCA reduces dimensionality (95%or 90% variance threshold), and k-means or agglomerative hierarchical clustering groups images by generating architecture.

[1]Based on six publications: overview and team working notes for the 2023, 2024, and 2025 editions of the ImageCLEFmedical GANs task, published in CLEF Working Notes (CEUR-WS) and ECIR proceedings.

Table 6.1 Subtask 1 dataset summary and cross-edition results. Best/worst refer to all participating teams.

Edition	Model	Best (all teams)	AIMultimediaLab	Worst
2023	Diffusion	F1 = 0.802 (VCMI)	F1 = 0.626 (#6/40)	0.093
2024	Diff.+cGAN	F1 = 0.667 (Inoue Koki)	F1 = 0.593 (#7/54)	0.019
2025	StyleGAN	κ = 0.148 (Neural Nexus)	κ = 0.036 (#10/33)	−0.176

6.3 Results

Subtask 1: Identify Training Data Fingerprints. The task is formulated as binary classification ("used for training" and "not used") and was evaluated with F1-score (2023–2024) and Cohen's kappa (2025).

AI Multimedia Lab approaches. In 2023, a DCGAN was trained on held-out real images and its outputs used as a reference; LBP texture features and VGG-16 deep features were fed to a binary SVM. LBP outperformed deep features on the test set (F1 of 0.626 vs. 0.585), suggesting texture captures more stable fingerprint characteristics. In 2024, two methods were proposed: (i) lung-region segmentation via U-Net followed by ResNet-50/DenseNet-121 feature extraction and hierarchical clustering, and (ii) autoencoder anomaly detection using reconstruction error (MSE or SSIM). The SSIM-based autoencoder achieved the best result on the cGAN dataset with an F1 0f F1 = 0.594. In 2025, a Siamese neural network trained with contrastive loss on synthetic-real image pairs achieved $\kappa = 0.036$ on the official test set.

Findings. In 2023-2024, multiple teams achieved clearly above-chance performance (F1 up to 0.80), confirming that diffusion models and cGANs leave detectable image-level fingerprints. The 2025 StyleGAN edition was markedly harder: all 33 participants achieved near-zero kappa, and the highest accuracy was only 0.574, suggesting StyleGAN is substantially more resistant to image-level membership inference. Whether this reflects a genuine architectural privacy improvement or the limitations of current analysis techniques remains an open question.

Subtask 2: Detect Generative Model Fingerprints. Introduced in 2024, this subtask asks whether synthetic images can be clustered by the generative model that produced them, without supervision. The training set contained synthetic CT images from three models and the test set contained images from four models. The official metric was the Adjusted Rand Index (ARI).

The AI Multimedia Lab pipeline extracted features from five ImageNet-pretrained CNNs (VGG-16, ResNet-50, MobileNetV2, EfficientNet, DenseNet-121) and LBP, then applied k-means and hierarchical clustering. On the development set, DenseNet-121 achieved a perfect ARI of 1.0. On the official test set, MobileNetV2 with hierarchical

clustering reached ARI = 0.9971 and DenseNet-121 reached ARI = 0.9965, placing the team 2nd out of 46 submissions (best team: ARI = 1.0 via t-SNE + k-means).
Findings. The striking contrast with Subtask 1 is the central finding. Whereas image-level membership inference was difficult, clustering by generative model was near-trivial. Different architectures imprint distinct, stable signatures in texture, frequency content, or structural artefacts that are reliably detectable, even when a new unseen model is introduced at test time.

Subtask 3: Identify Training Data Subsets. Introduced in 2025, this subtask shifts the analysis from image-level to dataset-level privacy: each synthetic image must be attributed to the specific subset of real images that trained the corresponding generative model. Five subsets ($t1$–$t5$) were used to train a diffusion mode. The test set contained synthetic images spanning thoracic, cervical, and abdominal CT, increasing anatomical variability. Accuracy was the official metric.

Two methods were proposed. Method 1 combined ImageNet CNN features (MobileNetV2, ResNet-50, EfficientNet, DenseNet-121) with k-means, hierarchical clustering, or SVM. The best result was obtained using ResNet-50 + SVM at 52.36%. Method 2 used a Siamese neural network (four convolutional blocks, 256-unit L2-normalised embedding, contrastive loss), assigning test images to the nearest subset centroid. On the development set, the SNN achieved 100% accuracy; on the official test set it reached **99.04%**, ranking 1st out of 11 submissions.
Findings. The result demonstrates that diffusion-generated images preserve strong statistical cues tied to their training subset, such as anatomical patterns, lesion prevalence, and acquisition characteristics, enabling near-perfect dataset-level attribution. This constitutes a significant and underappreciated privacy risk: even when individual image membership is hard to infer, cohort-level attributes can be fully recovered from synthetic images alone.

6.4 Cross-edition conclusions

Three consistent findings emerge: (i) diffusion-based models leave more detectable trainingdata traces than GANs; (ii) StyleGAN variants show the lowest susceptibility to both fingerprint detection and model attribution; (iii) attribution of synthetic images to training data subsets is feasible, demonstrating that privacy risks extend beyond simple membership inference. These findings motivate the development of privacy-by-design generative frameworks and highlight the inadequacy of assuming privacy preservation without empirical verification.

Chapter 7

Utility-Privacy Trade-offs in Synthetic Medical Image Generation

7.1 Introduction and motivation

This chapter presents a comprehensive comparative experimental study of six representative CT image synthesis architectures: DCGAN, cGAN, StyleGAN1, StyleGAN2, unconditional Diffusion, and conditional Diffusion (cDiffusion). The central research question is whether image quality, downstream utility, and privacy risk are correlated across architectures, or whether they represent independent dimensions requiring separate assessment.

