applied sciences

MDPI

Review

Efficient and Secure GANs: A Survey on Privacy-Preserving and Resource-Aware Models

Niovi Efthymia Apostolou [1], Elpida Vasiliki Balourdou [1], Maria Mouratidou [1], Eleni Tsalera [1,*], Ioannis Voyiatzis [1], Andreas Papadakis [2] and Maria Samarakou [1,*]

[1] Department of Informatics and Computer Engineering, School of Engineering, University of West Attica, 11521 Athens, Greece; ice23390085@uniwa.gr (N.E.A.); ice23390095@uniwa.gr (E.V.B.); ice23390183@uniwa.gr (M.M.); voyageri@uniwa.gr (I.V.)

[2] Department of Electrical and Electronic Engineering Educators, School of Pedagogical and Technological Education (ASPETE), 15122 Athens, Greece; apapadakis@aspete.gr

* Correspondence: etsalera@uniwa.gr (E.T.); marsam@uniwa.gr (M.S.)

Abstract

Generative Adversarial Networks (GANs) generate synthetic content to support applications such as data augmentation, image-to-image translation, and training models where data availability is limited. Nevertheless, their broader deployment is constrained by limitations in data availability, high computational and energy demands, as well as privacy and security concerns. These factors restrict their scalability and integration in real-world applications. This survey provides a systematic review of research aimed at addressing these challenges. Techniques such as few-shot learning, consistency regularization, and advanced data augmentation are examined to address data scarcity. Approaches designed to reduce computational and energy costs, including hardware-based acceleration and model optimization, are also considered. In addition, strategies to improve privacy and security, such as privacy-preserving GAN architectures and defense mechanisms against adversarial attacks, are analyzed. By organizing the literature into these thematic categories, the review highlights available solutions, their trade-offs, and remaining open issues. Our findings underline the growing role of GANs in artificial intelligence, while also emphasizing the importance of efficient, sustainable, and secure designs. This work not only concentrates the current knowledge but also sets the basis for future research.

Keywords: GANs; data limitations; energy—computational cost; privacy—security; deep learning; artificial intelligence

Academic Editor: Tobias Meisen

Received: 19 September 2025
Revised: 11 October 2025
Accepted: 16 October 2025
Published: 19 October 2025

Citation: Apostolou, N.E.; Balourdou, E.V.; Mouratidou, M.; Tsalera, E.; Voyiatzis, I.; Papadakis, A.; Samarakou, M. Efficient and Secure GANs: A Survey on Privacy-Preserving and Resource-Aware Models. *Appl. Sci.* **2025**, *15*, 11207. https://doi.org/10.3390/app152011207

1. Introduction

Deep Learning (DL) has emerged as the most effective and widely adopted approach due to its capability to model complex, high-dimensional data [1]. Generative Adversarial Networks (GANs), first introduced in 2014 [2], support generation of synthetic data required in machine learning (ML) applications. Before GANs, probabilistic methods like Variational Autoencoders (VAEs) and Restricted Boltzmann Machines offered useful but limited ways to learn data distributions. GANs stand out because of their unique training approach, which has expanded the potential for creating synthetic data. Unlike traditional statistical models, GANs learn complex data distributions through adversarial training, without needing explicit assumptions.

GANs consist of a pair of artificial neural networks that act competitively against each other: the generator (G) and the discriminator (D). They are trained in a minimax game where the generator creates forged copies of real data and the discriminator attempts to differentiate between original samples and the forgeries [3]. This adversarial framework drives both networks to improve in parallel, generating data of great fidelity in varied modalities such as audio, images, video, and text [4,5]. GANs have become essential in many scientific and industrial areas, ranging from computer vision [6,7], natural language processing [8], medical imaging [9], finance [10], and civil engineering [11]. Architectural improvements, such as Deep Convolutional GANs (DCGANs), Wasserstein GANs (WGANs), StyleGANs, and Progressive GANs, have been proposed to enhance stability, scalability, and output quality.

