Alessandro Ceresi[1], Borja Arroyo Galende[1], Javier Guinea-Pérez[1], Patricia A Apellániz[1],
Gustavo Hernández-Peñaloza[1], and Federico Álvarez[1]

[1]Affiliation not available

October 22, 2025

Abstract

The integration of Artificial Intelligence (AI) in healthcare holds significant promise, yet remains constrained by data scarcity, privacy and ethical concerns, and the inherent complexity of medical tasks. This review explores the intersection of deep generative modeling and federated learning, two complementary paradigms that enable collaborative and privacy-preserving innovation. We survey recent trends in deep generative models applied to healthcare, with a particular focus on federated settings. The review is structured around model architectures, providing a taxonomy-driven analysis that emphasizes solutions tailored to the heterogeneity of healthcare data. Although this review initially intended to cover multimodal applications, the current literature is largely limited to single-modality models. Among deep generative approaches, Generative Adversarial Networks dominate the field, particularly in domains such as medical image synthesis and data augmentation. In contrast, alternative architectures, including variational autoencoders, diffusion models, and autoregressive models, remain barely explored in federated healthcare scenarios. Notably, the number of related publications has increased substantially in recent years, rising from 1 in 2020 to 16 in 2024, reflecting the growing interest and momentum in this area. This paper highlights key challenges, open research questions, and future directions for the development of trustworthy, distributed generative AI in healthcare.

Deep Generative Models Meet Federated Learning: A Healthcare-Centered Review

Alessandro Ceresi, Borja Arroyo Galende, Javier Guinea-Pérez, Patricia A. Apellániz, Gustavo Hernández-Peñaloza, Federico Álvarez

***Abstract*—The integration of Artificial Intelligence (AI) in healthcare holds significant promise, yet remains constrained by data scarcity, privacy and ethical concerns, and the inherent complexity of medical tasks. This review explores the intersection of deep generative modeling and federated learning, two complementary paradigms that enable collaborative and privacy-preserving innovation. We survey recent trends in deep generative models applied to healthcare, with a particular focus on federated settings. The review is structured around model architectures, providing a taxonomy-driven analysis that emphasizes solutions tailored to the heterogeneity of healthcare data. Although this review initially intended to cover multimodal applications, the current literature is largely limited to single-modality models. Among deep generative approaches, Generative Adversarial Networks dominate the field, particularly in domains such as medical image synthesis and data augmentation. In contrast, alternative architectures, including variational autoencoders, diffusion models, and autoregressive models, remain barely explored in federated healthcare scenarios. Notably, the number of related publications has increased substantially in recent years, rising from 1 in 2020 to 16 in 2024, reflecting the growing interest and momentum in this area. This paper highlights key challenges, open research questions, and future directions for the development of trustworthy, distributed generative AI in healthcare.**

***Index Terms*—deep learning, federated learning, generative AI, healthcare, privacy.**

I. INTRODUCTION

Artificial Intelligence (AI) is revolutionizing healthcare by enabling faster and more accurate diagnoses, supporting informed clinical decision-making, and enhancing the optimization of healthcare delivery [1], [2]. Its capacity to process and learn from large-scale, complex medical data (such as Electronic Health Records (EHRs) [3], medical images [4], and signals from wearable devices [5]) is driving innovations across multiple domains, including early disease detection, patient monitoring, and personalized treatment planning. By supporting clinical workflows, AI contributes to more personalized, efficient, and effective care.

Despite these advances, the development and deployment of AI in healthcare remain hampered by critical challenges. These include the scarcity of labeled data, data imbalance, sampling bias, and the fragmented nature of medical records across institutions, affecting the quality of data [2]. Moreover, strict privacy regulations and ethical constraints further limit access to data, especially in cross-silo setups [6]. Such limitations are particularly acute in scenarios involving rare diseases, underrepresented patient subgroups, and low-resource clinical environments, where data are not only scarce but also heterogeneous, non-standardized, and highly sensitive [7]. These issues undermine the generalizability of AI models, reduce reproducibility, and restrict their applicability in real-world clinical settings.

These challenges motivate the use of Deep Generative Models (DGMs), which aim to learn the underlying distribution of healthcare data and produce synthetic samples that preserve their structure and statistical properties [8]. Beyond alleviating data scarcity, DGMs are particularly well-suited to understand the complexity and heterogeneity of medical data, which often spans multiple modalities, variable formats, and intricate temporal dynamics. By capturing nonlinear dependencies and multimodal correlations, DGMs can generate realistic synthetic data that support data augmentation [9], fill in missing information [10], and enable model development when real data are limited or restricted [11]. In the context of EHRs, for example, DGMs have been used to synthesize patient trajectories that reflect clinically plausible event sequences while preserving privacy [12]. These capabilities make DGMs valuable not only for enhancing model robustness but also for facilitating research and system development in privacy-sensitive or resource-constrained environments.

However, while DGMs can help mitigate many data-related challenges, their training process still relies on access to real patient data, which raises critical concerns about privacy, data governance, and regulatory compliance (GDPR [13], EHDS [14], or HIPAA [15]). These concerns are particularly pressing in clinical contexts, where data cannot be freely shared across institutions or jurisdictions. In this context, Federated Learning (FL) has emerged as a powerful paradigm to enable collaborative model training across multiple entities without exchanging raw data [16]. Instead, models are trained locally on decentralized datasets, and only model updates —typically gradients or weights —are exchanged. By adopting a decentralized computational model, this approach boosts

---

Manuscript received August 14, 2025; revised August 16, ???.

This work was supported by Synthetic Data Generation Framework for Integrated Validation of Use Cases and AI Healthcare Applications (SYNTHIA) project from Horizon Europe under Grant 101172872. Views and opinions expressed are, however, those of the authors only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the granting authority can be held responsible for them.

A. C, B. A., J. G, P. A, F. A are with Information Processing and Telecommunications Center, Escuela Técnica Superior de Ingeniería de Telecomunicación, Universidad Politécnica de Madrid, Madrid 28030 ES (e-mail: { alessandro.ceresi, borja.arroyog, javier.guinea.perez, patricia.alonsod, federico.alvarez }@upm.es).

G. H. is with Escuela Técnica Superior de Ingeniería de Sistemas Informáticos, Universidad Politécnica de Madrid, 28040 Madrid, ES (e-mail: gustavo.hernandez.penaloza@upm.es)

data privacy, ensuring that no single entity has full access to the complete dataset. Furthermore, FL enhances model robustness by exposing the model to diverse data distributions across clients. Additionally, to mitigate the risk of information disclosure during training, FL can be extended with secure multi-party computation [17] or homomorphic encryption [18], which effectively limit the amount of information handled by the aggregator, even in the presence of semi-trusted or honest-but-curious participants.

The combination of DGMs and FL presents a compelling and increasingly relevant solution to the dual challenges of data scarcity and privacy in healthcare. While DGMs can alleviate the lack of accessible data by synthesizing realistic and downstream task-tailored samples, FL ensures that these models can be trained collaboratively across distributed institutions without exposing sensitive patient information. This integration is not merely complementary in function, but also in architecture: the parameters of DGMs, such as Generative Adversarial Networks (GANs) [19] or Variational Autoencoders (VAEs) [20], can be optimized in a federated fashion, preserving data locality throughout the learning process. Together, they enable the secure and distributed generation of synthetic medical data, protecting confidentiality at both the training and output levels. This synergy holds significant potential for unlocking collaborative research, building robust benchmarking datasets, and supporting the development of clinical AI in privacy-sensitive and data-constrained environments.

Inspired by these complementary strengths, this review explores their application to the healthcare domain. We focus specifically on deep neural architectures, which represent the current state of the art in terms of scalability and capacity to model high-dimensional, multimodal medical data. Traditional generative methods, while useful in constrained scenarios, lack the flexibility and expressive power required to capture the full complexity of clinical information or to generate coherent outputs across diverse data modalities [21], [22]. Our goal is to provide a comprehensive and structured survey of recent works that apply DGMs in federated healthcare settings, highlighting architectural choices, methodological innovations, and application areas. In doing so, we identify common design patterns, key limitations in the existing literature, and open research questions that must be addressed to advance the development of trustworthy, privacy-preserving generative AI in clinical practice. To facilitate a focused and comparative analysis, we adopt a data-centric perspective and organize the review according to the main data formats encountered in healthcare applications: tabular data, sequential data, and medical imaging.

The rest of this paper is structured as follows. In Section II, we provide the necessary background on DGMs and FL, introducing key concepts and architectures relevant to their application in healthcare. Section III discusses the main challenges associated with different healthcare data modalities, highlighting how issues such as heterogeneity, privacy, and scarcity shape the design of generative solutions. Section IV outlines the methods used to identify and select relevant studies for this review. Section V presents the results of our analysis, offering a structured overview of existing works along dimensions such as model type, data characteristics, and task formulation. Finally, Section VI concludes the paper by summarizing key insights and outlining open challenges and future research directions.

II. BACKGROUND

This section lays the conceptual and technical groundwork for understanding the use of DGMs in healthcare, particularly when trained under privacy-preserving, distributed conditions. We begin by introducing the main families of DGMs and their relevance to medical data. We then discuss how these models can be trained collaboratively through FL to overcome data access constraints in clinical environments. Finally, we present key dimensions and methodological challenges involved in evaluating the fidelity, privacy, and utility of generated data in healthcare applications. These foundations support the subsequent analysis of modality-specific challenges and recent developments in the field.

A. Foundations of Deep Generative Models

In data-driven modeling, DGMs aim to learn the underlying data-generating process of a given dataset [23]. Specifically, they estimate the probability distribution of the observed data [24]. Given some data X, it is assumed that there is an underlying distribution from which these data were obtained, $X \sim p(X)$. The DGM then learns an approximation $\hat{p}(X)$ directly from the data.

DGMs can also be interpreted as graphical or hierarchical probabilistic models, where dependencies are defined between observed and latent variables [25]. These structures not only formalize the model architecture, but also provide a principled framework for sampling synthetic data [26]. Depending on the availability of data, DGMs can balance different degrees of prior assumptions. Models with strong prior structures are typically used in low-data regimes to inject domain knowledge, while models with weaker priors rely more heavily on learning from large datasets.

A wide range of DGM architectures has been developed in recent years. Within the scope of this review, we focus on four main families: GANs [19], Autoregressive (AR) models, VAEs [20], and Diffusion Models [27]. These approaches differ in their training objectives, generative mechanisms, and representational flexibility. A schematic summary of their training processes is provided in Figure 1. While other architectures, such as normalizing flows [28] or energy-based models [29], are used in the literature for different types of healthcare applications, in federated learning, the resulting research is limited to these four architectures.

These models have shown significant potential in healthcare, particularly in addressing challenges such as data scarcity, heterogeneity, and privacy constraints [30], [31]. They enable the generation of realistic synthetic data that can be used for dataset anonymization, augmentation, domain translation, or imputation of missing values [32], [33]. By preserving statistical and structural properties of the original data while

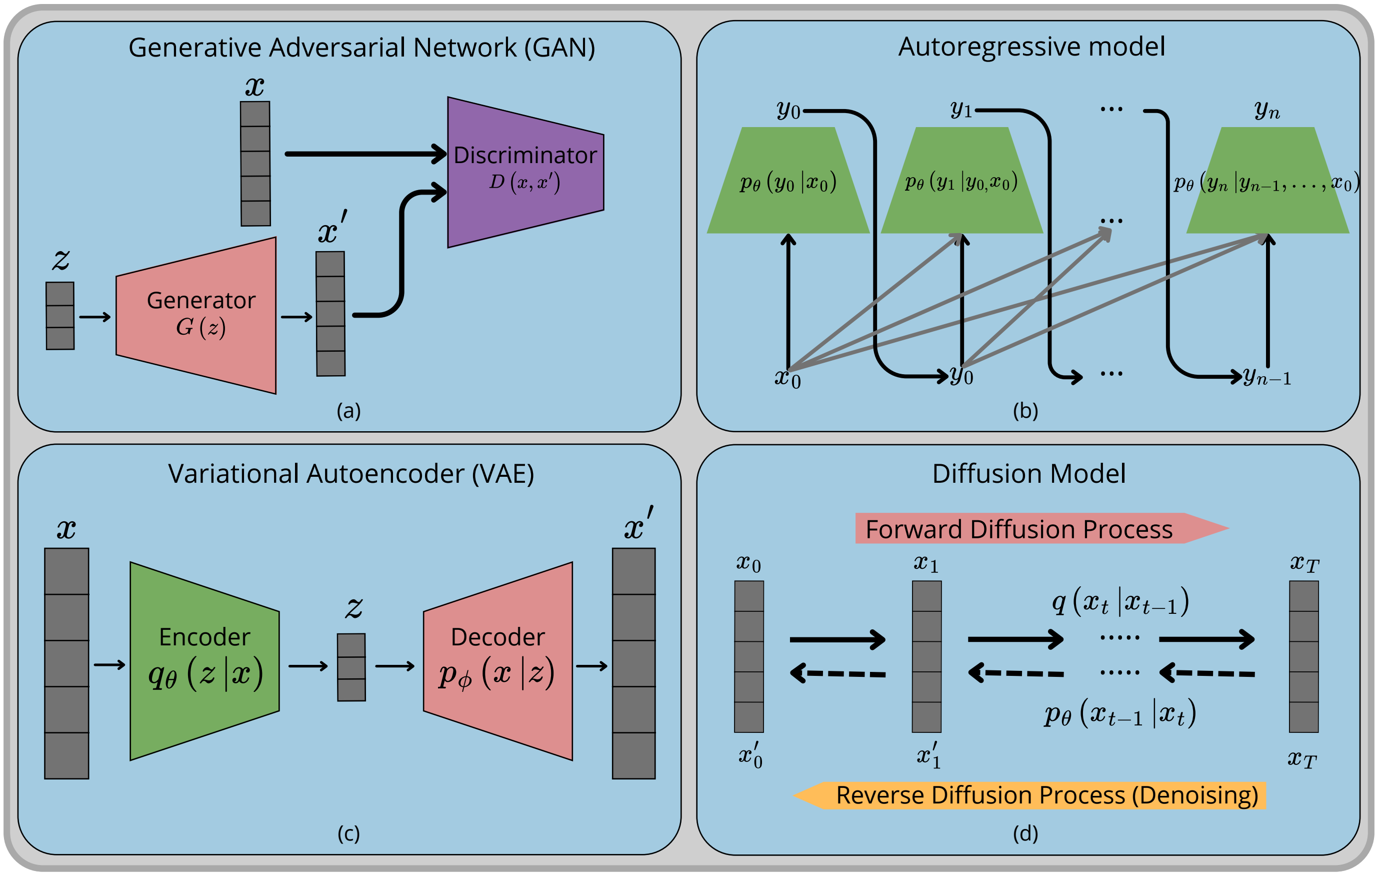

Fig. 1: Schematic representation of the core training processes of the four DGMs considered in this review: (a) **GANs**, where a generator produces synthetic samples that compete against a discriminator trained to distinguish real from fake data in an adversarial setup; (b) **AR models**, where data are generated sequentially, (i.e., each output y_t is conditioned on previous outputs and inputs), and the model learns $p_\theta(y_t|y_{t-1}, ..., x_0)$, with θ as the parameter set; (c) **VAEs**, where the input data are encoded into a latent space and reconstructed through a decoder, q_θ approximates the posterior over latent variables, and p_ϕ reconstructs the input from the latent code z (θ and ϕ are the trainable parameters of the encoder and decoder); (d) **Diffusion models**, generation involves a forward process that gradually adds noise to the data, and a reverse process learned to denoise it (q and p_θ denote the forward and reverse distributions, t is the timestep, and T the total number of steps).

protecting individual privacy [34], DGMs support the development of robust and generalizable AI models. They also facilitate clinical research in scenarios where real data are limited, restricted, or fragmented across institutions. Ultimately, DGMs can be tailored to simulate patient outcomes, support clinical decision-making, and personalize treatment planning [35], contributing to data-driven healthcare innovation while safeguarding patient confidentiality [36].