7.2 Methods and evaluation framework

Six architectures were trained on the same CT chest dataset. The evaluation framework integrates six complementary dimensions: (i) Fidelity and diversity metrics: FID, IS, α-precision, β-recall, density, and coverage; (ii) Anatomical plausibility: a UNet-based lung segmentation model verifies the presence, location, and area proportions of anatomical structures in synthetic images; (iii) Human perceptual study (Turing test): radiologists evaluate synthetic images in single-image and multi-image tasks, rating each as real, fake, or uncertain; (iv) Explainability (Grad-CAM): gradient-based class activation maps reveal which image regions discriminator networks use to distinguish real from synthetic samples; (v) Privacy assessment: membership inference attacks quantify the risk that synthetic images leak identifiable training-data information, and (vi) Visual embedding overlap: t-SNE visualisations of deep feature embeddings quantify distributional overlap and mode coverage between real and synthetic image sets.

7.3 Results

Image quality and distributional fidelity. Conditional Diffusion achieves the most balanced performance: low FID (56.4) and JSD (0.014), high precision (0.712), recall (0.375), density (0.576), and coverage (0.636). DCGAN exhibits near-zero precision, recall, density, and coverage indicating mode collapse. StyleGAN2 achieves moderate FID (57.9) but very low recall (0.0035) and limited coverage (0.182), confirming high visual quality restricted to a narrow region of the data manifold.

Human perceptual evaluation. In the single-image task, cGAN images were the easiest to detect as synthetic (93.3% correct), while StyleGAN1 was the most convincing (only 39.4% correctly detected, 59.2% misclassified as real). Strikingly, real CT images were correctly identified in only 38.1% of cases, highlighting both the inherent difficulty of the task and the high visual quality of the best generative models. In the multi-image task, overall discriminability improved for all models as side-by-side comparison revealed subtle inconsistencies invisible in single-image evaluation.

Privacy assessment. Membership inference results confirm that privacy risk is strongly architecture-dependent. Diffusion models (both unconditional and conditional) show the highest membership inference susceptibility, consistent with Chapter 6. StyleGAN variants ,despite achieving the highest perceptual realism, exhibit the lowest susceptibility, establishing them as the most privacy-safe architectures evaluated. Grad-CAM analysis reveals that discriminators identify GAN images through local texture artefacts and diffusion images through global structural inconsistencies.

7.4 Conclusions

No single architecture dominates across all evaluation dimensions: cDiffusion is the most balanced but least private; StyleGAN variants are the most realistic and most private but the least diverse; DCGAN suffers from mode collapse. These findings mandate multifaceted evaluation frameworks as standard practice for synthetic medical image validation and underscore that privacy-by-design approaches are needed for clinical deployment.

Chapter 8

Conclusions

This thesis investigates the application of AI and ML to medical data across multiple modalities and clinical contexts, contributing to the development of trustworthy and effective systems for healthcare. The contributions address synthetic medical image generation and privacy assessment, unsupervised learning for neurological disorder analysis, and machine learning for wearable-based motor symptom detection. Together, they advance the understanding of how AI methods can be designed, evaluated, and deployed responsibly under the constraints typical of clinical environments.

8.1 Obtained results

Part I introduces the theoretical foundations of the thesis. Chapter 2 establishes the conceptual background, covering the motivation for synthetic medical data generation, data privacy challenges in biomedical imaging, and the generative model architectures investigated throughout the thesis.

Part II presents the original scientific contributions. A systematic review of GAN-based methods for medical image synthesis establishes the landscape of architectures, datasets, and evaluation practices, identifying limited attention to privacy and downstream utility as key methodological gaps. A conditional deep convolutional GAN is trained to synthesize histopathological image patches across seven tissue categories, demonstrating that synthetic augmentation simultaneously improves downstream classifier performance and preserves patient privacy. Expert validation by histopathology specialists confirms image realism, and privacy analyses confirm the absence of identifiable training data fingerprints. Generalization is further validated on an independent breast cancer dataset.

A large-scale benchmarking campaign conducted over three consecutive editions (2023-2025) provides a controlled analysis of privacy behaviors across generative architectures. Recurring training data fingerprint detection experiments reveal that diffusion-based models consistently leave detectable traces of their training data, while StyleGAN-

based architectures demonstrate the lowest susceptibility to membership inference. Later editions further establish that generative models encode distinct model-specific fingerprints, and that attribution of synthetic images to training data subsets is feasible with high accuracy. A comprehensive comparative study of six CT image synthesis models (DCGAN, cGAN, StyleGAN1, StyleGAN2, Diffusion, and cDiffusion) further confirms that quality, utility, and privacy cannot be reliably inferred from one another, motivating the adoption of multi-faceted evaluation frameworks. Conditional diffusion achieves the most balanced performance across fidelity, diversity, and coverage, while StyleGAN variants attain the highest perceptual realism.

A systematic review of label-efficient methods for EEG-based epilepsy research examines how clustering, autoencoders, and contrastive learning handle limited annotations and inter-subject variability in clinical recordings. The review identifies the lack of standardized evaluation protocols as a key limitation in the field.