Despite these advancements, the literature shows significant gaps in three main areas: privacy preservation, computational and energy efficiency, and robustness to data scarcity [12,13]. Training advanced GANs often requires large computational resources, which leads to high energy use and carbon emissions [14]. In [15,16] it was reported that training large neural models on multi-GPU setups can consume between 100 and 500 kWh of electricity and emit up to 500 kg of CO_2 per training run. In [17] is reported that training StyleGAN2 for high-resolution image synthesis on 8 GPUs over two weeks consumes approximately 250–300 kWh of electricity. In addition, GANs trained on sensitive information, such as medical images, biomedical identifiers, or financial records, create significant privacy and security risks. Important ethical and security issues are posed, like deepfakes, identity theft, and unauthorized data use [18]. Several surveys examine GANs from complementary perspectives [13]. A thorough study in privacy and security [19] provides a GAN-based attacks/defenses taxonomy and map their use across application domains (medical imaging, network security) focusing on specific implementations. Also, [20] is a comprehensive study of privacy and utility metrics across GANs and VAEs, and attack types, with a focus on metric selection, evaluation protocols, and guidance on how to choose the appropriate measures. In contrast, this survey explicitly integrates three themes, analyzing their mutual trade-offs and deployment implications:

- Data Scarcity—We explore strategies for GANs to work well with restricted training datasets. This includes data augmentation techniques and specialized architectures aimed at low-data situations [21,22].
- Computational and Energy Efficiency—We look at ways to lower the training costs and carbon footprint of GANs. This aspect is crucial yet often overlooked in resource-aware AI [16,17].
- Privacy-Preserving Mechanisms—We review methods such as differential privacy, federated GANs, and cryptographic frameworks. These enhance security while reducing the risk of misuse [23,24].

Whereas previous studies generally focus on optimizing a single aspect (privacy, metrics, or applications), our review combines approaches that simultaneously tackle resource limitations and privacy concerns, recommending evaluation priorities (such as energy/carbon accounting) essential for practical implementation.

We combined trade-offs maps to illustrate how privacy methods (e.g., DCGAN) impact computational expenses and model accuracy, and similarly how compression/acceleration affects memorization and vulnerability. Subsequently, we aim to incorporate energy and carbon metrics into GAN assessment, moving beyond conventional evaluation criteria to include sustainability and privacy as key axes. This addresses a gap in prior studies, which have discussed metrics without considering energy, or explored applications and attacks without considering sustainability. Moreover, we examine which techniques can adapt to practical scenarios (small medical groups, non-image tabular information, mobile/IoT), where previous studies pay less attention to feasibility and scaling

limitations, or emphasize in metrics and effective measurement practices. Our goal is to suggest concrete design patterns for deploying GANs under practical constraints. By bringing together recent developments and highlighting promising research directions, this survey offers both a practical guide for real-world deployment and as a framework for addressing the ethical, computational, and methodological challenges of current models. To the best of our knowledge, it is the first survey to integrate efficiency and privacy considerations into a unified view of GANs.

Finally, it is important to recognize that the impact of GAN research extends beyond adversarial training itself. Insights gained from privacy-preserving and energy-efficient GANs increasingly inform adjacent areas such as diffusion models, large-scale language models, and hybrid generative architectures.

The rest of this paper is organized as follows: Section 2 provides an analysis on the functionality of GANs and the background, while Section 3 reviews related work on the evolution of GANs, their applications, limitations and unexplored research areas. Section 4 presents the methodology and materials used for this survey research. Section 5 analyzes strategies to address data scarcity, Section 6 examines methods for computational and energy efficiency, while Section 7 explores approaches for protecting privacy and security. Section 8 outlines the conclusions and limitations of this study and Section 9 provides directions for future research.

2. Background

A Generative Adversarial Network consists of two neural networks trained in opposition: the generator and the discriminator. The discriminator learns to distinguish real samples $x \sim p_{data}(x)$ from synthetic samples $\tilde{x} = G(z)$ generated by the generator, where $z \sim p_z(z)$ is drawn from a prior distribution such as a uniform or normal (Gaussian) distribution that defines the latent space from which the generator draws its input noise vectors. Formally, the discriminator estimates the probability $D(x)$ that a given sample is real, while the generator is trained to produce realistic outputs that minimize the discriminator's ability to identify them as fake. $V(D,G)$ represents the function of the adversarial objective between the generator and the discriminator, while $E_{x\sim p(x)}[\cdot]$ denotes the expected value over samples drawn from each distribution.