Generative Adversarial Networks (GANs): GANs consist of two neural networks, a generator and a discriminator, that are trained in a competitive framework. The generator creates synthetic data, while the discriminator attempts to distinguish between real and generated samples [19]. Through this adversarial process, the generator learns to produce increasingly realistic outputs without explicitly estimating the underlying data distribution. This ability to synthesize realistic samples has made GANs particularly attractive in medical imaging and structured data generation. They are often used for data augmentation, class balancing, or generating synthetic cohorts when real patient data are limited or sensitive. However, their performance is often challenged by mode collapse (where the generator produces a narrow set of outputs), failing to capture the full diversity of the data distribution. This issue is exacerbated in healthcare settings, where datasets are often small, heterogeneous, and imbalanced.

Auto-Regressive (AR) Models: AR models generate data sequentially, where each element is predicted based on previously generated elements and, optionally, on additional inputs. Formally, they model the joint distribution of a sequence as a product of conditional probabilities. AR models are commonly implemented using recurrent neural networks (RNNs) or, more recently, transformer-based architectures. Large language models (LLMs), such as the GPT family [37], are prominent examples of AR transformers. They have demonstrated remarkable capabilities in generating coherent, context-aware text, which has led to their adoption in tasks such as summarization, question answering, and dialogue generation. In healthcare, AR models are particularly relevant for modeling time-dependent data, such as patient trajectories, vital signs, or medication sequences, where preserving temporal consistency is essential [38]. Their ability to model sequential dependencies makes them well-suited for simulating longitudinal clinical data or forecasting future clinical events.

Variational Autoencoders (VAEs): VAEs are probabilistic generative models that learn a mapping from input data to a lower-dimensional latent space and back, enabling the generation of new samples from a learned distribution [20]. They combine variational inference with autoencoding by optimizing a tractable lower bound on the marginal likelihood of the data. The encoder maps the data into a distribution over latent variables, while the decoder reconstructs the data from samples drawn from that distribution. VAEs are particularly valuable in healthcare for handling noisy, sparse, or incomplete datasets, where their latent representations can be used for imputing missing values, simulating patient trajectories, or learning compressed embeddings for downstream tasks. Their principled approach to handling uncertainty makes them attractive for privacy-sensitive applications, and they are among the most widely explored generative architectures in biomedical research.

Diffusion Models: Diffusion models are a class of DGMs that learn data distributions by progressively corrupting input data with noise and then training a neural network to reverse this process [27]. This approach provides a stable and flexible training mechanism that has recently yielded state-of-the-art performance in image and audio generation. The forward process gradually adds noise over multiple timesteps, while the learned reverse process denoises the data to recover the original distribution. In healthcare, diffusion models are emerging as promising tools for generating high-fidelity synthetic data, particularly for complex modalities such as medical imaging and time series [39], [40]. They outperform traditional GANs in terms of stability and sample diversity, and their flexibility enables adaptation to a range of biomedical data types.

B. Federated Training of Deep Generative Models

Increasing concerns regarding the need for large, centralized datasets for training have accompanied the growing interest in generative modeling [41]. These concerns are especially critical in healthcare, where data are highly sensitive, regulated, and distributed across multiple institutions. Beyond the volume of data required, a key challenge lies in the secure, privacy-preserving use of this information across organizations that are unable (or unwilling) to share raw data due to legal and ethical constraints.

FL offers a compelling solution to this challenge by enabling collaborative training across decentralized nodes without transferring raw data [42]. In a typical FL setup, each data provider (e.g., a hospital or data silo) trains a local model using its private data. These local models are periodically merged by a central aggregator, which collects and combines model updates, such as parameters or gradients, using algorithms such as Federated Averaging (FedAvg) [43]. This paradigm preserves data locality while enabling the training of a shared global model.

This paradigm is particularly valuable in privacy-sensitive domains such as healthcare, where data often include confidential patient information [44]. FL provides a secure way to comply with strict data protection regulations, like GDPR 2016/679 [45] or HIPAA [46]. To further strengthen privacy guarantees, FL can be combined with techniques such as differential privacy (DP), which adds noise to sensitive data to obscure individual contributions [47], and secure multiparty computation, which allows encrypted collaboration between nodes [48].

However, applying FL to DGMs introduces unique technical challenges not present in discriminative tasks. While predictive models learn a conditional distribution $p(y|x)$, generative models aim to learn the joint distribution $p(x)$ (or related structured distributions), making them more sensitive to inconsistencies across nodes. In particular, FL for DGMs must contend with non-independent and identically distributed (non-IID) data, which is the norm in healthcare, as hospitals often treat different patient populations, use heterogeneous equipment, and follow varying clinical protocols. This heterogeneity can degrade both model convergence and the fidelity of generated samples [49]. Moreover, DGMs often require rich latent structures that generalize well across institutions. Maintaining consistency in the learned latent space under non-IID conditions is particularly difficult, and naive aggregation may lead to mode collapse, over-smoothing, or loss of diversity in generated outputs. These limitations are further amplified in federated settings by limited communication bandwidth, variable local computational resources, and the need for synchronization between nodes.

Despite these challenges, combining FL and DGMs has clear potential for healthcare: it enables privacy-preserving data synthesis, augmentation of underrepresented classes, generation of realistic rare cases, and simulation of patient trajectories, all without compromising patient confidentiality. This makes federated generative modeling a promising but still nascent research area in the development of trustworthy AI systems for clinical use [50].

C. Evaluation Dimensions for Deep Generative Models in Healthcare

Evaluating DGMs remains one of the most challenging and unsettled aspects of their development, particularly in sensitive domains such as healthcare [51]. Unlike discriminative models, for instance, which are typically assessed by their predictive performance on a test set (e.g., by estimating $p(Y|X)$ and computing classification or regression metrics), DGMs aim to learn the underlying joint distribution $p(X)$ and generate new samples from it. This difference makes direct evaluation more complex, especially when the output is not a single prediction but an entire synthetic dataset.

Currently, there is no universally accepted standard for evaluating DGMs [52], [53]. Instead, evaluation strategies are typically designed based on the intended use case, such as imputation, data augmentation, or anonymization, and must account for both technical accuracy and ethical implications [54]–[56]. In healthcare, where data are often scarce, heterogeneous, and sensitive, this evaluation becomes even more critical. To address this, most frameworks adopt a multi-dimensional evaluation approach based on three core aspects: fidelity, privacy, and utility [57], [58].

- **Fidelity** refers to the degree to which generated data replicate the statistical properties, structure, and complexity of the original data distribution. It ensures that the generated samples are realistic and coherent [59].
- **Privacy** assesses the risk of re-identification or information leakage from generated data, ensuring that these samples do not expose sensitive details about individuals in the original dataset [59]. In medical settings, this is crucial to ensure regulatory compliance and ethical integrity.
- **Utility** evaluates whether generated data can serve as a meaningful substitute for real data in downstream tasks, such as model training or exploratory analysis [60]. It reflects the practical value of these data in real-world clinical or research applications.

A balanced assessment across these three dimensions is essential for judging the overall quality and reliability of DGMs in healthcare. However, trade-offs often arise (for example, improving fidelity may compromise privacy [61]), making evaluation not only a technical challenge but also a matter of design choice aligned with specific clinical or regulatory priorities.

D. Privacy, Security, and Regulatory Considerations

Beyond technical design, deploying DGMs in healthcare also requires careful attention to privacy, security, and regulatory compliance. As these models increasingly approach clinical-grade applications, particularly in areas such as medical imaging or clinical decision support, understanding the legal and ethical landscape becomes essential. This subsection outlines the regulatory context and associated security challenges, along with key mitigation strategies.

Recent policy initiatives, most notably the FDA's executive summary *"Total Product Lifecycle Considerations (TPLC) for Generative AI-Enabled Devices"* [62] and the European Union's AI Act [63], underscore the necessity of holistic governance for DGMs under the domain of medical devices. These frameworks emphasize lifecycle-oriented governance, calling for rigorous design validation, quality assurance, and continuous post-deployment monitoring. In particular, the FDA highlights the unique challenges posed by generative AI in clinical applications, including variability of outputs and model drift, and emphasizes the need for robust oversight mechanisms.

Key regulatory challenges identified include managing:

- **Variable Outputs:** Complexity in validation processes due to variability in generative outputs.
- **Rapid Iteration:** Continuous oversight and evaluation are necessary due to the rapid evolution of the model.
- **Data Quality and Bias:** Ensuring unbiased, high-quality data sets throughout training and operational phases.

Privacy-enhancing technologies such as FL and DP are specifically endorsed by policy frameworks, including the EU AI Act, as proactive compliance mechanisms to mitigate privacy and security risks. FL enables decentralized collaborative training without direct data sharing, aligning with data minimization principles mandated by regulations like GDPR and HIPAA.

However, these same techniques introduce new security vulnerabilities that must be proactively managed:

- **Privacy Leakage:** DGMs risk inadvertently memorizing sensitive training data, making them susceptible to membership inference attacks [61], [64]. DP, although recommended, poses a privacy-utility trade-off [65].
- **Poisoning and Inversion in FL:** Malicious clients may introduce compromised model updates (poisoning) or reconstruct sensitive data from aggregated gradients (inversion attacks) [66], [67].
- **Model Stealing and Intellectual Property Risks:** DGMs represent valuable intellectual property, vulnerable to theft through unauthorized querying and replication. Current countermeasures, such as watermarking, are promising but still evolving [68].

Effectively mitigating these security risks while leveraging FL and privacy-preserving techniques ensures compliance with comprehensive lifecycle-based regulatory frameworks. Robust defense mechanisms, continuous threat monitoring, and adherence to international standards (such as ISO/IEC 27001, ISO/IEC 42001, and the NIST AI Risk Management Framework) become essential for successful regulatory alignment.

III. CHALLENGES ACROSS HEALTHCARE DATA MODALITIES

Traditionally, data have been treated as homogeneous entities. However, in practice, healthcare information exists in various modalities, each characterized by distinct structural and statistical properties. In this context, a data modality refers not only to the format in which data are represented, but also to the internal coherence and dependencies that define its semantics. For example, images encode spatially correlated information [69], while sequential data, such as clinical notes or physiological signals, unfold over time and require temporal modeling [70].

The goal of this section is to examine how these modality-specific characteristics impact the development and deployment of DGMs in healthcare. We structure the discussion around four major data categories: tabular, sequential, imaging, and multimodal. For each, we outline its clinical relevance, describe the structural and operational constraints that complicate generative modeling, and highlight the unique challenges that arise when such data are handled in privacy-sensitive, decentralized environments. This analysis provides the foundation for understanding both the opportunities and limitations of current approaches and for identifying modality-specific gaps in the literature.

A. Image data

Medical imaging plays a crucial role in clinical workflows, enabling the visualization of internal anatomical structures for diagnosis, monitoring, and treatment planning. Modalities such as Computed Tomography (CT), Magnetic Resonance Imaging (MRI), Positron Emission Tomography (PET), and digital pathology produce high-resolution data essential for understanding complex physiological conditions. In this context, the generation of high-quality synthetic medical images has emerged as a key area of research, offering a solution to data scarcity, class imbalance, and privacy concerns in model training [71]. However, utilizing medical image data, particularly in federated generative frameworks, introduces a variety of significant challenges that need to be addressed:

- **High Dimensionality and Storage Requirements**: Medical images are inherently high-dimensional and large, requiring substantial memory and compute resources. This poses a challenge for local training and communication in federated frameworks, particularly when working with 3D volumes or whole-slide images.
- **Cross-Institutional Variability**: Imaging data vary significantly across scanners, acquisition protocols, and institutions. Differences in resolution, contrast, noise profiles, and anatomical framing make it difficult to develop DGMs that generalize across domains and remain clinically reliable.
- **Semantic and Anatomical Consistency**: DGMs must preserve anatomical coherence and fine-grained structures, particularly when generating pathology or organ-level details. Ensuring semantic consistency across samples is critical for clinical applicability, but remains a significant modeling challenge.
- **Limited and Imbalanced Annotations**: Due to privacy restrictions, high annotation costs, and the low prevalence of certain pathologies, many institutions have access to only small or biased image datasets. This limits the diversity and representativeness of training data in federated scenarios [72].
- **Privacy and Security Risks**: Gradients derived from high-resolution images can leak sensitive information in federated setups. This raises concerns about membership inference and reconstruction attacks if communication is not adequately protected.
- **Local Computational Load**: Training large-scale DGMs, such as diffusion models or GANs for imaging, demands significant GPU resources. In federated environments, this can exceed the capabilities of some clinical nodes or require special scheduling and orchestration mechanisms.

B. Tabular data

A significant portion of the studies considered in this review relies on tabular data, which remains the dominant format for representing structured clinical information. This includes patient demographics, laboratory results, vital signs, diagnostic codes, treatment history, and genomic variants. Tabular data power a wide range of tasks, from risk prediction to clinical decision support and health economics modeling. Generating synthetic tabular data that preserve the complex statistical and clinical relationships present in real datasets is crucial for enabling privacy-preserving research and robust model development [73]–[75].

Despite their structured format, tabular healthcare data present unique challenges for generative modeling, especially in federated settings:

- **High Dimensionality and Sparsity**: Clinical datasets often include hundreds or thousands of features per patient, many of which are sparsely populated due to irregular data collection, missing values, or low-prevalence conditions [76]. This sparsity complicates the estimation of realistic joint distributions, especially in low-data or non-IID settings.
- **Mixed Data Types and Complex Dependencies**: Tabular data include numerical, ordinal, categorical, and binary variables, often with multimodal or skewed distributions. Capturing these heterogeneous types in a coherent generative framework requires models that can handle discrete-continuous mixtures while preserving inter-variable dependencies and logical constraints (e.g., medical codes compatible with age and sex).
- **Mode Collapse and Interpretability**: Tabular DGMs often struggle with full mode coverage, particularly in underrepresented subgroups or rare feature combinations. Moreover, evaluating the plausibility and interpretability of generated records remains challenging, especially when clinical knowledge is required to validate coherence.
- **Data Heterogeneity Across Institutions**: In federated environments, differences in EHR systems, data codification, variable naming, and units across institutions create difficulties for aligning features and ensuring consistent modeling. Without proper harmonization, models may learn site-specific patterns, leading to biased or non-generalizable outputs.
- **Privacy Risks in Small Cohorts**: Generating tabular data from rare disease cohorts or specialized subpopulations can increase the risk of re-identification, especially if DGMs overfit or memorize specific records. Federated frameworks help mitigate this but require additional safeguards to ensure privacy.

C. Sequential Data

Several healthcare data modalities are inherently sequential, capturing processes that evolve. These include biomedical time-series (e.g., electrocardiogram (ECG), electroencephalogram (EEG), or accelerometer signals), temporally indexed clinical markers (e.g., red blood cell count, medication history), unstructured text sequences (clinical narratives or event logs), and biological sequences (e.g., DNA or RNA). These data types are characterized by dependencies across time or position, variable sequence lengths, and often exhibit irregular sampling or censoring due to practical constraints in data collection.

Generative modeling of sequential data, especially in federated settings, presents a range of technical challenges:

- **Irregular Sampling and Temporal Correlation**: Many clinical sequences, such as lab measurements or symptom progressions, are sampled at non-uniform intervals. This irregularity breaks the assumptions of standard AR or RNN-based architectures, requiring specialized models that can handle time gaps or variable-frequency updates.

- **Censoring and Missing Events**: In longitudinal medical data, right-censoring and incomplete records are common, particularly in survival analysis and EHR logs. DGMs must learn to simulate realistic patient trajectories while accounting for partial observability and uncertainty in event timing.
- **Variable-Length and Structural Diversity**: Sequences differ substantially in length and structure (e.g., genomic data are discrete and positionally ordered, whereas biosignals are continuous and high-resolution). Designing models that generalize across this diversity is a core challenge.
- **Variable-Length and Heterogeneous Sequences**: Sequences differ in duration, resolution, and modality. For instance, biosignals are high-frequency and continuous, while genomic sequences are discrete and symbolic. Handling this diversity within a unified generative framework is challenging.
- **Realistic Trajectory Generation**: Synthetic sequences must maintain temporal continuity and biological plausibility. For example, generated ECG waveforms should reflect valid morphologies; simulated treatment sequences should follow realistic clinical progressions.
- **Temporal Drift and Non-IID Distributions**: Data collected from different hospitals or periods may reflect shifts in clinical practices or patient populations. Such drift can degrade generalization and increase training instability in federated setups.
- **Communication Costs and Compression Trade-offs**: High-frequency or long sequences produce large model updates. While latent compression (e.g., via VAEs) can reduce bandwidth, it may also discard clinically relevant temporal features [77].