An experimental study on the automatic detection of Freezing of Gait in Parkinson's disease uses 3D accelerometer data to systematically compare five deep learning frameworks: two AlexNet variants, LeNet, VGG-16, and an autoencoder. The study evaluates the effect of sensor placement, signal axis, segment length, and dataset balancing on detection performance. All frameworks achieve accuracies of up to 90.92%, with the autoencoder and LeNet providing the most consistent results. Thigh and trunk sensors yield the most discriminative signals, and magnitude-based inputs consistently produce the strongest performance, outperforming comparable CNN-based approaches from the literature by 0.8-7% in accuracy.

8.2 Original contributions

- [J1] develops a Conditional Generative Adversarial Network for generating realistic colorectal histopathology images across seven tissue classes. My main contribution consists of designing the cGAN architecture, validating the realism of the generated images through expert evaluation and benchmarking against state-othe-art methods, and assessing privacy preservation by verifying the absence of identifiable data leakage in the synthetic images. I further demonstrate the generalization capability of the proposed model and its positive impact on downstream classification performance using diverse colorectal cancer datasets.

- [J2] presents a systematic review of unsupervised and self-supervised learning approaches applied to EEG data for epilepsy research. The main contributions include a comprehensive analysis of methods for seizure detection, prediction, and patient characterization without reliance on labeled data, along with a critical examination of model architectures, feature learning strategies, datasets, and evaluation protocols. Key challenges such as data scarcity, high inter-subject

variability, and the lack of standardized benchmarks are identified, and future research directions for label-efficient learning in EEG-based epilepsy studies are outlined.

- [J3] investigates the use of deep learning techniques for the automatic detection and staging of membranous glomerulonephritis from electron microscopy images. The study proposes a two-stage framework based on Vision Transformer architectures, where the first stage detects glomerular membrane regions and the second stage performs disease staging. My contributions include participating in the development and evaluation of the proposed framework, contributing to the experimental design, and analyzing the model's performance on both internal and external datasets.

- [J4] presents a systematic review of Generative Adversarial Networks for medical image generation. It provides a modality-agnostic overview of the field, analyzing trends in dataset usage, pre-processing strategies, GAN architectures, and evaluation metrics. The review covers GAN applications for unconditional image generation, cross-modality synthesis, and image-to-image translation. By identifying key challenges and open research directions, this work establishes a foundational pillar of the thesis, with my contributions directly addressing the identified limitations and advancing the use of generative models in medical imaging.

- [J5] presents a comprehensive comparative study of generative models for medical image synthesis, focusing on the trade-offs between image quality, downstream utility, and privacy. The work evaluates six representative generative architectures under a unified experimental framework, combining quantitative metrics, embedding-based analyses, trait-based evaluations, and human perceptual studies. By systematically analysing model behavior across fidelity, diversity, anatomical coverage, and privacy-related criteria, the study highlights the limitations of single-metric evaluation and demonstrates how different generative models occupy distinct positions along the privacy-utility-quality spectrum.

- [C1, C4, C6, C7, C8, C9, B1, B2, B3] investigate the impact of generative models on biomedical image synthesis by proposing an evaluation framework to assess how synthetic images generated by different models encode, preserve, or reveal information about their training data. My contributions consist of organizing and coordinating a large-scale international benchmarking campaign conducted over three consecutive years, designing the evaluation tasks, and curating the datasets used for experimentation. I further analyzed the influence of different generative architectures on the privacy properties of synthetic biomedical images and provided comprehensive datasets of both real and generated images for each edition

of the benchmark. Across the three editions, the tasks evolved to address multiple research questions related to the influence of training data and generative model design on the characteristics and potential privacy risks of synthetic biomedical images. In addition, [J4] provides a comprehensive comparative analysis of synthetic images produced by multiple generative models, examining their behavior from both privacy and utility perspectives.

- [C2, C3, C5] propose novel frameworks for privacy assessment in biomedical image generation by addressing the tasks introduced in the international benchmarking campaign described above. These works investigate methods for identifying traces of training data within synthetic images and for analyzing whether generated samples reveal sensitive information about the original datasets. My contributions include designing and implementing deep learning–based solutions for detecting training data fingerprints in generated images and systematically evaluating their performance across different generative architectures and datasets. The proposed approaches consistently ranked among the top-performing methods in the benchmark evaluations, demonstrating their effectiveness for analyzing privacy risks in synthetic biomedical image generation.
- [C10, C11] investigate the automatic detection of freezing of gait (FoG) episodes in patients with Parkinson's disease using wearable sensor data. These works explore the application of deep learning methods to analyze motion signals captured by tri-axial accelerometers and to identify patterns associated with FoG events. My contributions consist of designing and evaluating deep convolutional neural network architectures for FoG detection from three-dimensional accelerometer signals, studying the influence of different network structures and signal representations on detection performance, and comparing the proposed approaches with existing methods for gait analysis in Parkinson's disease.

8.3 List of original publications

In this section, the list of all published papers and works submitted is presented. Fig. 8.1 shows a list of works published and available on Google Scholar on the 3rd of January 2026.

Journal papers:

(J1) **A.-G. Andrei**, M.G. Constantin, M. Graziani, H. Müller, B. Ionescu, "Privacy Preserving Histopathological Image Augmentation with Conditional Generative Adversarial Networks", Elsevier Pattern Recognition Letters, vol. 188, pp. 185-192, https://doi.org/10.1016/j.patrec.2024.12.014, ISSN: 0167-8655, 2025 (Q2 journal, Impact Factor: 3.3, WOS: 001407367700001).