The equation of the competitive training of the D and the G is given by [2] as follows:

$$\min_G \max_D V(D,G) = \mathrm{E}_{x\sim p_{data}(x)}[\log D(x)] + \mathrm{E}_{z\sim p_z(z)}[\log(1 - D(G(z)))] \quad (1)$$

The discriminator aims to maximize the probability $V(D,G)$ by correctly distinguishing real and generated data, whereas the generator minimizes it by providing samples that push $D(G(z))$ towards one [2]. Figure 1 illustrates the standard GAN framework, highlighting the core adversarial dynamics. The visualization represents the interaction between the generator and the discriminator as described above, emphasizing the simultaneous adversarial training of both networks.

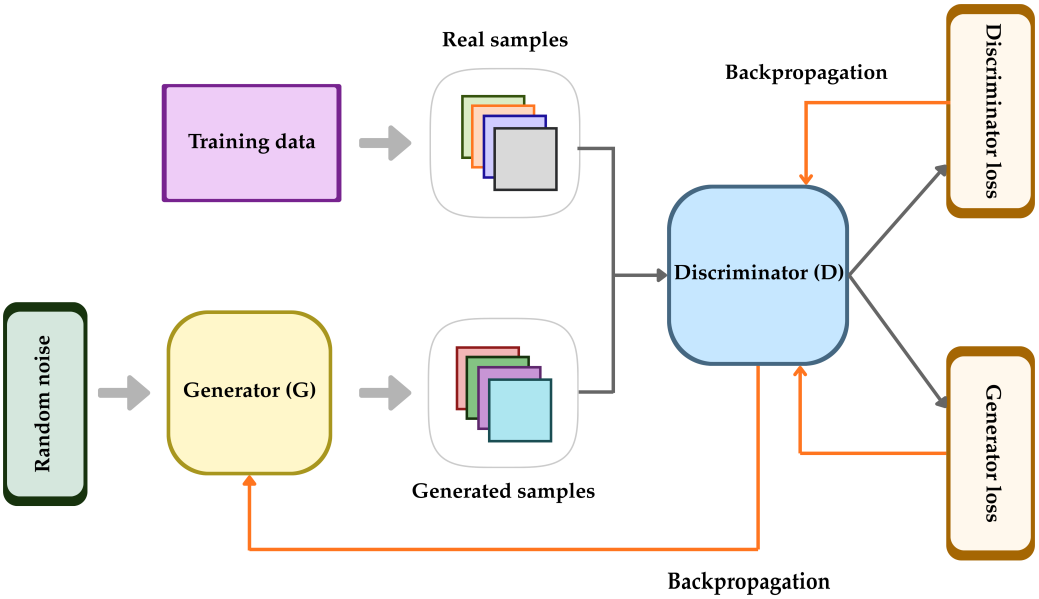

Figure 1. Schematic overview of Generative Adversarial Network (GAN). The generator transforms random noise into generated samples, which are evaluated by the discriminator against real data. Both networks are trained and updated through backpropagation according to their respective loss functions. Gray arrows visualize the forward propagation while orange arrows indicate backpropagation.

Early GAN architectures faced significant challenges, including training instability and mode collapse, and low image quality, which motivated a series of refinements between 2014 and 2016. Conditional GANs (CGANs) were proposed to control the generation process using auxiliary information [25], while DCGANs improved stability and scalability through convolutional architectures [3]. Additionally, LAPGAN (a conditional form of GAN integrated into the framework of a Laplacian pyramid) was designed to enable high-resolution image generation, addressing one of the key limitations early GAN models encountered [26]. The issue of instability continued to drive new designs from 2017 to 2018. The WGAN introduced a new distance metric that helped reduce model collapse and improved convergence [27]. WGAN-GP variant added a gradient penalty regularization, further stabilizing the training process. New evaluation metrics like Inception Score (IS) [28] and Fréchet Inception Distance (FID) [29] provided standardized benchmarks to assess generative quality.

From 2017 onward, GANs were rapidly adopted for image synthesis and translation. Pix2Pix [30] demonstrated supervised image-to-image translation using paired datasets, while CycleGAN [31] extended the method to unpaired data. StyleGAN [7] and its successors, StyleGAN2 [32] and StyleGAN3 [33], made significant advances by generating high-resolution, photorealistic images with remarkable fidelity.