D. Multimodal Integration

In real-world healthcare scenarios, information about a patient is rarely captured through a single type of data. Instead, multiple sources (each with different structures, semantics, and temporal resolutions) are often combined to build a more complete representation of clinical states. This is the case, for instance, with EHRs, which integrate structured tabular data (e.g., demographics, lab results), unstructured text (e.g., clinical notes), time-series signals (e.g., heart rate), and medical images [78], [79]. These heterogeneous, longitudinal records present a rich yet challenging landscape for AI models, particularly for DGMs.

Each modality presents its specific challenges (refer to previous sections), requiring specialized architectures tailored to their unique statistical properties [80]. Beyond the complexity within each modality, combining them introduces a new layer of difficulty: multimodal models must learn both intra-modality and cross-modality relationships, increasing architectural complexity and optimization difficulty [81]–[83]. In healthcare, this fusion holds the promise of revealing richer clinical patterns and improving prediction performance [84]. For instance, combining genomic features with imaging or biosignal data can uncover causal relationships that may inform disease progression or treatment response [85].

Multimodal generative modeling aims to synthesize or complete datasets that span two or more modalities. This can involve tasks such as domain translation (e.g., generating pathology images from genomic profiles), joint co-generation (e.g., producing aligned images and reports), or unified synthesis from latent representations. In healthcare, these models are particularly appealing because they enable the simulation of missing modalities, augmentation of rare multimodal cases, and causal exploration across biological systems.

However, developing multimodal DGMs, especially under federated constraints, introduces a series of challenges:

- **Modality Imbalance and Partial Availability**: Not all institutions collect the same modalities. Some patients may lack imaging, while others have no molecular data. This imbalance affects both training and downstream utility, particularly when client data are incomplete in federated setups.
- **Asynchrony and Resolution Mismatch**: Modalities may differ in spatial, temporal, or semantic resolution, making it difficult to align inputs for generation. Examples include combining high-resolution imaging with sparse clinical codes or continuous vital signs.
- **Joint Representation and Semantic Coherence**: Multimodal DGMs require shared latent spaces that capture intra- and inter-modality dependencies. Ensuring that generated modalities remain semantically consistent (e.g., an image matching its synthetic clinical report) is a critical and unresolved issue [83], [86].
- **Conditional and Flexible Generation**: Models must be capable of synthesizing one modality conditioned on others, even when combinations vary across patients or clients. This is particularly important in FL, where some sites may only support a subset of modalities [81].
- **Heterogeneous Client Capabilities**: Institutions differ in computational resources and data types. Training large-scale multimodal models locally may be infeasible at some nodes, requiring coordinated orchestration or fallback strategies.
- **Cross-Modal Privacy Risks**: While individual modalities may appear anonymized, their combination can amplify re-identification risks. Multimodal DGMs must address this compounded privacy concern in both centralized and federated deployments.

In summary, multimodal data offer the richest view of patient health but also represent one of the most complex modeling challenges for generative AI. Their inherent heterogeneity, incomplete structure, and high dimensionality demand architectures that are both flexible and privacy-aware, especially when used in distributed or cross-institutional settings such as FL.

IV. SEARCH METHODOLOGY

This section describes the methodology employed to identify, select, and analyze the scientific literature reviewed in this work. The process included the systematic collection of publications, keyword-based querying, and multi-stage filtering to ensure relevance to the use of DGM techniques in federated healthcare contexts.

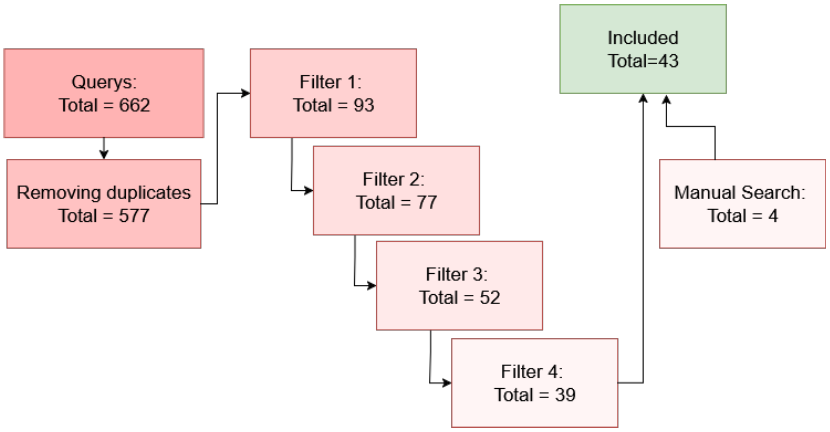

Fig. 2: General overview of the querying process.

TABLE I: Summary of the query design according to the keywords

Engine	Results	Query	Comments
WOS	353	TS=((Federated OR Federative) AND (Generativ* OR Synthesiz* OR Variational OR Adversarial OR GPT* OR GAN OR Generation OR VAE OR Autoregressive OR Diffusion) AND (Health* OR Medic*))	Also limited the publication year and document type to remove reviews.
IEEE	304	((((((”Abstract”:Generat*) OR (”Abstract”:Synthesiz*) OR (”Abstract”:Variational) OR (”Abstract”:Adversarial) OR (”Abstract”:Generation) OR (”Abstract”:GPT*) OR (”Abstract”:GAN) OR (”Abstract”:Autoregressive) OR (”Abstract”:VAE) OR (”Abstract”:Diffusion)) AND ((”Abstract”:Health*) OR (”Abstract”:Medic*)) AND (”Abstract”:Federat*)))))	Also limited publication year and document type to journal entry.
ACM	2	[Abstract:federat*] AND ([Abstract:health*] OR [Abstract:medic*]) AND [Keywords:generat*] AND [E-Publication Date:(01/01/2019 TO 12/31/2024)]	Did not include particular architectures as it added too much noise.
Scopus	7	(((”Index Terms”:generat*) AND ((”Abstract”:health*) OR (”Abstract”:medic*)) AND (”Abstract”:federat*))) AND PUBYEAR>2019 AND PUBYEAR<2025	Did not include particular architectures as it added too much noise.

An exhaustive investigation was conducted on several relevant search engines, including Web of Science[1], IEEExplore[2], Scopus[3], and ACM Library[4]. Queries were built using carefully selected keywords to ensure the exploration of all existing literature, allowing the identification of pertinent works for the work presented.

In addition to the previous indexed sources, we also manually included a small number of relevant *preprints* hosted on arXiv. Since arXiv[5] is not covered by the above digital libraries' indexing, these papers were added either because the authors were already aware of them or because they surfaced through supplementary, ad-hoc keyword searches on arXiv directly.

The overall process was carried out in a three-phase approach as presented in Figure 2. The words cover the main topics, and the search was aligned with a set of filters defined to aid in the identification of relevant works. The procedure is described in the following subsections.

A. Initial Phase: General Searching

Initially, general terms related to the use of DGMs were considered to define the scope of the study. Consequently, the terms generative, generation, etc., were included—also synthesize-rooted terms. To broaden the search and include papers in which a particular architecture was presented but none of the generation or synthesis terms appear, the most common architectures were included in the search. To limit the scope of the search to the healthcare domain, either medical or healthcare-related terms had to occur in the papers. Another search requisite was the appearance of terms related to FL. Thus, papers meeting all three search criteria: FL, data generation, and the healthcare domain, were listed. To limit our search to the state of the art, only papers that were published between 2020 and 2024 were included.

To unify the search criteria, a total of 4 similar queries were designed for the selected digital libraries. Each query was formulated to intersect three dimensions: research focus, specific healthcare applications, and methodology for AI-based data generation. Table I provides a summary of the query results, outlining the scope and outcomes of the search. The queries were formulated to be as similar as possible, given the

---

[1] Web of Science: https://www.webofscience.com. Accessed July 7th, 2025.
[2] IEEExplore: https://ieeexplore.ieee.org/. Accessed July 7th, 2025.
[3] Scopus: https://www.scopus.com/. Accessed July 7th, 2025.
[4] ACM Digital Library: https://dl.acm.org. Accessed July 7th, 2025.
[5] ArXiv: https://arxiv.org/. Accessed July 7th, 2025.

limitations of the query formulation tools available.

Using this approach, the study ensured a consistent and systematic identification of relevant literature on the topic. This strategy improved the thoroughness of the review and the precision in identifying studies related to generative AI. The resulting list of articles provided a solid foundation for subsequent analysis and synthesis, allowing a thorough understanding of current initiatives and their scope.

B. Second Phase: Research Filtering

Once the entire list was obtained, the process focused on identifying and discarding unrelated, out-of-scope, and duplicate papers. The objective was to refine the list and narrow it down to the most relevant papers for appraisal. Some unrelated papers complied with the search criteria, and other related papers were either not employing FL or not generating data per se. To discard these papers, the following criteria were applied:

- **Removal of Duplicates**: This filter removed all duplicate articles that appeared as some of the query engines index the same online libraries.
- **Relevance (Filter 1)**: This filter included an examination of the title and abstract of the article. Papers that met the search criteria but were deemed irrelevant to the topic were subsequently excluded from the review. For example, articles discussing next-generation mobile networks or articles where healthcare or FL were only mentioned in the motivation or as an example.
- **Healthcare Data (Filter 2)**: At this stage, a general inspection of each paper was conducted, and articles were excluded if their experimental setup did not include healthcare data.
- **Federation (Filter 3)**: In this step, papers were filtered based on the federation strategy they followed. Papers in which no federation was performed were excluded from the review.
- **DGMs (Filter 4)**: Lastly, papers in which generation was stated but no generative model or task was present were removed. This included papers that used non-generative models, such as autoencoders. Also excluded were papers with tasks such as entity extraction, summarization, and information retrieval.

The following section presents the results of the overall filtering process.

V. RESULTS AND DISCUSSION

The results of this review are organized based on the architecture of the DGMs proposed in the literature, reflecting the highly specialized nature of these models within FL environments for healthcare applications. This approach is motivated by the observation that the applications of DGMs in healthcare are remarkably diverse, spanning tasks such as synthetic data generation, data augmentation, dataset balancing, domain translation, and privacy-preserving data sharing. Each use case often demands tailored solutions, with the architecture of the DGM being central to addressing the specific requirements and challenges of the task. Given this strong dependence on architectural design, categorizing the reviewed studies by their DGM type, such as GANs, VAEs, diffusion models, or AR models, provides a clearer framework for understanding their methodologies, innovations, and limitations. This structure also facilitates comparisons across studies employing similar architectures, highlighting trends and gaps in the field, which will be discussed in the following sections. All the articles selected within the scope of this review are listed in Table II. In addition, an exhaustive description of the data sources can be found in Appendix A.

A. Generative Adversarial Networks (GANs)

GAN-based models remain one of the most frequently adopted generative approaches in healthcare, particularly in image synthesis. In the federated context, their application spans multiple modalities, including medical imaging, tabular data, and longitudinal sequences. This section reviews how recent works adapt the GAN framework to decentralized environments, the strategies proposed to mitigate known limitations such as mode collapse and training instability, and the effectiveness of these models across diverse tasks and clinical scenarios.

Owing to their malleability and ability to produce high-quality, sharp outputs, GANs have been leveraged across a broad range of healthcare data types. These include imaging, sequential, and structured tabular data, highlighting their versatility. Nevertheless, their primary strength lies in image generation, where their impact has been most pronounced. Accordingly, we begin by examining their use in medical imaging, organizing the works according to downstream tasks (augmentation, balancing, domain translation, and synthesis). In contrast, for tabular and sequential data, where the number of studies is more limited and the same architecture often supports multiple applications, the analysis is structured by data type rather than task.

Image Data

Data Augmentation: Many of the analyzed studies tackle the limited availability of image data by employing data augmentation techniques. In particular, GAN-based augmentation has become a cornerstone in federated healthcare settings, helping to address both data scarcity and heterogeneity in labeled medical datasets. In histopathology imaging for skin lesion classification, for instance, several works [90], [96], [96], [122] leverage DGMs to synthesize realistic lesion samples, enriching local datasets without breaching data privacy. The architectures presented are commonly trained on the ISIC skin lesion dataset and integrated with federated classifiers. Among these, the approach from the work of Cai et al. [90] employs a many-objective optimization-based GAN guided by evolutionary algorithms, optimizing synthetic image diversity, sharpness, and FID while maintaining FL integrity through FedAvg and gradient pruning. A similar strategy from Lan et al. [96], co-optimizes both the GAN and the federated CNN training objectives—global accuracy, AUC, training loss, and communication cost, by evolving multiple full pipeline configurations using population-based genetic evolution. In contrast, Selvaraj et al. [122] propose a co-training approach

TABLE II: Overview of federated applications of DGMs in Healthcare.
The generative tasks categories were defined as follow: *Synthesis* as default category, *Augmentation* when synthetic data are used to enrich a scarce training dataset, *Balancing* when synthetic data are generated to balance underrepresented classes of a dataset, *DomainTranslation* when synthetic data are generated in a domain different from that of the real data, *Imputation* when the generation process is employed to fill in missing or incomplete parts of a real dataset. Also, data type is encoded in the following categories: I=Imaging, T=Tabular, S=Sequential, M=Multimodal

Paper	**Architecture**	**Non-IID**	**Disease**	**Downstr. Task**	**Gen. Task**	**Data Type**	**Data Avail.**	**Code**	**Year**
Han et al. [87]	GAN		Pneumonia	Classification	Augmentation	I	✓	✓	2020
Kim et al. [88]	GAN		Parkinson	Segmentation	Synthesis	I	✓		2021
Zhang et al. [89]	GAN	✓	COVID-19	Classification	Balancing	I	✓		2021
Cai et al. [90]	GAN	✓	Skin cancer	Other	Augmentation	I	✓		2021
Weldon et al. [91]	GAN		Other	Generation	Synthesis	T	On request		2021
Brophy et al. [92]	GAN+AR		Other	Other	DomainTranslation	S	✓	✓	2021
Rajotte et al. [93]	GAN	✓	Skin cancer	Classification	Synthesis	I	✓		2021
Nguyen et al. [94]	GAN		COVID-19	Classification	Augmentation	I	✓	✓	2021
Yan et al. [95]	GAN	✓	Prostate cancer	Other	DomainTranslation	I	Partially		2021
Lan et al. [96]	GAN		Skin cancer	Other	Augmentation	I	✓		2022
Georgiadis et al. [97]	GAN	✓	COVID-19	Other	DomainTranslation	I	✓		2022
Fang et al. [98]	GAN	✓	Critical care ICU Diabetes Breast cancer Heart failure	Classification	Synthesis	T	Partially requiring certificate		2022
Jin et al. [99]	GAN		Skin cancer	Other	Synthesis	I	✓	✓	2022
Pennisi et al. [100]	GAN	✓	Tuberculosis	Classification	Synthesis	I	✓		2022
Elmas et al. [101]	GAN		Healthy	Generation	Synthesis	I	Partially	✓	2022
Shen et al. [102]	GAN	✓	Skin cancer	Classification	DomainTranslation	I	✓	✓	2022
Abdel-Basset et al. [103]	GAN	✓	COVID-19	Segmentation	DomainTranslation	I	✓		2022
Ul Alam et al. [104]	GAN		Sepsis	Classification	Imputation	T, S	With certificate		2023
Zhao et al. [105]	VAE	✓	Alzheimer	Classification	Augmentation	I	✓		2023
Lee et al. [106]	GAN	✓	COVID-19	Classification	Balancing	I	Partially	✓	2023
Yang et al. [77]	VAE	✓	Healthy	Classification	Augmentation	S	✓		2023
Peng et al. [107]	GAN		ASD Alzheimer	Classification	Augmentation	I	✓		2023
Wang et al. [108]	GAN		Healthy	Domain translation	DomainTranslation	I	✓	✓	2023
Peketi et al. [109]	GAN		Brain cancer	Segmentation	Augmentation	I	✓		2023
Tun et al. [110]	Diff		COVID-19	Generation	Synthesis	I	✓		2023
Jin et al. [111]	GAN		Skin cancer	Classification	Augmentation	I	✓		2023
Shaikh et al. [112]	GAN		Pulmonary diseases	Other	Augmentation	S	Partially		2023
Sattarov et al. [113]	Diff	✓	Diabetes	Classification	Synthesis	T	✓	✓	2024
Kundu et al. [114]	GAN	✓	Monkeypox	Classification	DomainTranslation	I	✓		2024
Rehman et al. [115]	GAN		Prostate cancer Lung cancer Breast cancer	Classification	Augmentation	T	✓	✓	2024
Chen et al. [116]	AR		Other	Other	Synthesis	I, S	✓		2024
Mazher et al. [117]	GAN	✓	Cardiac anomalies	Segmentation	DomainTranslation	I	✓		2024
Li et al. [118]	Diff		Healthy	Classification	Synthesis	I	✓		2024
Puppala et al. [119]	AR		Other	Chatbot	Synthesis	S		✓	2024
Sikandar et al. [120]	VAE		Stomach cancer	Classification	Balancing	S	✓		2024
Murmu et al. [121]	GAN	✓	Colon cancer Skin cancer	Classification	DomainTranslation	I	✓		2024
Selvaraj et al. [122]	GAN	✓	Skin cancer	Other	Augmentation	I	✓		2024
Chen et al. [123]	GAN		Skin cancer	Classification	Imputation	T, I	✓		2024
Alalwan et al. [124]	GAN		Brain cancer	Classification	Augmentation	I	✓		2024
Yan et al. [125]	AR	✓	Chest pathologies	Other	DomainTranslation	M	✓	✓	2024
Li et al. [126]	GAN+Diff	✓	Lung cancer	Other	DomainTranslation	I			2024
Nimeshika et al. [127]	GAN	✓	Alzheimer	Classification	Balancing	I	✓	✓	2024
Tariq Bdair et al. [128]	GAN		Brain cancer	I-t-I Translation	DomainTranslation	I		✓	2024