(J2) A.-M. Tautan, **A.-G. Andrei**, C.L. Smeralda, G. Vatti, S. Rossi, B. Ionescu, "Unsupervised Learning from EEG Data for Epilepsy: A Systematic Literature Review", Elsevier Artificial Intelligence in Medicine, vol. 162,ISSN: 0933-3657, DOI:https://doi.org/10.1016/j.artmed.2025.103095, 2025 (Q1 journal, Impact Factor: 6.2, WOS:001437200900001)

(J3) (Accepted) Mihai Gabriel Constantin, George Terinte-Balcan, Ioana Maria Lambrescu, Tudor Emanuel Fertig, Oana Bianca Lazăr, **Alexandra-Georgiana Andrei**, Nicoleta Petre, Bogdan Ionescu, Mihaela Gherghiceanu. Two-stage deep learning networks for the diagnosis and staging of membranous glomerulonephritis from electron microscopy images. Laboratory Investigation. ISSN: 0023-6837 (Q2 journal, Impact Factor: 4.2)

(J4) (Under review) **Andrei AG**, Ionescu B. Generative Adversarial Networks for Medical Image Generation: A Systematic Review. IEEE Journal of Biomedical and Health Informatics. ISSN: 2168-2208 (Q1 journal, Impact Factor: 6.7).

(J5) (In preparation) **Andrei AG**, M.G. Constantin, M. Dogariu, B. Ionescu, "Comparative Study of Generative Models: Utility and Privacy Implications".

Conference papers

(C1) **A.-G. Andrei**, A. Radzhabov, I. Coman, V. Kovalev, B. Ionescu, H. Müller, "Overview of the ImageCLEFmedical GANs 2023 Task: Identifying Training Data "Fingerprints" in Synthetic Biomedical Images Generated by GANs for Medical Image Security", ImageCLEF 2023, CLEF Conference and Labs of the Evaluation Forum, ISSN: 1613-0073, September 18–21, 2023, Thessaloniki, Greece (indexed in DBLP Computer Science Bibliography);

(C2) **A.-G. Andrei**, B. Ionescu, "AIMultimediaLab at ImageCLEFmedical GANs 2023: Determining "Fingerprints" of Training Data in Generated Synthetic Images", ImageCLEF 2023, CLEF Conference and Labs of the Evaluation Forum, ISSN: 1613-0073, September 18–21, 2023, Thessaloniki, Greece (indexed in DBLP Computer Science Bibliography);

(C3) **A.-G. Andrei**, M.G. Constantin, M. Dogariu, B. Ionescu, "AI Multimedia Lab at ImageCLEFmedical GANs 2024: Deep Learning Approaches for Analyzing Synthetic Medical Images", ImageCLEF 2024, CLEF Conference and Labs of the Evaluation Forum, CEUR-WS, 3740, ISSN: 1613-0073, September 9–12, 2024, Grenoble, France (indexed in DBLP Computer Science Bibliography);

(C4) **A.-G. Andrei**, A. Radzhabov, D. Karpenka, Y. Prokopchuk, V. Kovalev, B. Ionescu, H. Müller, "Overview of 2024 ImageCLEFmedical GANs Task–investigating

Generative Models' Impact on Biomedical Synthetic Images", ImageCLEF 2024, CLEF Conference and Labs of the Evaluation Forum, CEUR-WS, 3740, ISSN: 1613-0073, September 9–12, 2024, Grenoble, France (indexed in DBLP Computer Science Bibliography);

(C5) **A.-G. Andrei**, M.G. Constantin, M. Dogariu, L. Ştefan, B. Ionescu, "AI Multimedia Lab at ImageCLEFmedical GANs 2025: Identifying Real-image Usage in Generated Medical Images", ImageCLEF 2025, CLEF Conference and Labs of the Evaluation Forum, CEUR-WS, 4038, ISSN: 1613-0073, September 9–12, 2025, Madrid, Spain (indexed in DBLP Computer Science Bibliography);

(C6) **A.-G. Andrei**, M.G. Constantin, M. Dogariu, A. Radzhabov, L.D. Ştefan, Y. Prokopchuk, V. Kovalev, H. Müller, B. Ionescu, "Overview of ImageCLEFmedical 2025 GANs Task: Training Data Analysis and Fingerprint Detection", ImageCLEF 2025, CLEF Conference and Labs of the Evaluation Forum, CEUR-WS, 4038, ISSN: 1613-0073, September 9–12, 2025, Madrid, Spain (indexed in DBLP Computer Science Bibliography);

(C7) B. Ionescu, H. Müller, A.M. Drăgulinescu, A. Idrissi-Yaghir, A. Radzhabov, A.G.S. de Herrera, **A. Andrei**, et al. "Advancing Multimedia Retrieval in Medical, Social Media and Content Recommendation Applications with ImageCLEF 2024", in Advances in Information Retrieval: 46th European Conference on Information Retrieval, ECIR 2024 Proceedings, March 24–28, Glasgow, UK, 2024 (indexed in Web of Science);

(C8) B. Ionescu, H. Müller, A.-M. Drăgulinescu, A. Popescu, A. Idrissi-Yaghir, A.G. Seco de Herrera, **A. Andrei**, et al. "ImageCLEF 2023 Highlight: Multimedia Retrieval In Medical, Social Media And Content Recommendation Applications", in Springer Lecture Notes in Computer Science, 13982, pp. 557–567, ISBN: 978-3-031-28240-9, DOI: https://doi.org/10.1007/978-3-031-28241-6_64, ECIR 2023 Proceedings, April 2 – 6, Dublin, Ireland, 2023 (indexed in Web of Science);