Figure 2 illustrates the chronological evolutions of important GAN architectures. Each model represents a significant improvement resolving distinct limitations such as stability, scalability, mode robustness, application and quality.

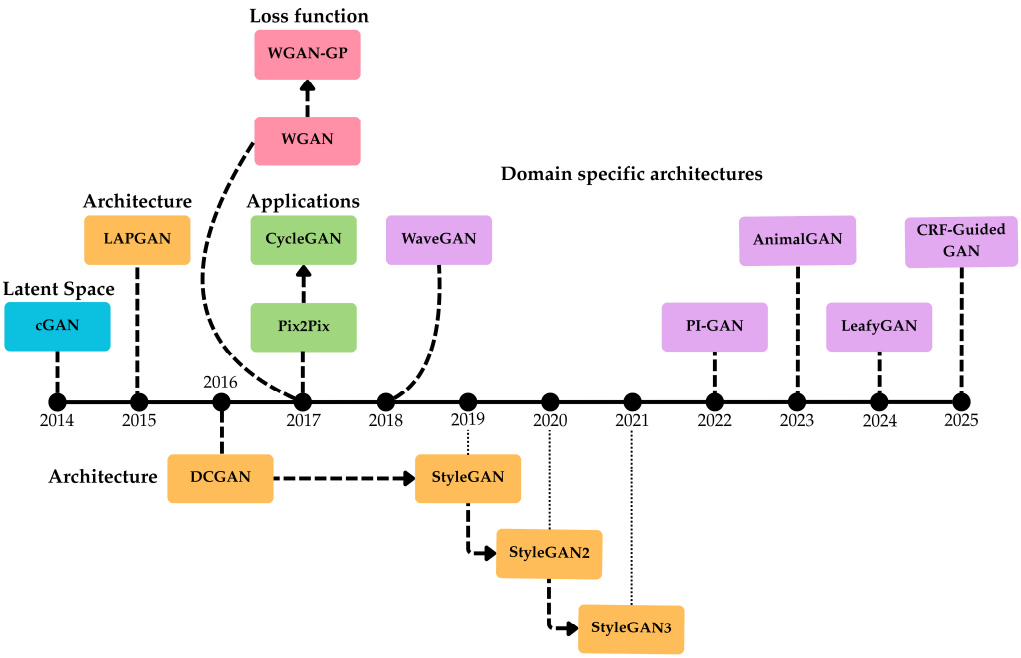

Figure 2. Evolution of fundamental GAN architectures from 2014 to 2025.

In addition to current algorithmic and hardware enhancements, recent advancements in control theory indicate alternative methods for stabilizing and accelerating GAN training. Models like state-filtered disturbance rejection control (SFDRC) [34] introduce dynamic state filters that target and mitigate high-frequency disturbances, similar to removing stochastic gradient noise in adversarial learning. Thus, incorporating SFDRC-based filtering enhances convergence stability and energy efficiency.

Similarly, multilayer neurocontrol systems with active disturbance rejection control (ADRC) [35] address high-order uncertain systems by combining neural approximations with adaptive disturbance compensation. This model resembles the generator-discriminator interaction: the controller (discriminator) assesses and mitigates external disturbances (training instabilities) as the plant (generator) acquires the target distribution. Integrating ADRC principles into GAN optimization may produce adaptive gradient controllers capable of autonomously modifying learning rates, filtering out stochastic noise, and ensuring convergence despite data shortages or adversarial disturbances.

Over the last five years, research has increasingly shifted toward domain-specific applications, where GANs are designed for specialized tasks. In the medical imaging domain, Conditional Random Field (CRF)-Guided GAN is a recent approach that seeks to produce 3D high-resolution medical image synthesis with reduced memory requirements, thereby improving both image fidelity and computational efficiency [36]. Several GAN variants have also been developed for agriculture. Leafy-GAN aids plant disease detection, AnimalGAN produces synthetic animal-testing data, and Plant Identification GAN (PI-GAN) supports plant classification. These models contribute to sustainable farming by enabling early disease detection, minimizing animal testing, and improving crop monitoring [37,38].

In the field of audio synthesis, GANs have shown considerable performance. WaveGAN generates waveforms trained on speech datasets, allowing the production of sound effects, bird vocalizations, and musical instruments, broadening GAN applications into acoustics [39].