that uses local diagnostic models as discriminators in GAN training, generating images that are exchanged between clients to personalize and diversify local training sets without explicit parameter sharing. While the three above-presented methods are characterized by improved cancer diagnostic performance, the last one by Jin et al. [111] questions the robustness of GAN architectures against backdoor attacks. The study proposes a robust FedGAN framework that detects and mitigates backdoor poisoning in collaborative medical image synthesis. Beyond dermatology, augmentation techniques extend to other imaging modalities and diagnostic contexts. While implementations appear simplistic, e.g., training GANs offline for augmentation prior to federated CNN training in brain MRI classification, the study of Alalwan et al. [124] still underscores the growing interest in GAN-assisted augmentation as a privacy-preserving enabler of robust federated diagnostics. Peng et al. [107] introduce the novel FedNI framework in the neuroimaging domain. FedNI employs a federated GAN to augment local datasets, represented as disjoint small-scale graphs partitioned from a global population graph, for downstream disease prediction tasks. In brain tumor segmentation, the Wasserstein GAN-based architecture FLWGAN [109] is used to generate synthetic brain MRIs locally. FLWGAN supplements real samples to train 3D U-Nets across federated clients, effectively avoiding mode collapse and demonstrating strong performance on public tumor segmentation benchmarks. Many efforts have been directed toward chest X-ray images in recent years, particularly due to their central role in diagnosing respiratory conditions during the COVID-19 pandemic. This surge of interest has made chest radiographs a key focus for data augmentation strategies using GANs, especially in privacy-sensitive and data-scarce environments. Han et al. [87], for instance, anticipated the pandemic with a study on pulmonary diseases. Their work validates the utility and privacy of synthetic data through radiograph similarity tests with clinician input, confirming that GANs can augment without duplicating real patient data. In the context of COVID-19 diagnosis, Nguyen et al. [94] introduce a framework for chest X-ray data augmentation using a GAN model within a federated edge-cloud architecture incorporating differential privacy and blockchain. This approach ensures both scalability and data protection, reinforcing the value of GANs for secure and effective medical image synthesis in pandemic response efforts.

Data Balancing: Among the various applications of data augmentation, data balancing serves as a targeted strategy to mitigate the effects of non-IID data distributions and site-specific class imbalances. Generating synthetic samples for underrepresented classes helps harmonize local datasets across federated clients, facilitating more stable and equitable model training. Similarly to the previously analyzed works, the FedDPGAN framework by Zhang et al. [89] employs the same data modality for the diagnosis of COVID-19. It combines federated GAN training with DP mechanisms, enabling clients with severely non-IID class distributions to rebalance datasets before classification through local ResNets, specifically applied to chest X-ray diagnosis of COVID-19 pneumonia. Furthermore, Lee et al. [106] introduce a multi-site CT scan analysis for COVID-19 diagnosis, spreading federated COVID-19 data balancing to image volumes analysis. Their framework leverages a GAN with ResNet-based architecture trained independently at each site to reflect the imaging heterogeneity inherent to localized acquisition protocols, producing balanced augmentation without compromising privacy. Extending the dataset balancing to neurological imaging, the work on Alzheimer's disease classification by Nimeshika et al. [127] explores the integration of split FL with GAN-based class balancing in non-IID settings. The framework applies generative oversampling locally while employing a split learning strategy to decouple feature extraction and classification across clients, aiming to preserve privacy while improving model performance under heterogeneous data distributions.

Domain Translation: Federated GANs have also proven to be useful in normalizing stylistic and protocol-induced variations across medical imaging sites. In histopathology, Shen et al.'s [102] orchestral stain-normalization GAN shares a single generator across clients, each with a unique discriminator trained on its local stain style, using adversarial, pattern-preserving L1, and temporal self-distillation losses to produce uniformly stained tissue images without centralizing data. Similarly, Yan et al. [95] apply per-site CycleGANs to translate each client's scans into the style of the site with the simplest image distribution: that target site first trains a WGAN (with a noisy discriminator) to synthesize reference-style images, which are then shared so all other clients can map their data into that common domain before standard FedAvg training. In the neuroimaging realm, Wang et al.'s [108] FedMed-GAN benchmarks unsupervised cross-modality synthesis: translating among MRI weightings of T1, T2, and PD by federating a CycleGAN-style architecture. Bdair et al. [128] add spatial self-attention to a federated CycleGAN for MRI to CT translation, synchronizing the UNet-based generators across four sites via FedAvg to yield privacy-preserving CTs from paired 2D slices. In the field of CT imaging Georgiadis et al. [97] prestandardize heterogeneous COVID-19 CT scans via locally trained CycleGANs that map each site's raw images into a small, shared target style (determined by a jointly held reference set) before federated segmentation; Li et al. [126] extend this by coarsely translating CT scans from different reconstruction methods via CycleGAN, then refining outputs with a score-based diffusion model to address distribution shifts in large-scale multi-center studies; Abdel-Basset et al. [103] encode multi-institutional chest radiographs into feature vectors, translate these features via a CycleGAN-like module into a unified latent space, and decode them for COVID-19 segmentation. In the area of skin lesions, Kundu et al. [114]generate balanced monkeypox and nonmonkeypox skin images locally via CycleGAN augmentation, then federate a classifier on the harmonized data. In MRI images, Mazher et al. [117]fuse spatial–temporal GAN translations with transformer-based segmentation to handle non-IID 4D cardiac MRI sequences. Finally, Murmu et al.'s [121] CusIAFL framework integrates a Pix2Pix GAN for secure image-to-image translation, adding GAN-generated encryption keys in diverse image datasets (including colon cancer and skin lesions).

Further Applications: Several federated GANs focus on image synthesis without an immediate downstream augmentation, balancing, or translation goal. Rajotte et al. [93] create a PrivGAN extension that employs a centralized discriminator to coordinate across biased, non-IID sites, enabling small or skewed clients to "borrow" diversity from larger peers while preserving privacy. Elmas et al.'s [101] FedGIMP tackles MRI reconstruction by federating a generator—conditioned on a random vector plus a one-hot site index—against local discriminators: the generator learns site-specific image priors that are iteratively refined against each client's raw k-space data to recover fully sampled images. Pennisi et al. [100] skip federation during synthesis yet highlight the power of latent-space manipulation: they train StyleGAN2 locally, cluster its latent codes via LPIPS, interpolate between cluster centroids to generate new images, aggregate everyone's synthetics into a global dataset, and train a central classifier. Kim et al. [88] use a federated encoder that obfuscates patient identity: each site trains an encoder to fool a local discriminator (trying to link images to subjects) while preserving segmentation accuracy in a downstream network that only ever sees anonymized feature maps. Finally, Jin et al. [99] expose the vulnerability of federated synthesis by showing how a few malicious clients can poison a global GAN with trigger patches; they propose combining outlier-detection (Isolation Forest) to downweight bad updates and a WGAN-GP loss to smooth gradients, jointly neutralizing such attacks.

Tabular Data

While GANs have shown great promise in synthesizing image and signal data, their application to tabular clinical data in federated settings remains comparatively limited. Most research on federated generative modeling has focused on medical imaging, where data structure and evaluation are more standardized. Nonetheless, tabular data are of paramount importance in healthcare, given the widespread use of tabular EHRs, and several recent studies have begun to explore this domain. Weldon et al. [91] train local GANs on multi-institutional EHR data (MIMIC-III) and aggregate the outputs into a central model, showing comparable quality to centralized training but highlighting privacy and representativeness concerns. Fang et al. [98] propose DPCT-GAN, a conditional GAN trained with FedAvg and local differential privacy across nine datasets, supporting controlled synthetic data generation with strong privacy guarantees. Focusing on imputation, MIFC-SVFL [123] uses a CGAN in a vertical FL setting to impute metadata and image embeddings treated as tabular features. Finally, Rehman et al. [115] propose FedCSCD-GAN, which combines anonymization and class balancing to enhance cancer diagnosis performance across distributed tabular datasets.

Sequential Data

Beyond imaging, several works have adapted GANs to sequential clinical data within a federated framework. Brophy et al. [92] authors recast photoplethysmogram (PPG)-to arterial blood pressure (ABP) inference as a CycleGAN task on time series: each generator is a stack of recurrent layers plus a final dense output, discriminators are CNNs, and both are trained under FedAvg on local PPG–ABP pairs. A separate study on respiratory-disease classification uses GANs at each node to augment lung-sound recordings before federating a downstream classifier; by enriching rare classes locally, the system achieves more balanced training data and improved multiclass performance under a probability-ranked, weighted-aggregation scheme. Finally, FedSepsis integrates a GAN-based imputation model (GAIN) into a multimodal IoMT pipeline for early sepsis detection: on edge devices, clinical notes are embedded via ClinicalBERT, vitals are modeled with an LSTM predictor, and missing values are filled by GAIN; models are synchronized across Raspberry Pi and Jetson Nano nodes via FedAvg (and FedOpt), yielding a deployable, privacy-preserving system that handles heterogeneous EHR streams without sacrificing practical throughput.

B. AR models

One of the main challenges in federating AR models is data throughput. AR models can have a large number of parameters, making updating all parameters unfeasible in federated environments. A common solution is to employ a static model and use lightweight adapters between some layers that are then federated [90].

Two recent works apply federated LLM methods to clinical text tasks. Puppala et al. [119] build a personalized medical chatbot by fine-tuning a frozen GPT-2 at each site using LoRA adapters [129] and FedAvg; after aggregation, they feed local articles into the shared model to extract question–answer–source triplets, which are stored in a vector search database so end users retrieve contextually relevant answers without ever sharing raw documents. Yan et al. [125] also built an LLM for report generation and question answering, although their work tackles vision–language fusion. They freeze a LLaMA-2 text encoder while fine-tuning a Swin Transformer visual encoder (via LoRA), and then they project image embeddings into the text space. To support hospitals missing one modality, each client computes local embedding prototypes (k-means centers) for text and images, sharing them along with an interprototype correlation matrix, thereby aligning cross-modal features during training.

Also dealing with multimodal data and text generation, [116] Chen et al. federate a multimodal report generator for chest X-rays. Each site trains an image feature extractor, a disease-label classifier, and a transformer-based report module that conditions on class predictions, visual features, and any existing text notes. Models are aggregated via a mix of FedAvg, FedProx [49] (which penalizes divergence from the global model), and a performance-weighted FedSW scheme [130]. This produces a unified encoder-decoder pipeline capable of crafting coherent radiology reports without exposing raw images or text.

In a different data domain, [92] Brophy et al. demonstrate how adversarial and AR ideas can combine to address physiological signal generation. They perform domain translation from photoplethysmogram (PPG) to arterial blood pressure

(ABP) using a CycleGAN-style on time series data. To adapt the architecture, each generator is a stack of recurrent networks (LSTM layers) plus a final fully-connected output, while each discriminator is a convolutional network. Clients train only with FedAvg on local PPG–ABP pairs; no raw signal leaves the hospital.

By definition, the next outputs of AR models are always conditioned on the previous outputs of the model. Sequential data of this nature are common in healthcare settings, where longitudinal studies give updates on clinical markers over time. There are currently no federated architectures exploring the generation of longitudinal clinical data, but it has been proven that it is possible to employ AR models with local training [131].

C. VAEs architecture

Despite the widespread application of VAEs in centralized settings, their presence in FL environments, where data privacy, distributional heterogeneity, and communication efficiency are critical, is surprisingly limited.

Among the reviewed literature, only three healthcare studies stand out for explicitly focusing on a VAE architecture within a federated framework. Zhao et al. [105] address the issue of data scarcity within individual data silos in FL networks by introducing a data augmentation pipeline for brain MRI slices in Alzheimer's patients. In their work, clients share latent space distributions produced by locally trained VAE models, and sample from the aggregated distribution to enrich each client's dataset. This privacy-preserving approach enables the federated training of a classification model, leading to improved diagnostic performance.

Yang et al. [77] introduced the FedVAE framework, which is particularly interesting for its application to sequential data, a domain where time-dependent relationships are crucial. The FedVAE study explores this challenge by applying its framework to the MobiACT [132] and MobiFALL [133] public datasets, both designed to capture human activity over time. By operating over sequences and learning latent representations that encapsulate temporal dynamics, the FedVAE framework ensures that essential time-dependent features are not lost during the compression process. Another important advantage of the use of a VAE lies in the possibility of performing dimensionality reduction directly at the client nodes. Given the substantial volume of sensor-based time-series data stored locally, the VAE encodes high-dimensional data into compact latent representations before communication with the central server, thereby reducing the overall communication overhead. Furthermore, the generative capabilities of the VAE are leveraged to augment minority classes within local datasets, contributing to improved class balance and enhancing the performance of downstream classification models.

Meanwhile, Sikandar et al. [120] explored the use of a VAE architecture on amino acid sequences from cancer patients to enhance the performance of AI-based diagnostic models. Despite the inherently sequential nature of the data, their approach involved a preprocessing step that transformed the sequences into a tabular format, thereby discarding sequential dependencies. The study primarily focused on addressing local class imbalance through the generation of synthetic samples, aiming to improve the generalization on downstream classification tasks.

D. Diffusion models

As previously noted in the VAEs section, most research on diffusion models has been conducted in centralized environments, and their integration into FL frameworks remains very limited. So far, the healthcare literature features only a few notable studies, all of which have been published in the last two years.

Li et al. [126], with their study on CT scans of COVID-19 patients, address the challenge of generating images within a unified domain independent of the different reconstruction algorithms used to acquire the CT scans. They employ a combination of a CycleGAN model, used to generate coarse CTs in a unified domain, and a Denoising Diffusion Probabilistic Model (DDPM) [27] to refine the generated CT, obtaining the definitive image. Moreover, by testing synthetic data produced with evaluation metrics and segmentation model performance, they confirm the potential of GAN-based models to increase the robustness of ML models through augmentation. Nevertheless, they prize diffusion models for their suitability for increasing the trustworthiness and reliability of synthesized medical data.