(C9) B. Ionescu, H. Müller, D.-C. Stanciu, A. Idrissi-Yaghir, A. Radzhabov, A.G. Seco de Herrera, **A. Andrei**, et al. "ImageCLEF 2025: Multimedia Retrieval in Medical, Social Media and Content Recommendation Applications", in Advances in Information Retrieval: 47th European Conference on Information Retrieval, ECIR 2025 Proceedings, April 6–10, Lucca, Italy, 2025 (indexed in Web of Science);

(C10) **A.-G. Andrei**, A.M. Tăuțan, B. Ionescu, "Freezing of Gait Detection in Parkinson's Disease: Comparison of Deep Learning Frameworks", 18th IEEE International Symposium on Medical Measurements and Applications - MeMeA 2024, 26–28

June, Eindhoven, Netherlands, 2024, DOI: 10.1109/MeMeA60663.2024.10596827, ISSN: 2837-5882 (indexed in Web of Science);

(C11) A.-M. Tăuțan,**A.-G. Andrei**, B. Ionescu, "Freezing of Gait Detection for Parkinson's Disease Patients using Accelerometer data: Case study", IEEE International Conference on E-Health and Bioengineering - EHB, WOS: 000646194100095, October 29-30, Iași, Romania, 2020, (indexed in Web of Science).

Book chapters:

(B1) B. Ionescu, H. Müller, ..., **A.-G. Andrei**, ..., A. Popescu, "Overview of the ImageCLEF 2023: Multimedia Retrieval in Medical, Social Media and Internet Applications", in Springer Lecture Notes in Computer Science, 14163, pp. 370–396, ISBN: 978-3-031-42447-2, DOI: https://doi.org/10.1007/978-3-031-42448-9_25, CLEF 2023 Proceedings, 18-21 September, Thessaloniki - Greece, 2023;

(B2) B. Ionescu, H. Müller, ..., **A.-G. Andrei**, Y. Prokopchuk, ..., B. Stein. "Overview of the ImageCLEF 2024: Multimedia Retrieval in Medical Applications", in Springer Lecture Notes in Computer Science, 14959, pp. 140–164, ISBN: 978-3-031-71907-3, DOI: https://doi.org/10.1007/978-3-031-71908-0_7, CLEF 2024 Proceedings, 9-12 September, Grenoble - France, 2024;

(B3) B. Ionescu, H. Müller, D.-C. Stanciu, **A.-G. Andrei**, ... B. Stein, Overview of ImageCLEF 2025: Multimedia Retrieval in Medical, Social Media and Content Recommendation Applications", in Springer Lecture Notes in Computer Science, 16089, pp. 290–314, ISBN: 978-3-032-04353-5, DOI: https://doi.org/10.1007/978-3-032-04354-2_17, CLEF 2025 Proceedings, 9-12 September, Madrid - Spain, 2025.

8.4 Perspectives for further developments

The contributions of this thesis open several directions for future research, spanning synthetic medical image generation, privacy-aware model design, clinical analysis, and the broader integration of AI into healthcare practice.

Within the domain of synthetic medical image generation, the findings demonstrate that privacy risk is strongly architecture-dependent and cannot be assumed from image quality metrics alone. A natural direction for future work is the incorporation of privacy-preserving mechanisms, such as differential privacy, adversarial regularization, or membership inference-aware training objectives directly into the generative model pipeline, rather than evaluating privacy post-hoc. This shift from privacy assessment to privacy by design represents an important step toward making synthetic medical data genuinely suitable for sharing and reuse in clinical contexts.

Evaluation frameworks for generative models also require further development. The results of this thesis show that visual realism, distributional fidelity, anatomical plausibility, and privacy resistance are complementary dimensions that can vary substantially across architectures. Future work should focus on establishing standardized, multi-faceted benchmarks that integrate task-driven utility metrics, anatomy-aware assessment, and resistance to data leakage attacks. Such benchmarks would enable more reliable comparisons across generative architectures and provide clearer guidance for practitioners selecting models for privacy-sensitive applications.