Despite these advances, challenges related to data limitations, computational costs, and privacy concerns, which are critically examined in Section 3, still persist.

3. Related Work and Problem Statement

Despite the rapid advancements, several research areas remain insufficiently explored. Non-visual domain, such as electroencephalogram (EEG) signal generation [40] and alphanumeric character recognition [41], and time-series analysis [42], have received limited attention, due to data scarcity and high labeling costs. Sequential and time-series domains often depend on binary adversarial feedback for learning, which provides insufficient signals for the network to accurately reflect the temporal dynamics in the training data [43]. In graph-based data, face inherent challenges, as standard embedding algorithms fail to preserve a graph's topological structure and differentiate heterogeneous nodes and edges [44]. The computational and energy costs of training large-scale GANs, along with data scarcity constrain many potential applications [13]. Structured latent spaces and disentangled representations have been proposed to improve controllability [18,45]. However, they often add extra training complexity and do not solve the basic efficiency problems. Moreover, the "black-box" nature of GANs restricts interpretability of the learnt features [46], and complicates secure data management. Current evaluation metrics poorly capture perceptual and semantic fidelity. This leads to limited trust and reproducibility in GAN-generated data. In summary, while GANs have matured significantly in image synthesis and domain-specific applications, their underexplored non-visual areas and latent space structures represent promising paths for future research. Beyond image synthesis, GANs have shown strong potential across diverse scientific domains. In the following, we review the related works aligned with the three main dimensions of our research, introducing the problem statement for each case.

3.1. Data Scarcity

A crucial aspect of successful GAN training is the use of extensive, high-quality datasets [47,48]. Modern and high-quality GANs require about 10^5 to 10^6 images to operate properly [49]. Therefore, limited datasets interfere with the training process of the neural networks [49–52]. This difficulty occurs as the generator attempts to approximate the real data distribution while avoiding memorization of the provided data [53]. The training process of both the discriminator and the generator is unstable, implying that the discriminator can determine whether the data are authentic, whereas the generator does not receive sufficient training, resulting in the creation of repetitive or low-diversity outputs [54]. In many scientific fields, such as biology and plant science, providing samples is impractical due to significant costs and the scarcity of biological entities [55]. In medicine, this issue seems important, as GANs offer limited utility for rare medical conditions [56]. Similarly, specific languages and dialects lack sufficient data for neural networks, leading to often unreliable sound analysis [57,58].

Data augmentation is a prominent method to effectively mitigate overfitting [59]. However, the use of augmented data appears to imbalance the training process of GANs [60]. This is justified by the fact that a GAN trained on enhanced data would only learn and repeat the augmentations, not the distribution of the original data. Undesirable augmentations in the generated samples compromise the accuracy of the outputs. For example, if noise is introduced into the training samples to augment them, the generator may encode these augmented characteristics, embedding noise into its outputs regardless of the original data [61]. Developing approaches that preserve the benefits of augmentation without contaminating the generative process remains an important direction for improving the reliability and applicability of GANs.

Earlier approaches mainly relied on data augmentations [62], transfer learning, and advanced machine learning techniques aiming to improve the efficiency of the model under challenging conditions [63]. More recent and advanced studies explore the range of older methods, such as geometric data augmentation [64] and improved transfer learning

techniques. New mechanisms have been developed introducing consistency regularization [65], meta learning, metric learning and few-shot classifications to stabilize the performance of GAN training under few-shot conditions [66]. Despite these advances, most investigations review each approach individually with narrow experimental contexts often to specific architectures. A structured evaluation is presented in Section 5.

3.2. Energy Consumption

Another research concern is the computational and energetical demands of GANs. The complexity of the adversarial training has motivated the development of lightweight architectures and hardware–based approaches such as memristor accelerators [67] and Field-Programmable Gate Array (FPGA) [68]. Currently, several surveys attempt to accurately measure the environmental impact of GANs while also researching the possible causes. The research towards computational efficiency and energy consumption is limited, while existing surveys typically treat algorithmic and hardware methods separately. The requirement for solutions is addressed in Section 6 while a comparative analysis is included for the most efficient approach. The reduction of CO_2 emissions has not been sufficiently reviewed nor researched; therefore, this work does not provide a possible solution.