The work by Tun et al. [110] targets the synthetic generation of chest X-ray images. This study implements a DDPM with a U-Net-based denoising model trained across distributed clients on a dataset of roughly 2,500 chest X-ray images. The main objective was to assess the feasibility of diffusion models within federated settings, specifically testing varying numbers of clients and introducing non-IID data configurations to simulate real-world conditions. The results highlight one of the challenges of diffusion models: despite their generative strength, they require large-scale datasets to perform optimally. In this case, the limited data size constrained the model's generative quality, making the process computationally expensive relative to the output.

Furthermore, the study by Sattarov et al. [113] shifts the focus to tabular data. FedTabDiff also tackles synthetic generation tasks in different federated configurations using a large dataset comprising health records from over 100,000 diabetes patients, including both categorical and numerical variables. To accommodate the tabular nature of the data, the authors replace the typical U-Net architecture for image domains with an MLP network as the denoising backbone. Unlike the previous vision task, FedTabDiff demonstrates more robust generative performance, largely attributed to the abundance of training data, which better suits the data-hungry nature of diffusion models.

Lastly, Li et al. [118] investigated an extension of the DDPM that operates on latent representations of real data. This type of architecture, known as a Latent Diffusion Model (LDM) [134], enables more efficient training and sampling by performing the diffusion process in a lower-dimensional latent representation. The study focuses on generating a synthetic dataset of histopathology images, selected for their high informativeness and visual interpretability. A large number of samples is initially generated and subsequently distilled into a smaller, representative subset using a graph-based selection strategy. While the architectural advantages of LDM, particularly its computational efficiency, make it well-suited for this task, the approach still requires access to a sufficiently large training dataset.

E. Statistical analysis

As a last step to better understand the trends and characteristics of DGMs applied to healthcare within FL architectures, we conducted a statistical analysis on the selection of papers. Table III and Table II provide a summary of the main characteristics of each study: the former focuses on the healthcare application, while the latter highlights the architectural aspects. The analysis revealed several key insights that may be useful for researchers who want to navigate the present topic.

To provide an overview of the interest in the topic addressed by this review, we analyzed the number of publications per year among the selected studies. As illustrated in Figure 3, there is a clear upward trend in the number of articles published over time. This increase reflects a growing interest and importance of the topic within the scientific community, as well as an expanding body of research contributing to its development.

Furthermore, from a general analysis of the reviewed literature, an important observation is the consistent adoption of a centralized FL topology, where a central server coordinates the training process. Notably in their work, Nguyen et al. [94] introduce the only example of decentralized design in the field of healthcare applications, with a blockchain-based peer-to-peer FL to enhance the security against attacks on the central node. In a few studies, instead of exchanging model weights as in standard FL processes, clients and the server exchanged latent representations or directly generated synthetic samples, which required the use of tailored aggregation methods. Overall, apart from these outliers, the uniformity in both the architectural design and the aggregation strategy points to an early consensus within the community. However, it also underscores a limited examination of alternative approaches, suggesting that the exploration in this field is still in its early stages, offering significant opportunities for further research.

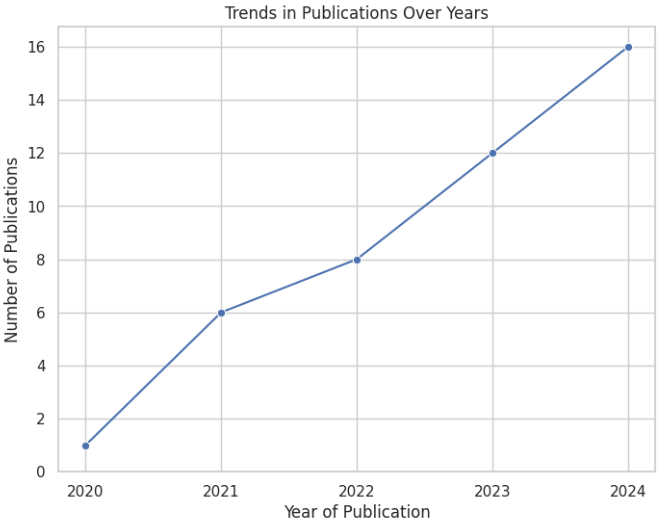

Fig. 3: The graph represents the trend of the studies published each year in the range selected for the present review

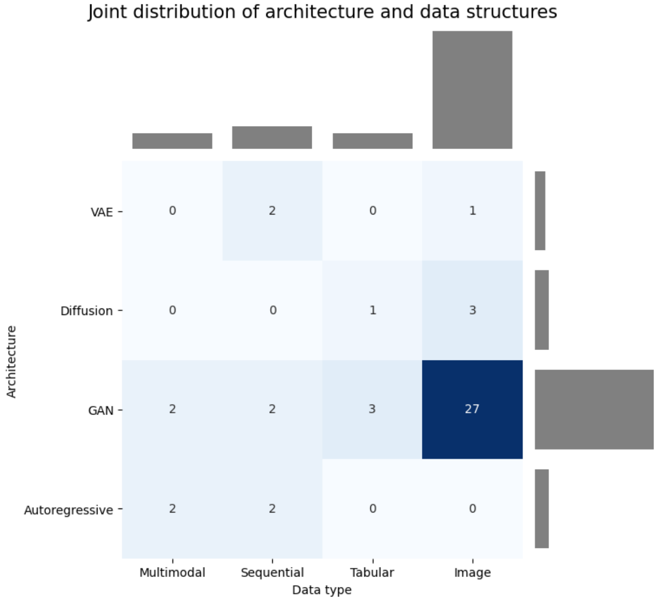

Fig. 4: The central heatmap displays the frequency of co-occurrence between different model architectures (e.g., GAN, VAE) and data types (e.g., Image, Tabular). Marginal histograms summarize the overall usage frequency of each architecture (top) and each data type (right). Color gradients in the central histograms reflect relative usage intensity. This visualization highlights common pairings and reveals preferences or gaps in current research practices.

Moreover, the analysis delves into the main focus of this work, the DGM architecture. The most commonly employed architecture was GAN, which accounted for the majority of occurrences, as we can see from Figure 4. The other architectures, VAE, diffusion models, and AR models, were used less frequently. This highlights the dominance of GANs in healthcare generative tasks, likely due to their versatility in data synthesis and augmentation.

The prevalence of GANs likely reflects a skewed distribution of data structures in FL applications, which are highly related to the architectural requirements of these types of models. In this regard, between the studies employing a federated framework, image data was by far the most prevalent type, followed by sequential data and tabular data, as emerges from Figure 4. This reflects the significant focus on medical imaging tasks, such as disease diagnosis and augmentation of imaging datasets. Such type of data, in fact, is often the scarcest due to various technical challenges associated with medical imaging [135], which often necessitates relying on FL to enable the use of data-hungry models, such as DGMs, through various institutions without requiring centralized access to large-scale imaging datasets.

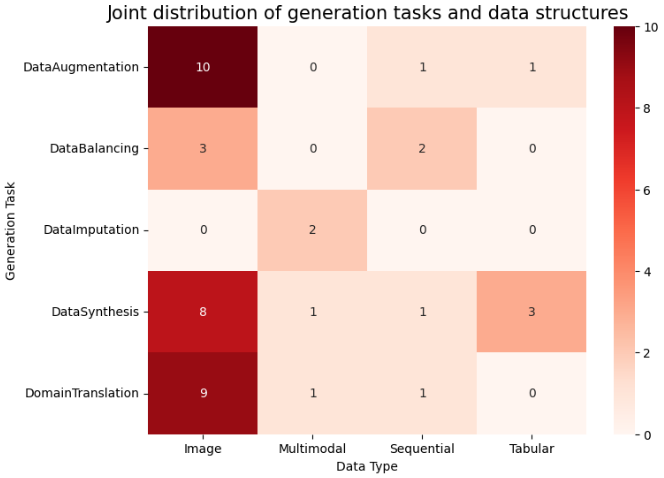

Fig. 5: The central heatmap shows the frequency of co-occurrence between different generative tasks and the data structures they target. The visualization highlights dominant pairings, such as image-based data for augmentation, while also revealing underexplored areas in federated generative modeling for healthcare.

Furthermore, synthetic data generation is a paradigm that can support a wide range of tasks (defined in Table II); among them, augmentation and domain translation tasks are the most common within the scope of this work. The abundance of augmentation pipelines in the reviewed literature emphasizes the role of DGMs in addressing data scarcity and imbalance in healthcare datasets. Domain translation tasks also stand out among the others. In many of the referenced studies, it is leveraged to overcome non-IID contexts between the client datasets, for which DGMs are revealed as powerful tools. The other tasks, which are imputation and balancing, were less frequent, but showcased innovative applications, highlighting the unexplored potential of DGMs in federated healthcare settings. Imputation, in particular, was applied only in studies involving multiple data modalities, as illustrated in Figure 5, which presents an analysis of task types in relation to the underlying data structures. This reveals an interesting trend: imputation is predominantly associated with heterogeneous data scenarios, where integrating information across modalities is essential for model performance and consistency.

These findings provide a comprehensive overview of the current landscape of DGMs in healthcare FL. They highlight both the strengths, such as the widespread adoption of GANs and the focus on image data, and the gaps, such as limited pipelines addressing multimodal data, and the exploration of other architectures. Addressing these gaps will be crucial for advancing the field and ensuring the practical applicability of these models in real-world healthcare settings.

VI. CONCLUSION AND FUTURE LINES

This review systematically examined the integration of DGMs within FL frameworks in healthcare, highlighting prevalent trends, strengths, and critical limitations. Synthetic data generation has emerged as a powerful paradigm to address real-world constraints on sensitive data usage, facilitating downstream healthcare tasks such as augmentation, domain translation, imputation, and class balancing. Our analysis revealed a pronounced dominance of GAN architectures, driven largely by their versatility and suitability for prevalent healthcare data modalities, notably image-based datasets.

However, this popularity also underscores significant biases and gaps in the current literature. The strong focus on image data and GAN architectures has resulted in comparatively limited exploration and development of alternative DGMs such as VAE, AR, and diffusion models. Consequently, the field exhibits a clear imbalance, indicating substantial opportunities for innovation and diversification. VAEs offer benefits for privacy-aware compression and can reduce communication costs in federated settings, especially for time-series or omics data. Diffusion models produce high-quality samples but remain underused due to their computational demands and limited use beyond synthetic data generation. AR models, well-suited for sequential and multimodal tasks, have recently been adapted to federated environments via lightweight fine-tuning (e.g., LoRA), enabling clinical applications like chatbots and report generation. Yet, their potential for modeling longitudinal synthetic data remains largely unexplored.

Beyond architecture-specific considerations, a major bottleneck is the absence of standardized benchmarks and robust validation protocols. The lack of shared datasets, metrics, and evaluation pipelines makes it difficult to compare methods objectively, assess generalization, or guide model selection. Moreover, few studies explicitly address regulatory constraints, such as data protection laws, fairness, or auditability (factors that are critical for clinical adoption).

Looking forward, future research should concretely focus on establishing standardized evaluation benchmarks and comparative studies across multiple DGM architectures. Regulatory frameworks concerning data privacy, security, and compliance should be explicitly considered in future methodological developments, shaping how DGMs and FL systems can be safely and effectively deployed in clinical settings. Furthermore, researchers should actively explore AR and diffusion models for sequential and tabular data types respectively, as well as alternative architectures such as flow-based and energy-based DGMs, diversifying beyond GAN-centric solutions. Additionally, addressing multimodal data scenarios, an area notably lacking in current literature, will significantly enrich the capability and applicability of FL-based DGMs in healthcare. Addressing these areas will not only enrich the landscape of DGMs in healthcare FL but also pave the way for compliant, secure, and clinically impactful generative solutions.

Appendix A
Summary of Data Sources

TABLE III: Overview of healthcare datasets used in generative FL studies analyzed. Underlined words denote the name of a dataset

Paper	**Data Source**
Han et al. [87]	ChestX-ray:112 120 Chest X-rays from 30 805 subjects CheXpert: 224 316 Chest X-rays from 65 240 subjects CXR8: 30 000 Chest X-rays
Kim et al. [88]	773 T1 weighted MRIs from 350 subjects
Zhang et al. [89]	3600 Chest X-rays
Cai et al. [90]	ISIC dataset: 10016 skin lesion samples
Weldon et al. [91]	Patient ICD-9 diagnoses: 46 520 subjects and 1 071 variables
Brophy et al. [92]	Training: 144 000 eight-second segments (320 hours) Validation: 55 000 eight-second segments (122 hours) Test: 900 ten-second segments (150 minutes, one patient)
Rajotte et al. [93]	HAM10000: 10 000 histological images aprox
Nguyen et al. [94]	DarkCOVID: 620 Chest X-rays ChestCOVID: 950 Chest X-rays
Yan et al. [95]	LocalPCa: 135 DWI images PROSTATEx: 188 DWI images
Lan et al. [96]	ISIC: 10 015 skin lesion samples
Georgiadis et al. [97]	Coronacase: 20 Chest CT images MedSeg: 100 Chest CT images
Fang et al. [98]	5 clinical datasets: eICU: 200 000 records MIMIC-III: 40 000 records Diabetes130-US: 100 000 records Breast Cancer: 569 records Heart Failure: 300 records
Jin et al. [99]	ISIC-2018: 10 015 skin lesion samples
Pennisi et al. [100]	Montgomery County:130 Chest X-rays Shenzhen Hospital: 662 Chest X-rays
Elmas et al. [101]	T1-, T2-, PD-weighted MRIs originated from: In-house dataset - 10 subjects IXI - 600 subjects fastMRI Brain - 6 970 MRIs fastMRI Knee - 1 500 MRIs BRATS - 1 200 subjects
Shen et al. [102]	Cancer Genome Atlas (TCGA): 100 histopathology slides CRC-VAL-HE-7K: 25 WSIs NCT-CRC-HE-100K: 86 WSIs
Abdel-Basset et al. [103]	COVID-19 CT scans: 829 annotated slices from 9 CT scans COVID-19-CT-Seg: $>$ 1800 annotated slices from 20 CT scans MosMedData: annotated scans from 50 CT scans
Ul Alam et al. [104]	Training: 14 334 - Testing: 1 824 - Tuning: 1 796 TOTAL: 17,954 samples (Very imbalanced (90%-10%))
Zhao et al. [105]	6400 MRI slices
Lee et al. [106]	Multisite set with 1 555 661 CT scans from 10 275 subjects
Yang et al. [77]	Sensor-based recordings: 13 daily activities from 57 subjects

Continued on next page

Paper	**Data Source**
Peng et al. [107]	ABIDE: 1 029 fMRIs ADNI: 911 fMRIs
Wang et al. [108]	T1, T2, and FLAIR MRIs for 580 subjects
Peketi et al. [109]	MSD: 484 3D brain T1, T2 and FLAIR MRIs
Tun et al. [110]	2 482 CT scans
Jin et al. [111]	ISIC: 3 297 skin lesion samples ChestX dataset: 5 856 Chest X-rays
Shaikh et al. [112]	ICBHI: 920 audio samples from 126 subjects JUT dataset: audio samples from 16 subjects
Sattarov et al. [113]	101 767 subjects 48 variables
Kundu et al. [114]	Monkeypox Skin Lesion Dataset (MSLD) Monkeypox Skin Images Dataset (MSID)
Rehman et al. [115]	Prostate cancer: 100 subjects 9 variables Lung cancer: 309 subjects 15 variables Breast cancer: 569 subjects 32 variables
Chen et al. [116]	IU-Xray: 2 955 articulable reports 5 910 frontal/lateral chest X-ray images
Mazher et al. [117]	M&Ms: 375 Cardiac MRIs ACDC Challenge: 150 Cardiac MRIs
Li et al. [118]	PathMNIST: 107 180 colon biopsies
Puppala et al. [119]	Unstructured text and prompts
Sikandar et al. [120]	IntOGen and UniProt: 858 samples of 1104 amino acid sequences
Murmu et al. [121]	LC25000: 10 000 histopathological images HAM10000: 10 015 images Satellite: 6 000 images Sandstone: 4 000 images Intel Image: 25 000 images Sea: 4,000 images
Selvaraj et al. [122]	PAD-UFES-20: 2 298 skin smartphone images HAM10000: 10 015 skin smartphone images PH2: 200 skin smartphone images
Chen et al. [123]	PAD-UFES: 2 298 records ISIC 2019: 25 331 records
Alalwan et al. [124]	Br35H: 3 060 MRI slices
Yan et al. [125]	IU-Xray (radiology report generation) 3 955 anonymized radiology reports 7 470 chest X-ray images VQA-RAD (visual QA) 314 radiology images 2 248 question–answer pairs Slake-English (visual QA) 642 images 7 033 question–answer pairs
Li et al. [126]	213 CT scans from two reconstruction processes
Nimeshika et al. [127]	MRI images dataset: 224 Mild Dementia patients 16 Moderate Dementia patients 800 Non-Dementia patients
Tariq Bdair et al. [128]	SynthRAD2023:180 paired MRI-CT brain scans