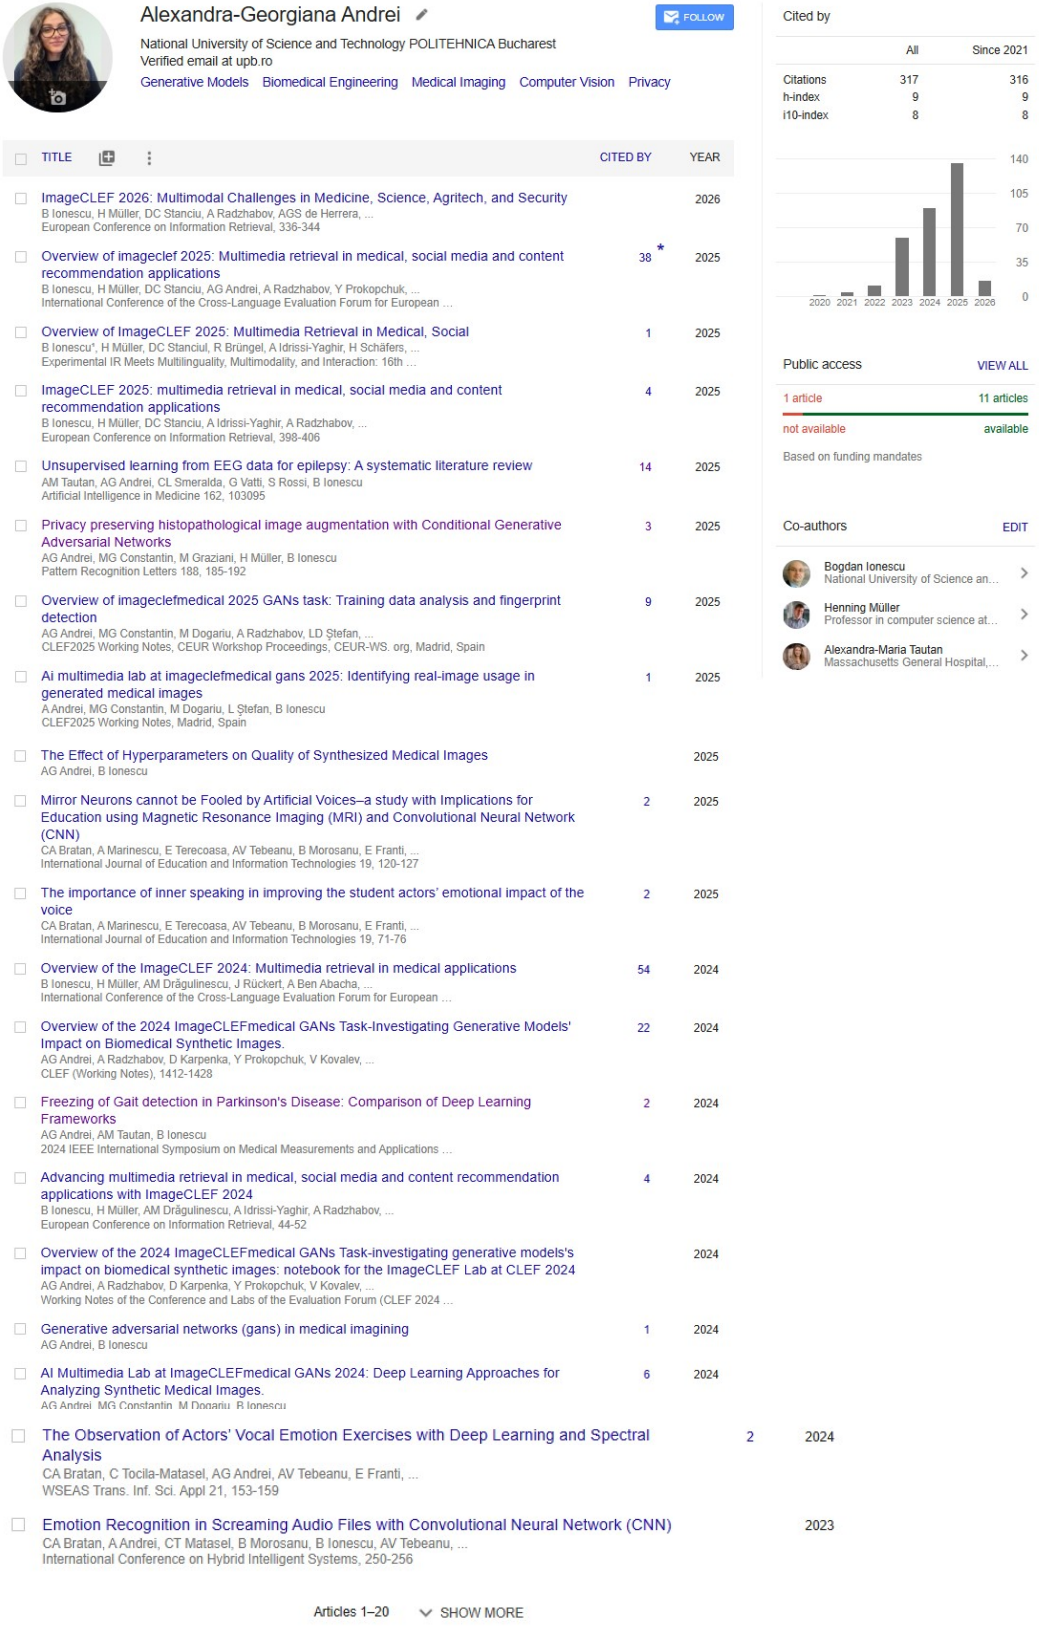
In this context, community-driven benchmarking initiatives such as the ImageCLEFmedical GANs challenge provide an important experimental platform for studying privacy and security aspects of synthetic medical images. The 2026 edition of this challenge introduces tasks focusing on detecting whether specific real images were used in the training of generative models, identifying the latent vectors corresponding to generated images, and generating privacy-preserving CT slices. These tasks explicitly investigate the presence of training-data fingerprints, latent-space memorization, and unintended reconstruction of patient-specific anatomy in synthetic images. Participating in and extending such evaluation frameworks represents a valuable direction for future work, as they allow systematic investigation of privacy leakage mechanisms and facilitate the development of generative models that balance realism with strong privacy guarantees.

The extension of the proposed frameworks to additional imaging modalities, including MRI, ultrasound, and fundus photography, and to more complex generation tasks such as volumetric 3D synthesis and multi-modal image translation represents another important direction. Similarly, the investigation of foundation models and large-scale generative architectures, such as latent diffusion models, in the context of medical image privacy remains largely unexplored and warrants dedicated study.

In the domain of neurological disorder analysis, several avenues for future work emerge from the contributions of Chapter 4. For EEG-based epilepsy research, the development of standardized evaluation benchmarks and the systematic comparison of unsupervised learning methods across datasets would help consolidate a currently fragmented literature. For Freezing of Gait detection, future work should address the generalization challenges identified in this thesis, including the small number of subjects available in existing datasets and the imbalance between FoG and normal gait episodes. Directions include leave-one-subject-out evaluation for more individualized analysis, the integration of convolutional and recurrent architectures within autoencoder frameworks to enhance temporal feature extraction, and the potential use of generative models to augment the scarce FoG episode data available for training.