The increasing use of high-demand AI technologies has led to substantial rises in energy consumption and computational costs [16]. Deep generative models focus on improving quality and accuracy, resulting in significant energy use and CO_2 emissions [69]. Although, convolutional neural networks (CNNs) have been used in several documented techniques to address this problem [70], their direct application is limited by two factors: (i) the adversarial two-network training dynamics require adaptive data flow methods, such as DiffBlock (described in Section 6.1), for efficiency, and (ii) GANs introduce unique optimization challenges due to their complex loss functions [71–73].

A widely discussed factor of energy inefficiencies is the dual-network architecture of GAN models. Both generator and discriminator run forward and backward propagation while constantly exchanging data. This raises energy demands and increases memory traffic [74]. Frequent data exchanges between memory and processors increase energy use and slow down training in contrast to single-network deep learning models [75]. Energy estimation is essential for avoiding significant financial losses and reducing the carbon footprint [76].

Environmental impact factors include (i) the training procedure's duration, (ii) the geographical location of the server, (iii) the electricity grid it uses and (iv) the type and model of hardware being used for the training [77]. GAN training functions by consuming increasing amounts of energy while also intensifying the emissions of carbon dioxide [78,79]. The authors in [15] revealed that estimates of energy and CO_2 emissions vary widely based on methodology (Figure 3). Partial proxies (e.g., GPU-hours × TDP) disregard CPU and memory usage, as well as temporal fluctuations in power, while coarse regional or national averages neglect variations in grid carbon intensity. In contrast, real-time monitoring that integrates wall-plug power with grid carbon data provides more reliable results. Without such comprehensive accounting, energy use and emissions are easily miscalculated. Figure 3 compares various methods for estimating CO_2 emissions and energy consumption during GAN training. It illustrates how methodological choices influence the reported values. The left panel depicts CO_2 emissions, while the right panel shows the corresponding energy consumption.

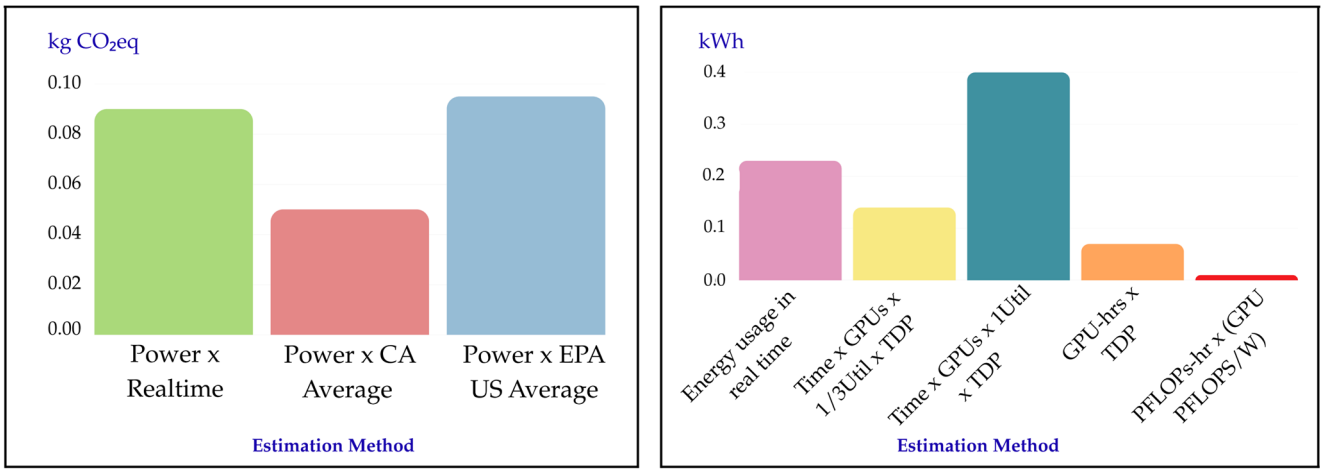

Figure 3. In the left chart, the assessment of carbon emissions is compared when using distinct estimation methods. Similarly, on the right chart the energy consumed in kWh is estimated with five different methods. In both charts, each method yielded substantially different results, which is entirely expected given regional variations in grid carbon intensity and methodological assumptions. This serves as an illustration that relying on limited information could result in an exaggeration or underestimation of carbon emissions. This research was conducted in [15].