REFERENCES

[1] M. Bekbolatova, J. Mayer, C. W. Ong, and M. Toma, "Transformative Potential of AI in Healthcare: Definitions, Applications, and Navigating the Ethical Landscape and Public Perspectives," *Healthcare*, vol. 12, no. 2, p. 125, Jan. 2024, number: 2 Publisher: Multidisciplinary Digital Publishing Institute. [Online]. Available: https://www.mdpi.com/2227-9032/12/2/125

[2] A. Al Kuwaiti, K. Nazer, A. Al-Reedy, S. Al-Shehri, A. Al-Muhanna, A. V. Subbarayalu, D. Al Muhanna, and F. A. Al-Muhanna, "A Review of the Role of Artificial Intelligence in Healthcare," *Journal of Personalized Medicine*, vol. 13, no. 6, p. 951, Jun. 2023, number: 6 Publisher: Multidisciplinary Digital Publishing Institute. [Online]. Available: https://www.mdpi.com/2075-4426/13/6/951

[3] S. Mullankandy, S. Mukherjee, and B. S. Ingole, "Applications of AI in Electronic Health Records, Challenges, and Mitigation Strategies," in *2024 International Conference on Computer and Applications (ICCA)*, Dec. 2024, pp. 1–7. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/10927863

[4] I. Castiglioni, L. Rundo, M. Codari, G. Di Leo, C. Salvatore, M. Interlenghi, F. Gallivanone, A. Cozzi, N. C. D'Amico, and F. Sardanelli, "AI applications to medical images: From machine learning to deep learning," *Physica Medica*, vol. 83, pp. 9–24, Mar. 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1120179721000946

[5] F. Subhan, A. Mirza, M. B. M. Su'ud, M. M. Alam, S. Nisar, U. Habib, and M. Z. Iqbal, "AI-Enabled Wearable Medical Internet of Things in Healthcare System: A Survey," *Applied Sciences*, vol. 13, no. 3, p. 1394, Jan. 2023, number: 3 Publisher: Multidisciplinary Digital Publishing Institute. [Online]. Available: https://www.mdpi.com/2076-3417/13/3/1394

[6] S. M. Williamson and V. Prybutok, "Balancing Privacy and Progress: A Review of Privacy Challenges, Systemic Oversight, and Patient Perceptions in AI-Driven Healthcare," *Applied Sciences*, vol. 14, no. 2, p. 675, Jan. 2024, number: 2 Publisher: Multidisciplinary Digital Publishing Institute. [Online]. Available: https://www.mdpi.com/2076-3417/14/2/675

[7] M. Wojtara, E. Rana, T. Rahman, P. Khanna, and H. Singh, "Artificial intelligence in rare disease diagnosis and treatment," *Clinical and Translational Science*, vol. 16, no. 11, pp. 2106–2111, Nov. 2023.

[8] P. Zhang and M. N. Kamel Boulos, "Generative AI in Medicine and Healthcare: Promises, Opportunities and Challenges," *Future Internet*, vol. 15, no. 9, p. 286, Sep. 2023, number: 9 Publisher: Multidisciplinary Digital Publishing Institute. [Online]. Available: https://www.mdpi.com/1999-5903/15/9/286

[9] M. Frid-Adar, E. Klang, M. Amitai, J. Goldberger, and H. Greenspan, "Synthetic data augmentation using GAN for improved liver lesion classification," in *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, Apr. 2018, pp. 289–293, iSSN: 1945-8452. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/8363576

[10] J. Yoon, J. Jordon, and M. Schaar, "GAIN: Missing Data Imputation using Generative Adversarial Nets," in *Proceedings of the 35th International Conference on Machine Learning*. PMLR, Jul. 2018, pp. 5689–5698, iSSN: 2640-3498. [Online]. Available: https://proceedings.mlr.press/v80/yoon18a.html

[11] M. K. Baowaly, C.-C. Lin, C.-L. Liu, and K.-T. Chen, "Synthesizing electronic health records using improved generative adversarial networks," *Journal of the American Medical Informatics Association: JAMIA*, vol. 26, no. 3, pp. 228–241, Mar. 2019.

[12] J. Li, B. J. Cairns, J. Li, and T. Zhu, "Generating synthetic mixed-type longitudinal electronic health records for artificial intelligent applications," *npj Digital Medicine*, vol. 6, no. 1, pp. 1–18, May 2023, publisher: Nature Publishing Group. [Online]. Available: https://www.nature.com/articles/s41746-023-00834-7

[13] European Union, "Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation)," May 2016. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2016/679/oj

[14] Directorate-General for Health and Food Safety (DG SANTE), European Commission, "Regulation (EU) 2025/327 of the European Parliament and of the Council on the European Health Data Space," 2025. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2025/327/oj

[15] U.S. Department of Health and Human Services, Office for Civil Rights, "Notice of Proposed Rulemaking to Update the HIPAA Security Rule," 2025. [Online]. Available: https://www.hhs.gov/hipaa/for-professionals/security/hipaa-security-rule-nprm/index.html

[16] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y. Arcas, "Communication-Efficient Learning of Deep Networks from Decentralized Data," in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*. PMLR, Apr. 2017, pp. 1273–1282, iSSN: 2640-3498. [Online]. Available: https://proceedings.mlr.press/v54/mcmahan17a.html

[17] O. Goldreich, "Secure multi-party computation," vol. 78, no. 110, pp. 1–108, 1998.

[18] M. Ogburn, C. Turner, and P. Dahal, "Homomorphic Encryption," *Procedia Computer Science*, vol. 20, pp. 502–509, Jan. 2013. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1877050913011101

[19] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative Adversarial Nets," in *Advances in Neural Information Processing Systems*, vol. 27. Curran Associates, Inc., 2014. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2014/hash/f033ed80deb0234979a61f95710dbe25-Abstract.html

[20] D. P. Kingma and M. Welling, "Auto-Encoding Variational Bayes," *CoRR*, Dec. 2013. [Online]. Available: https://www.semanticscholar.org/paper/5f5dc5b9a2ba710937e2c413b37b053cd673df02

[21] E. Choi, S. Biswal, B. Malin, J. Duke, W. F. Stewart, and J. Sun, "Generating Multi-label Discrete Patient Records using Generative Adversarial Networks," in *Proceedings of the 2nd Machine Learning for Healthcare Conference*. PMLR, Nov. 2017, pp. 286–305, iSSN: 2640-3498. [Online]. Available: https://proceedings.mlr.press/v68/choi17a.html

[22] S. Sai, A. Gaur, R. Sai, V. Chamola, M. Guizani, and J. J. P. C. Rodrigues, "Generative AI for Transformative Healthcare: A Comprehensive Study of Emerging Models, Applications, Case Studies, and Limitations," *IEEE Access*, vol. 12, pp. 31 078–31 106, 2024. [Online]. Available: https://ieeexplore.ieee.org/document/10440330

[23] H. Gm, M. K. Gourisaria, M. Pandey, and S. S. Rautaray, "A comprehensive survey and analysis of generative models in machine learning," *Computer Science Review*, vol. 38, p. 100285, Nov. 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1574013720303853

[24] A. Lamb, "A Brief Introduction to Generative Models," Feb. 2021. [Online]. Available: https://arxiv.org/abs/2103.00265v1

[25] Z. Ou, "A Review of Learning with Deep Generative Models from Perspective of Graphical Modeling," 2018, version Number: 4. [Online]. Available: https://arxiv.org/abs/1808.01630

[26] T. White, "Sampling Generative Networks," Dec. 2016, arXiv:1609.04468. [Online]. Available: http://arxiv.org/abs/1609.04468

[27] J. Ho, A. Jain, and P. Abbeel, "Denoising Diffusion Probabilistic Models," Dec. 2020, arXiv:2006.11239. [Online]. Available: http://arxiv.org/abs/2006.11239

[28] D. J. Rezende and S. Mohamed, "Variational Inference with Normalizing Flows," Jun. 2016, arXiv:1505.05770. [Online]. Available: http://arxiv.org/abs/1505.05770

[29] D. H. Ackley, G. E. Hinton, and T. J. Sejnowski, "A learning algorithm for boltzmann machines," *Cognitive Science*, vol. 9, no. 1, pp. 147–169, Jan. 1985. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0364021385800124

[30] R. Challen, J. Denny, M. Pitt, L. Gompels, T. Edwards, and K. Tsaneva-Atanasova, "Artificial intelligence, bias and clinical safety," *BMJ Quality & Safety*, vol. 28, no. 3, pp. 231–237, Mar. 2019, publisher: BMJ Publishing Group Ltd Section: Narrative review. [Online]. Available: https://qualitysafety.bmj.com/content/28/3/231

[31] M. Frid-Adar, I. Diamant, E. Klang, M. Amitai, J. Goldberger, and H. Greenspan, "GAN-based Synthetic Medical Image Augmentation for increased CNN Performance in Liver Lesion Classification," Mar. 2018, arXiv:1803.01229. [Online]. Available: http://arxiv.org/abs/1803.01229

[32] R. J. Chen, M. Y. Lu, T. Y. Chen, D. F. K. Williamson, and F. Mahmood, "Synthetic data in machine learning for medicine and healthcare," *Nature Biomedical Engineering*, vol. 5, no. 6, pp. 493–497, Jun. 2021, publisher: Nature Publishing Group. [Online]. Available: https://www.nature.com/articles/s41551-021-00751-8

[33] J. Yoon, D. Jarrett, and M. van der Schaar, "Time-series Generative Adversarial Networks," in *Advances in Neural Information Processing Systems*, vol. 32. Curran Associates, Inc., 2019. [Online]. Available: https://papers.nips.cc/paper_files/paper/2019/hash/c9efe5f26cd17ba6216bbe2a7d26d490-Abstract.html

[34] B. Beaulieu-Jones, Z. Wu, C. Williams, R. Lee, S. Bhavnani, J. Byrd, and C. Greene, "Privacy-Preserving Generative Deep Neural Networks Support Clinical Data Sharing," *Circulation: Cardiovascular Quality and Outcomes*, vol. 12, Jul. 2019.

[35] N. Razavian, J. Marcus, and D. Sontag, "Multi-task Prediction of Disease Onsets from Longitudinal Laboratory Tests," in *Proceedings of the 1st Machine Learning for Healthcare Conference*. PMLR, Dec. 2016, pp. 73–100, iSSN: 1938-7228. [Online]. Available: https://proceedings.mlr.press/v56/Razavian16.html

[36] E. J. Topol, "High-performance medicine: the convergence of human and artificial intelligence," *Nature Medicine*, vol. 25, no. 1, pp. 44–56, Jan. 2019, publisher: Nature Publishing Group. [Online]. Available: https://www.nature.com/articles/s41591-018-0300-7

[37] "Introducing ChatGPT," Mar. 2024. [Online]. Available: https://openai.com/index/chatgpt/

[38] M. Dalal, A. C. Li, and R. Taori, "Autoregressive Models: What Are They Good For?" Oct. 2019. [Online]. Available: https://arxiv.org/abs/1910.07737v1

[39] G. Müller-Franzes, J. M. Niehues, F. Khader, S. T. Arasteh, C. Haarburger, C. Kuhl, T. Wang, T. Han, S. Nebelung, J. N. Kather, and D. Truhn, "Diffusion Probabilistic Models beat GANs on Medical Images," *Scientific Reports*, vol. 13, no. 1, p. 12098, Jul. 2023, arXiv:2212.07501 [eess]. [Online]. Available: http://arxiv.org/abs/2212.07501

[40] Y. Zhong, X. Wang, J. Wang, X. Zhang, Y. Wang, M. Huai, C. Xiao, and F. Ma, "Synthesizing Multimodal Electronic Health Records via Predictive Diffusion Models," Jun. 2024, arXiv:2406.13942 [cs]. [Online]. Available: http://arxiv.org/abs/2406.13942

[41] I. A. Majeed, S. Kaushik, A. Bardhan, V. S. K. Tadi, H.-K. Min, K. Kumaraguru, and R. D. Muni, "Comparative assessment of federated and centralized machine learning," Feb. 2022, arXiv:2202.01529. [Online]. Available: http://arxiv.org/abs/2202.01529

[42] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated Machine Learning: Concept and Applications," *ACM Trans. Intell. Syst. Technol.*, vol. 10, no. 2, pp. 12:1–12:19, 2019. [Online]. Available: https://dl.acm.org/doi/10.1145/3298981

[43] S. Ahn, S. Kim, Y. Kwon, J. Park, J. Youn, and S. Cho, "Communication-Efficient Diffusion Strategy for Performance Improvement of Federated Learning with Non-IID Data," Jan. 2023, arXiv:2207.07493 [cs]. [Online]. Available: http://arxiv.org/abs/2207.07493

[44] P. M. Mammen, "Federated Learning: Opportunities and Challenges," Jan. 2021, arXiv:2101.05428 [cs]. [Online]. Available: http://arxiv.org/abs/2101.05428

[45] N. Truong, K. Sun, S. Wang, F. Guitton, and Y. Guo, "Privacy preservation in federated learning: An insightful survey from the GDPR perspective," *Computers & Security*, vol. 110, p. 102402, Nov. 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0167404821002261

[46] L. O. Gostin, "National Health Information PrivacyRegulations Under the Health Insurance Portability and Accountability Act," *JAMA*, vol. 285, no. 23, pp. 3015–3021, Jun. 2001. [Online]. Available: https://doi.org/10.1001/jama.285.23.3015

[47] Z. Ji, Z. C. Lipton, and C. Elkan, "Differential Privacy and Machine Learning: a Survey and Review," Dec. 2014, arXiv:1412.7584. [Online]. Available: http://arxiv.org/abs/1412.7584

[48] Y. Lindell, "Secure multiparty computation," *Communications of the ACM*, vol. 64, no. 1, pp. 86–96, Jan. 2021. [Online]. Available: https://dl.acm.org/doi/10.1145/3387108

[49] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated Optimization in Heterogeneous Networks," Dec. 2018. [Online]. Available: https://arxiv.org/abs/1812.06127v5

[50] C. Zhang, Y. Xie, H. Bai, B. Yu, W. Li, and Y. Gao, "A survey on federated learning," *Knowledge-Based Systems*, vol. 216, p. 106775, Mar. 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0950705121000381

[51] M. Rujas, R. Martín Gómez del Moral Herranz, G. Fico, and B. Merino-Barbancho, "Synthetic data generation in healthcare: A scoping review of reviews on domains, motivations, and future applications," *International Journal of Medical Informatics*, vol. 195, p. 105763, Mar. 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S138650562400426X

[52] H. Thanh-Tung and T. Tran, "Toward a Generalization Metric for Deep Generative Models," May 2021, arXiv:2011.00754 [cs]. [Online]. Available: http://arxiv.org/abs/2011.00754

[53] L. Regenwetter, A. Srivastava, D. Gutfreund, and F. Ahmed, "Beyond Statistical Similarity: Rethinking Metrics for Deep Generative Models in Engineering Design," *Computer-Aided Design*, vol. 165, p. 103609, Dec. 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0010448523001410