More broadly, the integration of synthetic data and AI-driven analysis tools into real clinical and research settings requires closer alignment with medical expertise and regulatory frameworks. Future developments should investigate how synthetic

datasets can be validated, documented, and shared responsibly, and how regulatory requirements such as GDPR can be operationalized within generative model workflows. Interdisciplinary collaboration between AI researchers, clinicians, and regulatory bodies will be essential to ensure that the methods developed in this and related work translate into trustworthy, deployable tools for medical imaging research and clinical practice.



Alexandra-Georgiana Andrei

FOLLOW

National University of Science and Technology POLITEHNICA Bucharest
Verified email at upb.ro

Generative Models | Biomedical Engineering | Medical Imaging | Computer Vision | Privacy

TITLE	CITED BY	YEAR
ImageCLEF 2026: Multimodal Challenges in Medicine, Science, Agritech, and Security B Ionescu, H Müller, DC Stanciu, A Radzhabov, AGS de Herrera, ... European Conference on Information Retrieval, 336-344		2026
Overview of imageclef 2025: Multimedia retrieval in medical, social media and content recommendation applications B Ionescu, H Müller, DC Stanciu, AG Andrei, A Radzhabov, Y Prokopchuk, ... International Conference of the Cross-Language Evaluation Forum for European ...	38 *	2025
Overview of ImageCLEF 2025: Multimedia Retrieval in Medical, Social B Ionescu*, H Müller, DC Stanciul, R Brüngel, A Idrissi-Yaghir, H Schäfers, ... Experimental IR Meets Multilinguality, Multimodality, and Interaction: 16th ...	1	2025
ImageCLEF 2025: multimedia retrieval in medical, social media and content recommendation applications B Ionescu, H Müller, DC Stanciu, A Idrissi-Yaghir, A Radzhabov, ... European Conference on Information Retrieval, 398-406	4	2025
Unsupervised learning from EEG data for epilepsy: A systematic literature review AM Tautan, AG Andrei, CL Smeralda, G Vatti, S Rossi, B Ionescu Artificial Intelligence in Medicine 162, 103095	14	2025
Privacy preserving histopathological image augmentation with Conditional Generative Adversarial Networks AG Andrei, MG Constantin, M Graziani, H Müller, B Ionescu Pattern Recognition Letters 188, 185-192	3	2025
Overview of imageclefmedical 2025 GANs task: Training data analysis and fingerprint detection AG Andrei, MG Constantin, M Dogariu, A Radzhabov, LD Ştefan, ... CLEF2025 Working Notes, CEUR Workshop Proceedings, CEUR-WS. org, Madrid, Spain	9	2025
Ai multimedia lab at imageclefmedical gans 2025: Identifying real-image usage in generated medical images A Andrei, MG Constantin, M Dogariu, L Ştefan, B Ionescu CLEF2025 Working Notes, Madrid, Spain	1	2025
The Effect of Hyperparameters on Quality of Synthesized Medical Images AG Andrei, B Ionescu		2025
Mirror Neurons cannot be Fooled by Artificial Voices–a study with Implications for Education using Magnetic Resonance Imaging (MRI) and Convolutional Neural Network (CNN) CA Bratan, A Marinescu, E Terecoasa, AV Tebeanu, B Morosanu, E Franti, ... International Journal of Education and Information Technologies 19, 120-127	2	2025
The importance of inner speaking in improving the student actors' emotional impact of the voice CA Bratan, A Marinescu, E Terecoasa, AV Tebeanu, B Morosanu, E Franti, ... International Journal of Education and Information Technologies 19, 71-76	2	2025
Overview of the ImageCLEF 2024: Multimedia retrieval in medical applications B Ionescu, H Müller, AM Drăgulinescu, J Rückert, A Ben Abacha, ... International Conference of the Cross-Language Evaluation Forum for European ...	54	2024
Overview of the 2024 ImageCLEFmedical GANs Task-Investigating Generative Models' Impact on Biomedical Synthetic Images. AG Andrei, A Radzhabov, D Karpenka, Y Prokopchuk, V Kovalev, ... CLEF (Working Notes), 1412-1428	22	2024
Freezing of Gait detection in Parkinson's Disease: Comparison of Deep Learning Frameworks AG Andrei, AM Tautan, B Ionescu 2024 IEEE International Symposium on Medical Measurements and Applications ...	2	2024
Advancing multimedia retrieval in medical, social media and content recommendation applications with ImageCLEF 2024 B Ionescu, H Müller, AM Drăgulinescu, A Idrissi-Yaghir, A Radzhabov, ... European Conference on Information Retrieval, 44-52	4	2024
Overview of the 2024 ImageCLEFmedical GANs Task-investigating generative models's impact on biomedical synthetic images: notebook for the ImageCLEF Lab at CLEF 2024 AG Andrei, A Radzhabov, D Karpenka, Y Prokopchuk, V Kovalev, ... Working Notes of the Conference and Labs of the Evaluation Forum (CLEF 2024 ...		2024
Generative adversarial networks (gans) in medical imagining AG Andrei, B Ionescu	1	2024
AI Multimedia Lab at ImageCLEFmedical GANs 2024: Deep Learning Approaches for Analyzing Synthetic Medical Images. AG Andrei, MG Constantin, M Dogariu, B Ionescu	6	2024
The Observation of Actors' Vocal Emotion Exercises with Deep Learning and Spectral Analysis CA Bratan, C Tocila-Matasel, AG Andrei, AV Tebeanu, E Franti, ... WSEAS Trans. Inf. Sci. Appl 21, 153-159	2	2024
Emotion Recognition in Screaming Audio Files with Convolutional Neural Network (CNN) CA Bratan, A Andrei, CT Matasel, B Morosanu, B Ionescu, AV Tebeanu, ... International Conference on Hybrid Intelligent Systems, 250-256		2023