For the estimation of CO_2 emissions (left panel) three approaches are considered. Power × Realtime calculates the total kilograms of CO_2 equivalent (kg CO_2eq) by integrating the measured wall-plug power P(t) over time and converting each period using the time-resolved grid carbon intensity c(t). The Power × CA Average approach multiplies the total measured energy by the California average carbon factor, whereas Power × EPA US Average multiplies the total measured energy by the U.S. national average carbon factor.

In terms of energy consumption (right panel), several methods are employed. The method proposed by [15] measures system power in real time (GPU + CPU + system) and converts it to kWh. A second proxy method uses wall-clock runtime, assuming 33% GPU utilization, ignoring CPU/memory and temporal variation. A similar variant assumes full 100% GPU utilization. Another proxy multiplies accumulated GPU-hours (time × utilization) by TDP. Finally, the PFLOPs-hours/(GPU PFLOPs/W) method calculates volume (PFLOPs-hours) and applies a performance-per-watt ratio to estimate energy.

These inconsistencies emphasize that the reported energy and carbon metrics should not remain purely descriptive but also enhance the development of future models. Accurately measuring energy consumption and carbon emissions can guide trade-offs between the complexity of architecture, training duration, and deployment feasibility. For example, models with lower energy requirements might more effectively balance accuracy, responsiveness, and sustainability in real-world applications. Establishing standardized reporting practices is essential to ensure that such metrics consistently influence design decisions.

In addition to training, on-device inference poses additional constraints. Interactive applications based on GANs (such as mobile image editing, VR/AR headsets, or voice synthesis) require low latency but are constrained on memory, processing and battery life [80,81]. These limitations hinder practical use, highlighting the need for solutions that lower energy consumption without sacrificing fidelity or responsiveness. The complexity of the adversarial training has motivated the development of lightweight architectures and hardware–based approaches such as memristor accelerators [67] and FPGA [68]. Currently, several surveys attempt to accurately measure the environmental impact of GANs while also researching the possible causes. The research towards computational efficiency and energy consumption is minimal, while existing surveys typically treat algorithmic

and hardware methods separately. The requirement for solutions is addressed in Section 6, while a comparative analysis is included for the most efficient approach. The reduction of CO_2 emissions has not been sufficiently reviewed nor researched; therefore, this work does not provide a possible solution.

3.3. *Data Privacy and Security Vulnerabilities*

GANs are a major target for privacy attacks and security breaches [82]. Sensitive client data from facial recognition systems and medical analysis models are the main targets of cyber violations [19]. Privacy of neural networks is at risk from Membership Inference Attacks (MIAs) [83]. MIAs focus on the information stored in the model's memory by exploiting overfitting, a phenomenon that occurs when the generated data closely resembles the real training data [84]. Confidential data is revealed and compromised, making the issue especially significant in sensitive sectors such as healthcare, power systems, and finance services [85,86]. Various types of privacy assaults include model inversion and data reconstruction attacks. Model inversion attacks leak training-data details by reconstructing realistic inputs from a model's outputs enabling the attacker to replicate the original data [87]. Medical data and biometrics, such as facial features and fingerprints, which are frequently used for personally identifiable information (PII), are the primary targets of model inversion attacks [88]. Data reconstruction attacks attempt to recover the actual data used for training rather than attempting to synthesize input data [89]. The objective of both privacy attacks is the acquisition of sensitive information from the input data.

Figure 4 illustrates a model inversion (MI) attack where several Models Under Attack (MUAs) assist in generating realistic data samples. A generator creates candidate images from noise vectors and target class identifiers (ID), refining them iteratively until MUAs classify them as genuine. When an auxiliary dataset with similar distribution is available, a discriminator further improves realism. Even without such data, the attack remains viable in a data-free setup, demonstrating its robustness. This figure emphasizes how adversaries can reconstruct private data without having access to the training inputs.