[54] N. Chen, A. Klushyn, R. Kurle, X. Jiang, J. Bayer, and P. Smagt, "Metrics for Deep Generative Models," in *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*. PMLR, Mar. 2018, pp. 1540–1550, iSSN: 2640-3498. [Online]. Available: https://proceedings.mlr.press/v84/chen18e.html

[55] J.-O. Palacio-Niño and F. Berzal, "Evaluation Metrics for Unsupervised Learning Algorithms," May 2019. [Online]. Available: https://arxiv.org/abs/1905.05667v2

[56] O. Boursalie, R. Samavi, and T. E. Doyle, "Evaluation Metrics for Deep Learning Imputation Models," in *AI for Disease Surveillance and Pandemic Intelligence: Intelligent Disease Detection in Action*, A. Shaban-Nejad, M. Michalowski, and S. Bianco, Eds. Cham: Springer International Publishing, 2022, pp. 309–322. [Online]. Available: https://doi.org/10.1007/978-3-030-93080-6_22

[57] M. Hernandez, G. Epelde, A. Alberdi, R. Cilla, and D. Rankin, "Synthetic data generation for tabular health records: A systematic review," *Neurocomputing*, vol. 493, pp. 28–45, Jul. 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0925231222004349

[58] A. Figueira and B. Vaz, "Survey on Synthetic Data Generation, Evaluation Methods and GANs," *Mathematics*, vol. 10, no. 15, p. 2733, Jan. 2022, number: 15 Publisher: Multidisciplinary Digital Publishing Institute. [Online]. Available: https://www.mdpi.com/2227-7390/10/15/2733

[59] O. Mendelevitch and M. D. Lesh, "Fidelity and Privacy of Synthetic Medical Data," Jun. 2021, arXiv:2101.08658. [Online]. Available: http://arxiv.org/abs/2101.08658

[60] K. El Emam, "Seven Ways to Evaluate the Utility of Synthetic Data," *IEEE Security & Privacy*, vol. 18, no. 4, pp. 56–59, Jul. 2020, conference Name: IEEE Security & Privacy. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/9138552

[61] B. A. Galende, P. A. Apellániz, J. Parras, S. Zazo, and S. Uribe, "Membership Inference Attacks and Differential Privacy: A Study Within the Context of Generative Models," *IEEE Open Journal of the Computer Society*, vol. 6, pp. 801–811, 2025. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/11008817

[62] "Executive Summary for the Digital Health Advisory Committee Meeting: Total Product Lifecycle Considerations for Generative AI-Enabled Devices," U.S. Food & Drug Administration, Silver Spring, MD, Executive Summary, Nov. 2024. [Online]. Available: https://www.fda.gov/media/174897/download

[63] European Union, "Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)," Jul. 2024. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2024/1689/oj

[64] J. Hayes, L. Melis, G. Danezis, and E. De Cristofaro, "LOGAN: Membership Inference Attacks Against Generative Models," *Proceedings on Privacy Enhancing Technologies*, vol. 2019, no. 1, pp. 133–152, Jan. 2019. [Online]. Available: https://petsymposium.org/popets/2019/popets-2019-0008.php

[65] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, "Deep Learning with Differential Privacy," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. Vienna Austria: ACM, Oct. 2016, pp. 308–318. [Online]. Available: https://dl.acm.org/doi/10.1145/2976749.2978318

[66] E. Bagdasaryan, A. Veit, Y. Hua, D. Estrin, and V. Shmatikov, "How To Backdoor Federated Learning," in *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*. PMLR, Jun. 2020, pp. 2938–2948. [Online]. Available: https://proceedings.mlr.press/v108/bagdasaryan20a.html

[67] B. Hitaj, G. Ateniese, and F. Perez-Cruz, "Deep Models Under the GAN: Information Leakage from Collaborative Deep Learning," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. Dallas Texas USA: ACM, Oct. 2017, pp. 603–618. [Online]. Available: https://dl.acm.org/doi/10.1145/3133956.3134012

[68] F. Tramèr, F. Zhang, A. Juels, M. K. Reiter, and T. Ristenpart, "Stealing Machine Learning Models via Prediction {APIs}," 2016, pp. 601–618. [Online]. Available: https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/tramer

[69] J. G. Berryman, "Measurement of spatial correlation functions using image processing techniques," *Journal of Applied Physics*, vol. 57, no. 7, pp. 2374–2384, Apr. 1985. [Online]. Available: https://doi.org/10.1063/1.334346

[70] E. Alvarez-Lacalle, B. Dorow, J.-P. Eckmann, and E. Moses, “Hierarchical structures induce long-range dynamical correlations in written texts,” *Proceedings of the National Academy of Sciences*, vol. 103, no. 21, pp. 7956–7961, May 2006, publisher: Proceedings of the National Academy of Sciences. [Online]. Available: https://www.pnas.org/doi/abs/10.1073/pnas.0510673103

[71] R. Patel, A. K. Verma, N. K. Sahu, and P. Rai, “Generative AI Techniques in Medical Imaging Analysis: A Systematic Review,” in *Real-Time Data Decisions With AI and ChatGPT Techniques*. IGI Global, 2025, pp. 91–126. [Online]. Available: https://www.igi-global.com/chapter/generative-ai-techniques-in-medical-imaging-analysis/www.igi-global.com/chapter/generative-ai-techniques-in-medical-imaging-analysis/357191

[72] G. A. Kaissis, M. R. Makowski, D. Rückert, and R. F. Braren, “Secure, privacy-preserving and federated machine learning in medical imaging,” *Nature Machine Intelligence*, vol. 2, no. 6, pp. 305–311, Jun. 2020, publisher: Nature Publishing Group. [Online]. Available: https://www.nature.com/articles/s42256-020-0186-1

[73] J. A. Bartell, S. B. Valentin, A. Krogh, H. Langberg, and M. Bøgsted, “A primer on synthetic health data,” Jul. 2024, arXiv:2401.17653. [Online]. Available: http://arxiv.org/abs/2401.17653

[74] H. Y. J. Kang, M. Ko, and K. S. Ryu, “Tabular Transformer Generative Adversarial Network for Heterogeneous distribution in healthcare,” Mar. 2024, iSSN: 2693-5015. [Online]. Available: https://www.researchsquare.com/article/rs-4134206/v1

[75] G. O. Ghosheh, C. L. Thwaites, and T. Zhu, “Synthesizing Electronic Health Records for Predictive Models in Low-Middle-Income Countries (LMICs),” *Biomedicines*, vol. 11, no. 6, p. 1749, Jun. 2023, number: 6 Publisher: Multidisciplinary Digital Publishing Institute. [Online]. Available: https://www.mdpi.com/2227-9059/11/6/1749

[76] C. Ma, S. Tschiatschek, K. Palla, J. M. Hernández-Lobato, S. Nowozin, and C. Zhang, “EDDI: Efficient Dynamic Discovery of High-Value Information with Partial VAE,” May 2019, arXiv:1809.11142 [cs]. [Online]. Available: http://arxiv.org/abs/1809.11142

[77] H. Yang, M. Ge, K. Xiang, X. Bai, and H. Li, “FedVAE: Communication-Efficient Federated Learning With Non-IID Private Data,” *IEEE Systems Journal*, vol. 17, no. 3, pp. 4798–4808, Sep. 2023, conference Name: IEEE Systems Journal. [Online]. Available: https://ieeexplore.ieee.org/document/10132861

[78] K. Häyrinen, K. Saranto, and P. Nykänen, “Definition, structure, content, use and impacts of electronic health records: A review of the research literature,” *International Journal of Medical Informatics*, vol. 77, no. 5, pp. 291–304, May 2008. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1386505607001682

[79] R. S. Evans, “Electronic Health Records: Then, Now, and in the Future,” *Yearbook of Medical Informatics*, vol. 25, pp. S48–S61, Mar. 2018, publisher: Georg Thieme Verlag KG. [Online]. Available: https://www.thieme-connect.com/products/ejournals/abstract/10.15265/IYS-2016-s006

[80] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” *Nature*, vol. 521, no. 7553, pp. 436–444, May 2015, publisher: Nature Publishing Group. [Online]. Available: https://www.nature.com/articles/nature14539

[81] Y. Zhang, M. Chen, J. Shen, and C. Wang, “Tailor Versatile Multi-Modal Learning for Multi-Label Emotion Recognition,” *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 8, pp. 9100–9108, Jun. 2022, number: 8. [Online]. Available: https://ojs.aaai.org/index.php/AAAI/article/view/20895

[82] G. M. Dimitri, “A Short Survey on Deep Learning for Multimodal Integration: Applications, Future Perspectives and Challenges,” *Computers*, vol. 11, no. 11, p. 163, Nov. 2022, number: 11 Publisher: Multidisciplinary Digital Publishing Institute. [Online]. Available: https://www.mdpi.com/2073-431X/11/11/163

[83] T. Baltrušaitis, C. Ahuja, and L.-P. Morency, “Challenges and applications in multimodal machine learning,” in *The Handbook of Multimodal-Multisensor Interfaces: Signal Processing, Architectures, and Detection of Emotion and Cognition - Volume 2*. Association for Computing Machinery and Morgan & Claypool, Oct. 2018, vol. 21, pp. 17–48. [Online]. Available: https://doi.org/10.1145/3107990.3107993

[84] J. Lipkova, R. J. Chen, B. Chen, M. Y. Lu, M. Barbieri, D. Shao, A. J. Vaidya, C. Chen, L. Zhuang, D. F. K. Williamson, M. Shaban, T. Y. Chen, and F. Mahmood, “Artificial intelligence for multimodal data integration in oncology,” *Cancer Cell*, vol. 40, no. 10, pp. 1095–1110, Oct. 2022, publisher: Elsevier. [Online]. Available: https://www.cell.com/cancer-cell/abstract/S1535-6108(22)00441-X

[85] C. Cui, H. Yang, Y. Wang, S. Zhao, Z. Asad, L. A. Coburn, K. T. Wilson, B. A. Landman, and Y. Huo, “Deep multimodal fusion of image and non-image data in disease diagnosis and prognosis: a review,” *Progress in Biomedical Engineering*, vol. 5, no. 2, p. 022001, Apr. 2023, publisher: IOP Publishing. [Online]. Available: https://dx.doi.org/10.1088/2516-1091/acc2fe

[86] M. Lorenzi, M. Deprez, I. Balelli, A. L. Aguila, and A. Altmann, “Integration of Multimodal Data,” in *Machine Learning for Brain Disorders*, O. Colliot, Ed. New York, NY: Springer US, 2023, pp. 573–597. [Online]. Available: https://doi.org/10.1007/978-1-0716-3195-9_19

[87] T. Han, S. Nebelung, C. Haarburger, N. Horst, S. Reinartz, D. Merhof, F. Kiessling, V. Schulz, and D. Truhn, “Breaking medical data sharing boundaries by using synthesized radiographs,” *Science Advances*, vol. 6, no. 49, p. eabb7973, Dec. 2020, publisher: American Association for the Advancement of Science. [Online]. Available: https://www.science.org/doi/10.1126/sciadv.abb7973

[88] B. N. Kim, J. Dolz, P.-M. Jodoin, and C. Desrosiers, “Privacy-Net: An Adversarial Approach for Identity-Obfuscated Segmentation of Medical Images,” *IEEE Transactions on Medical Imaging*, vol. 40, no. 7, pp. 1737–1749, Jul. 2021, conference Name: IEEE Transactions on Medical Imaging. [Online]. Available: https://ieeexplore.ieee.org/document/9376890

[89] L. Zhang, B. Shen, A. Barnawi, S. Xi, N. Kumar, and Y. Wu, “FedDPGAN: Federated Differentially Private Generative Adversarial Networks Framework for the Detection of COVID-19 Pneumonia,” *Information Systems Frontiers: A Journal of Research and Innovation*, vol. 23, no. 6, pp. 1403–1415, 2021.

[90] X. Cai, Y. Lan, Z. Zhang, J. Wen, Z. Cui, and W. Zhang, “A Many-Objective Optimization Based Federal Deep Generation Model for Enhancing Data Processing Capability in IoT,” *IEEE Transactions on Industrial Informatics*, vol. 19, no. 1, pp. 561–569, Jan. 2023. [Online]. Available: https://ieeexplore.ieee.org/document/9472967

[91] J. Weldon, T. Ward, and E. Brophy, “Generation of Synthetic Electronic Health Records Using a Federated GAN,” Sep. 2021, arXiv:2109.02543 [cs]. [Online]. Available: http://arxiv.org/abs/2109.02543

[92] E. Brophy, M. De Vos, G. Boylan, and T. Ward, “Estimation of Continuous Blood Pressure from PPG via a Federated Learning Approach,” *Sensors*, vol. 21, no. 18, p. 6311, Jan. 2021, number: 18 Publisher: Multidisciplinary Digital Publishing Institute. [Online]. Available: https://www.mdpi.com/1424-8220/21/18/6311

[93] J.-F. Rajotte, S. Mukherjee, C. Robinson, A. Ortiz, C. West, J. M. L. Ferres, and R. T. Ng, “Reducing bias and increasing utility by federated generative modeling of medical images using a centralized adversary,” in *Proceedings of the Conference on Information Technology for Social Good*, ser. GoodIT ’21. New York, NY, USA: Association for Computing Machinery, Sep. 2021, pp. 79–84. [Online]. Available: https://dl.acm.org/doi/10.1145/3462203.3475875

[94] D. C. Nguyen, M. Ding, P. N. Pathirana, A. Seneviratne, and A. Y. Zomaya, “Federated Learning for COVID-19 Detection with Generative Adversarial Networks in Edge Cloud Computing,” *IEEE Internet of Things Journal*, vol. 9, no. 12, pp. 10 257–10 271, Jun. 2022, arXiv:2110.07136 [cs, eess]. [Online]. Available: http://arxiv.org/abs/2110.07136

[95] Z. Yan, J. Wicaksana, Z. Wang, X. Yang, and K.-T. Cheng, “Variation-Aware Federated Learning With Multi-Source Decentralized Medical Image Data,” *IEEE Journal of Biomedical and Health Informatics*, vol. 25, no. 7, pp. 2615–2628, Jul. 2021. [Online]. Available: https://ieeexplore.ieee.org/document/9268161

[96] Y. Lan, L. Xie, X. Cai, and L. Wang, “A many-objective evolutionary algorithm based on integrated strategy for skin cancer detection,” *KSII Transactions on Internet and Information Systems (TIIS)*, vol. 16, no. 1, pp. 80–96, 2022, publisher: Korean Society for Internet Information. [Online]. Available: https://koreascience.kr/article/JAKO202209059045674.page

[97] A. Georgiadis, V. Babbar, F. Silavong, S. Moran, and R. Otter, “ST-FL: style transfer preprocessing in federated learning for COVID-19 segmentation,” in *Medical Imaging 2022: Imaging Informatics for Healthcare, Research, and Applications*, vol. 12037. SPIE, Apr. 2022, pp. 13–27. [Online]. Available: https://www.spiedigitallibrary.org/conference-proceedings-of-spie/12037/1203704/ST-FL--style-transfer-preprocessing-in-federated-learning-for/10.1117/12.2611096.full

[98] M. L. Fang, D. S. Dhami, and K. Kersting, “DP-CTGAN: Differentially Private Medical Data Generation Using CTGANs,” in *Artificial Intelligence in Medicine*, M. Michalowski, S. S. R. Abidi, and S. Abidi, Eds. Cham: Springer International Publishing, 2022, pp. 178–188.

[99] R. Jin and X. Li, “Backdoor Attack is a Devil in Federated GAN-Based Medical Image Synthesis,” in *Simulation and Synthesis in Medical*

Imaging, C. Zhao, D. Svoboda, J. M. Wolterink, and M. Escobar, Eds. Cham: Springer International Publishing, 2022, pp. 154–165.

[100] M. Pennisi, F. Proietto Salanitri, S. Palazzo, C. Pino, F. Rundo, D. Giordano, and C. Spampinato, "GAN Latent Space Manipulation and Aggregation for Federated Learning in Medical Imaging," in *Distributed, Collaborative, and Federated Learning, and Affordable AI and Healthcare for Resource Diverse Global Health*, S. Albarqouni, S. Bakas, S. Bano, M. J. Cardoso, B. Khanal, B. Landman, X. Li, C. Qin, I. Rekik, N. Rieke, H. Roth, D. Sheet, and D. Xu, Eds. Cham: Springer Nature Switzerland, 2022, pp. 68–78.