Articles 1–20 SHOW MORE

Cited by

	All	Since 2021
Citations	317	316
h-index	9	9
i10-index	8	8

Public access VIEW ALL

1 article not available | 11 articles available

Based on funding mandates

Co-authors EDIT

Bogdan Ionescu
National University of Science an...

Henning Müller
Professor in computer science at...

Alexandra-Maria Tautan
Massachusetts General Hospital,...

Fig. 8.1 Google Scholar list of published works with citation counts (screenshot taken on 3 of May 2026).

References

[1] A. C. Quiros, R. Murray-Smith, K. Y. (2019). Pathologygan: Learning deep representations of cancer tissue. *arXiv preprint arXiv:1907.02644*.

[2] Andrei, A.-G., Constantin, M. G., Dogariu, M., and Ionescu, B. (2024). Ai multimedia lab at imageclefmedical gans 2024: Deep learning approaches for analyzing synthetic medical images. In *CLEF2024 Working Notes*, CEUR Workshop Proceedings, Grenoble, France. CEUR-WS.org.

[3] Andrzejak, R. G., Lehnertz, K., Mormann, F., Rieke, C., David, P., and Elger, C. E. (2001). Indications of nonlinear deterministic and finite-dimensional structures in time series of brain electrical activity: Dependence on recording region and brain state. *Physical Review E*, 64(6):061907.

[4] Arjovsky, M., Chintala, S., and Bottou, L. (2017). Wasserstein generative adversarial networks. In *International conference on machine learning*, pages 214–223. PMLR.

[5] Bachlin, M., Plotnik, M., Roggen, D., Maidan, I., Hausdorff, J. M., Giladi, N., and Troster, G. (2009). Wearable assistant for parkinson's disease patients with the freezing of gait symptom. *IEEE Transactions on Information Technology in Biomedicine*, 14(2):436–446.

[6] Brock, A., Donahue, J., and Simonyan, K. (2018). Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*.

[7] Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative adversarial nets. *Advances in neural information processing systems*, 27.

[8] Gurina, T. S. and Simms, L. (2020). Histology, staining.

[9] Heremans, E., Nieuwboer, A., and Vercruysse, S. (2013). Freezing of gait in parkinson's disease: where are we now? *Current neurology and neuroscience reports*, 13(6):350.

[10] Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. (2017). Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30.

[11] Ho, J., Jain, A., and Abbeel, P. (2020). Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851.

[12] Ho, J. and Salimans, T. (2021). Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*.

[13] Isola, P., Zhu, J.-Y., Zhou, T., and Efros, A. A. (2017). Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134.

[14] Kaabachi, B., Despraz, J., Meurers, T., Otte, K., Halilovic, M., Kulynych, B., Prasser, F., and Raisaro, J. L. (2025). A scoping review of privacy and utility metrics in medical synthetic data. *NPJ digital medicine*, 8(1):60.

[15] Karras, T., Aila, T., Laine, S., and Lehtinen, J. (2017). Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*.

[16] Karras, T., Laine, S., and Aila, T. (2019). A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410.

[17] Kather, J. N., Halama, N., and Marx, A. (2018). 100,000 histological images of human colorectal cancer and healthy tissue. Zenodo.

[18] Liu, Q., Shakya, R., Jovanovic, J., Khalil, M., and de la Hoz-Ruiz, J. (2025). Ensuring privacy through synthetic data generation in education. *British Journal of Educational Technology*, 56(3):1053–1073.

[19] Min, G. and Oh, J. (2025). Can synthetic data protect privacy? *IEEE Access*.

[20] Mirza, M. and Osindero, S. (2014). Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*.

[21] Organization, W. H. (2023). Colorectal cancer. https://www.who.int/news-room/fact-sheets/detail/colorectal-cancer. Accessed: 10 November 2025.

[22] Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, S. E., et al. (2021). The prisma 2020 statement: an updated guideline for reporting systematic reviews. *bmj*, 372.

[23] Parkinson's Foundation (n.d.). Statistics. Accessed: 2025-03-09.

[24] Radford, A., Metz, L., and Chintala, S. (2015). Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*.

[25] Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., and Chen, X. (2016). Improved techniques for training gans. *Advances in neural information processing systems*, 29.

[26] Sanchez-Serrano, P., Rios, R., and Agudo, I. (2025). A decision framework for privacy-preserving synthetic data generation. *Computers and Electrical Engineering*, 126:110468.

[27] Shoeb, A. H. (2009). *Application of machine learning to epileptic seizure onset detection and treatment*. PhD thesis, Massachusetts Institute of Technology.

[28] Stadler, T., Oprisanu, B., and Troncoso, C. (2020). Synthetic data-a privacy mirage. *arXiv preprint arXiv:2011.07018*.

[29] World Health Organization (2019). Epilepsy: A public health imperative: Summary. https://apps.who.int/iris/handle/10665/325440. Accessed: Jan. 30, 2023.

[30] Zhu, J.-Y., Park, T., Isola, P., and Efros, A. A. (2017). Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232.