[101] G. Elmas, S. U. H. Dar, Y. Korkmaz, E. Ceyani, B. Susam, M. Ozbey, S. Avestimehr, and T. Çukur, "Federated Learning of Generative Image Priors for MRI Reconstruction," *IEEE Transactions on Medical Imaging*, vol. 42, no. 7, pp. 1996–2009, Jul. 2023, conference Name: IEEE Transactions on Medical Imaging. [Online]. Available: https://ieeexplore.ieee.org/document/9943293

[102] Y. Shen, A. Sowmya, Y. Luo, X. Liang, D. Shen, and J. Ke, "A Federated Learning System for Histopathology Image Analysis With an Orchestral Stain-Normalization GAN," *IEEE Transactions on Medical Imaging*, vol. 42, no. 7, pp. 1969–1981, Jul. 2023, conference Name: IEEE Transactions on Medical Imaging. [Online]. Available: https://ieeexplore.ieee.org/document/9947057

[103] M. Abdel-Basset, H. Hawash, and M. Abouhawwash, "Collaborative Screening of COVID-19-like Disease from Multi-Institutional Radiographs: A Federated Learning Approach," *Mathematics*, vol. 10, no. 24, p. 4766, Jan. 2022, number: 24 Publisher: Multidisciplinary Digital Publishing Institute. [Online]. Available: https://www.mdpi.com/2227-7390/10/24/4766

[104] M. U. Alam and R. Rahmani, "FedSepsis: A Federated Multi-Modal Deep Learning-Based Internet of Medical Things Application for Early Detection of Sepsis from Electronic Health Records Using Raspberry Pi and Jetson Nano Devices," *Sensors*, vol. 23, no. 2, p. 970, Jan. 2023, number: 2 Publisher: Multidisciplinary Digital Publishing Institute. [Online]. Available: https://www.mdpi.com/1424-8220/23/2/970

[105] L. Zhao and J. Huang, "A distribution information sharing federated learning approach for medical image data," *Complex & Intelligent Systems*, vol. 9, no. 5, pp. 5625–5636, Oct. 2023. [Online]. Available: https://doi.org/10.1007/s40747-023-01035-1

[106] E. H. Lee, B. Kelly, E. Altinmakas, H. Dogan, M. Mohammadzadeh, E. Colak, S. Fu, O. Choudhury, U. Ratan, F. Kitamura, H. Chaves, J. Zheng, M. Said, E. Reis, J. Lim, P. Yokoo, C. Mitchell, G. Houshmand, M. Ghassemi, R. Killeen, W. Qiu, J. Hayden, F. Rafiee, C. Klochko, N. Bevins, F. Sazgara, S. S. Wong, M. Moseley, S. Halabi, and K. W. Yeom, "AI Models Close to your Chest: Robust Federated Learning Strategies for Multi-site CT," Apr. 2023, arXiv:2303.13567 [cs]. [Online]. Available: http://arxiv.org/abs/2303.13567

[107] L. Peng, N. Wang, N. Dvornek, X. Zhu, and X. Li, "FedNI: Federated Graph Learning With Network Inpainting for Population-Based Disease Prediction," *IEEE Transactions on Medical Imaging*, vol. 42, no. 7, pp. 2032–2043, Jul. 2023. [Online]. Available: https://ieeexplore.ieee.org/document/9815303

[108] J. Wang, G. Xie, Y. Huang, J. Lyu, F. Zheng, Y. Zheng, and Y. Jin, "FedMed-GAN: Federated domain translation on unsupervised cross-modality brain image synthesis," *Neurocomputing*, vol. 546, p. 126282, Aug. 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0925231223004058

[109] D. Peketi, V. Chalavadi, C. K. Mohan, and Y. W. Chen, "FLWGAN: Federated Learning with Wasserstein Generative Adversarial Network for Brain Tumor Segmentation," in *2023 International Joint Conference on Neural Networks (IJCNN)*, Jun. 2023, pp. 1–8, iSSN: 2161-4407. [Online]. Available: https://ieeexplore.ieee.org/document/10191202

[110] Y. L. Tun, C. M. Thwal, J. S. Yoon, S. M. Kang, C. Zhang, and C. S. Hong, "Federated Learning with Diffusion Models for Privacy-Sensitive Vision Tasks," in *2023 International Conference on Advanced Technologies for Communications (ATC)*, Oct. 2023, pp. 305–310, iSSN: 2162-1039. [Online]. Available: https://ieeexplore.ieee.org/document/10318858

[111] R. Jin and X. Li, "Backdoor attack and defense in federated generative adversarial network-based medical image synthesis," *Medical Image Analysis*, vol. 90, p. 102965, Dec. 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1361841523002256

[112] A. Abdul Sattar Shaikh, M. S. Bhargavi, and P. Kumar C, "Weighted aggregation through probability based ranking: An optimized federated learning architecture to classify respiratory diseases," *Computer Methods and Programs in Biomedicine*, vol. 242, p. 107821, Dec. 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S016926072300487X

[113] T. Sattarov, M. Schreyer, and D. Borth, "FedTabDiff: Federated Learning of Diffusion Probabilistic Models for Synthetic Mixed-Type Tabular Data Generation," Jan. 2024, arXiv:2401.06263 [cs]. [Online]. Available: http://arxiv.org/abs/2401.06263

[114] D. Kundu, M. M. Rahman, A. Rahman, D. Das, U. R. Siddiqi, M. G. R. Alam, S. K. Dey, G. Muhammad, and Z. Ali, "Federated Deep Learning for Monkeypox Disease Detection on GAN-Augmented Dataset," *IEEE Access*, vol. 12, pp. 32 819–32 829, 2024. [Online]. Available: https://ieeexplore.ieee.org/document/10445472

[115] A. Rehman, H. Xing, L. Feng, M. Hussain, N. Gulzar, M. Adnan Khan, A. Hussain, and D. Saeed, "FedCSCD-GAN: A secure and collaborative framework for clinical cancer diagnosis via optimized federated learning and GAN," *Biomedical Signal Processing and Control*, vol. 89, p. 105893, Mar. 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1746809423013265

[116] J. Chen and R. Pan, "Medical report generation based on multimodal federated learning," *Computerized Medical Imaging and Graphics*, vol. 113, p. 102342, Apr. 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0895611124000193

[117] M. Mazher, I. Razzak, A. Qayyum, M. Tanveer, S. Beier, T. Khan, and S. A. Niederer, "Self-supervised spatial–temporal transformer fusion based federated framework for 4D cardiovascular image segmentation," *Information Fusion*, vol. 106, p. 102256, Jun. 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1566253524000344

[118] Z. Li and B. Kainz, "Image Distillation for Safe Data Sharing in Histopathology," in *Medical Image Computing and Computer Assisted Intervention – MICCAI 2024*, M. G. Linguraru, Q. Dou, A. Feragen, S. Giannarou, B. Glocker, K. Lekadir, and J. A. Schnabel, Eds. Cham: Springer Nature Switzerland, 2024, pp. 459–469.

[119] S. Puppala, I. Hossain, M. J. Alam, and S. Talukder, "SCAN: A HealthCare Personalized ChatBot with Federated Learning Based GPT," Jul. 2024, pp. 1945–1951.

[120] M. Sikandar, I. U. Din, and A. Almogren, "Integrating Generative AI and Federated Learning for Privacy Preserved Sequence-Based Stomach Adenocarcinoma Detection," *IEEE Transactions on Consumer Electronics*, vol. 70, no. 3, pp. 5278–5285, Aug. 2024. [Online]. Available: https://ieeexplore.ieee.org/document/10587004

[121] A. Murmu, P. Kumar, N. Rao Moparthi, S. Namasudra, and P. Lorenz, "Reliable Federated Learning With GAN Model for Robust and Resilient Future Healthcare System," *IEEE Transactions on Network and Service Management*, vol. 21, no. 5, pp. 5335–5346, Oct. 2024. [Online]. Available: https://ieeexplore.ieee.org/document/10586897

[122] A. K. Selvaraj, S. B. Prathiba, A. Deepak Kumar, R. Dhanalakshmi, T. R. Gadekallu, and G. Srivastava, "Co-Training-Based Personalized Federated Learning With Generative Adversarial Networks for Enhanced Mobile Smart Healthcare Diagnosis," *IEEE Transactions on Consumer Electronics*, vol. 70, no. 3, pp. 6131–6139, Aug. 2024. [Online]. Available: https://ieeexplore.ieee.org/document/10679979

[123] S. Chen, T. Jin, Y. Xia, and X. Li, "Metadata and Image Features Co-Aware Semi-Supervised Vertical Federated Learning With Attention Mechanism," *IEEE Transactions on Vehicular Technology*, vol. 73, no. 2, pp. 2520–2532, Feb. 2024. [Online]. Available: https://ieeexplore.ieee.org/document/10246395

[124] N. Alalwan, A. Alwadain, A. I. Alzahrani, A. H. Al-Bayatti, A. Abozeid, and R. M. Abd El-Aziz, "Advancements in brain tumor identification: Integrating synthetic GANs with federated-CNNs in medical imaging analysis," *Alexandria Engineering Journal*, vol. 105, pp. 105–119, Oct. 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1110016824006872

[125] B. Yan, Q. Chen, Y. Chen, X. Jiang, W. Huang, B. Wang, Z. Wang, C. Gao, and T. Zhang, "Buffalo: Biomedical Vision-Language Understanding with Cross-Modal Prototype and Federated Foundation Model Collaboration," in *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, ser. CIKM '24. New York, NY, USA: Association for Computing Machinery, Oct. 2024, pp. 2775–2785. [Online]. Available: https://dl.acm.org/doi/10.1145/3627673.3679627

[126] F. Li, Y. Wang, O. Beyan, M. Schöneck, and L. L. Caldeira, "Voxel-Wise Medical Image Generalization for Eliminating Distribution Shift," *ACM Trans. Knowl. Discov. Data*, vol. 18, no. 7, pp. 167:1–167:16, Jun. 2024. [Online]. Available: https://dl.acm.org/doi/10.1145/3643034

[127] G. N. Nimeshika and S. D, "Enhancing Alzheimer's disease classification through split federated learning and GANs for imbalanced datasets," *PeerJ Computer Science*, vol. 10, p. e2459,

Nov. 2024, publisher: PeerJ Inc. [Online]. Available: https://peerj.com/articles/cs-2459

[128] T. Bdair, H. Saadeh, B. Qaqish, A. Sulaq, and M. Rawashdeh, "Medical Image-to-Image Translation with Spatial Self-Attention for Radiotherapy in Federated Learning," in *2024 Fifth International Conference on Intelligent Data Science Technologies and Applications (IDSTA)*, Sep. 2024, pp. 103–110. [Online]. Available: https://ieeexplore.ieee.org/document/10746985

[129] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, "LoRA: Low-Rank Adaptation of Large Language Models," Oct. 2021. [Online]. Available: https://openreview.net/forum?id=nZeVKeeFYf9

[130] X. Zhao and D. Shen, "FedSW: Federated learning with adaptive sample weights," *Information Sciences*, vol. 654, p. 119873, Jan. 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0020025523014585

[131] B. Theodorou, C. Xiao, and J. Sun, "Synthesize high-dimensional longitudinal electronic health records via hierarchical autoregressive language model," *Nature Communications*, vol. 14, no. 1, p. 5305, Aug. 2023, publisher: Nature Publishing Group. [Online]. Available: https://www.nature.com/articles/s41467-023-41093-0

[132] G. Vavoulas, C. Chatzaki, T. Malliotakis, M. Pediaditis, and M. Tsiknakis, "The MobiAct Dataset: Recognition of Activities of Daily Living using Smartphones," Jun. 2025, pp. 143–151. [Online]. Available: https://www.scitepress.org/Link.aspx?doi=10.5220/0005792401430151

[133] G. Vavoulas, M. Pediaditis, E. G. Spanakis, and M. Tsiknakis, "The MobiFall dataset: An initial evaluation of fall detection algorithms using smartphones," in *13th IEEE International Conference on BioInformatics and BioEngineering*, Nov. 2013, pp. 1–4. [Online]. Available: https://ieeexplore.ieee.org/document/6701629

[134] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-Resolution Image Synthesis with Latent Diffusion Models." IEEE Computer Society, Jun. 2022, pp. 10 674–10 685. [Online]. Available: https://www.computer.org/csdl/proceedings-article/cvpr/2022/694600k0674/1H1iFsO7Zuw

[135] A. Alabduljabbar, S. U. Khan, A. Alsuhaibani, F. Almarshad, and Y. N. Altherwy, "Medical imaging datasets, preparation, and availability for artificial intelligence in medical imaging," *Journal of Alzheimer's Disease Reports*, vol. 8, no. 1, pp. 1471–1483, Jan. 2024, publisher: SAGE Publications. [Online]. Available: https://doi.org/10.3233/ADR-240129

B Biography Section

Alessandro Ceresi He holds a degree in Physics and a Master's in Applied Physics from the University of Bologna (Alma Mater Studiorum). He is currently working in the SSR (Signals, Systems and Radiocommunications) department at UPM as a researcher. His work spans statistical data analysis and the application of artificial intelligence, particularly in the medical field.

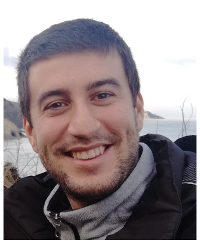

Borja Arroyo Galende holds a degree in Natural Environment Engineering from ETSIMFMN (UPM) and a degree in Computer Engineering from ETSI Informática (UNED), and a master's degree in Engineering and Data Science from ETSI Informática (UNED). He is currently developing his thesis in the context of machine learning in federated environments. He also works as a software developer/architect at GATV.

Javier Guinea-Pérez received his B.S. in telecommunication engineering from the Universidad Politécnica de Madrid (UPM), Madrid, Spain in 2018. In 2020 he obtained a double M.Sc. in Computer Science at Kungliga Tekniska Hogskolan, Stockholm, Sweden and the Delft University of Technology, Delft, Netherlands. He is currently working on his Ph.D. thesis in the SSR (Signals, Systems and Radiocommunications) department at UPM. His research interests include deep learning, medical imaging and biostatistics.

Patricia A. Apellániz received the Bachelor's degree in Telecommunication Engineering from the Universidad Autónoma de Madrid, Madrid, Spain, in 2018, and the Master's degree in Telecommunication Engineering from the Universidad Politécnica de Madrid (UPM), Madrid, in 2020. In 2025, she obtained a Ph.D. in deep learning algorithms for medical applications. She is currently a researcher at UPM. Her research interests include audio and image signal processing, with applications in various deep learning domains.

Gustavo Hernandez-Peñaloza (Member, IEEE) received the Telecom Engineering degree from Universidad Santo Tomás, in 2007, the Master of Science degree in telecommunication technologies, system and networks from the Universidad Politécnica de Valencia (UPV),in 2009, and the Master of Business Administration (M.B.A.) and Ph.D. degrees (cum laude) from the Universidad Politécnica de Madrid (UPM),in 2019. From 2010 to 2013, he was an Associate Research Fellow with the Universidad de Valencia (UV). He is currently working as a Postdoctoral Fellow with UPM. He has been involved in multiple national and European funded projects in technical and management activities. His main research interest includes artificial intelligence (AI) applied to healthcare..

Federico Álvarez García (M'07) received the Telecom Engineer degree (Hons.) in 2003 and Ph.D. degree (cum laude) in 2009, both by the Universidad Politécnica de Madrid (UPM) where he is currently working as assistant professor. From 2003 he is working for the research group in the Visual Telecommunications Applications group (GATV) of UPM. He has been participating with different managerial and technical responsibilities in several national and EU projects, being coordinator of 5 EU projects in last 6 years. Dr. Álvarez had participated in national and international standardization fora (DVB, CENELEC TC206, etc.); he is a member of the program committee of several scientific conferences and the author and co-author of 60+ papers and several books, book chapters and patents in the field of ICT networks and audiovisual technologies.