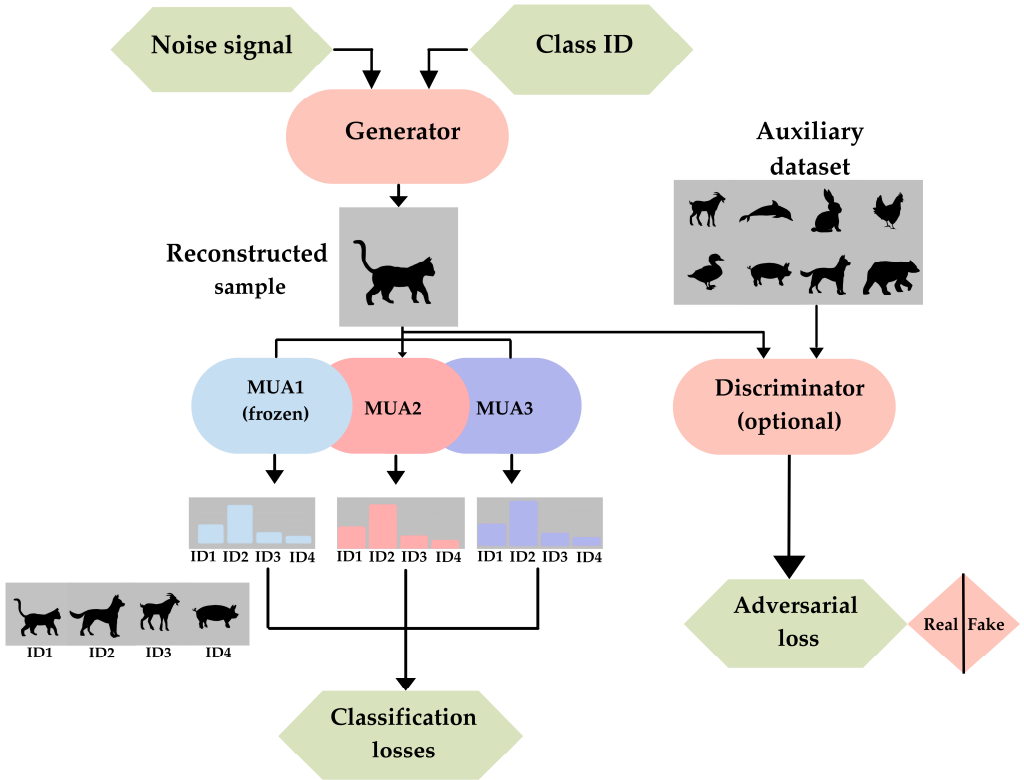

Figure 4. Mechanism of Model Inversion attack used to reconstruct sensitive training data. The generator creates candidate images form noise vectors and target class IDs, which are iteratively refined until Models Under Attacks classify them as realistic. When available, an auxiliary dataset enhances realism via a discriminator; otherwise, the attack proceeds in a fully free-data setting, demonstrating its effectiveness.

GANs may also be exploited for Distributed Denial of Service (DDoS) attacks, posing significant cybersecurity concerns [90]. These attacks typically overwhelm transport and application layers by flooding networks with excessive traffic, ultimately causing host crashes and service outages [22]. Intrusion Detection Systems (IDS) are commonly used to mitigate these attacks by filtering out malicious traffic. However, IDS are less effective against GAN-generated adversarial traffic, which includes polymorphic DDoS attacks and WGAN -based frameworks [91]. Polymorphic adversarial attacks constantly change their signatures to avoid detection, whereas [92] demonstrated that WGAN-based frameworks can generate functionally valid yet discreet adversarial traffic by altering non-functional features. In these frameworks, the discriminator mimics IDS behavior, providing feedback to the generator to improve evasion.

Multiple approaches have been presented to prevent data leakage, including privacy-preserving objectives and distributed learning frameworks [93]. Research on this area explores the development of secure but accurate models while training on sensitive datasets [94]. At the same time, the balance between maintaining the model's accuracy and quality while providing a privacy aware GAN remains an open research question. Section 7.1. provides promising solutions that aim to develop efficient and privacy-aware models.

4. Methodology and Materials

We employ a systematic literature review (SLR) approach based on PRISMA [95] guidelines and adapted from [96] to provide a structured and reproducible analysis of research on Generative Adversarial Networks (GANs) with a focus on efficiency and privacy-preserving methods. The review procedure comprises four main stages: literature identification, screening, categorization, and synthesis.

For literature retrieval, we implemented a multi-stage search strategy across several databases to ensure comprehensive coverage. Peer-reviewed journals and conference papers were sourced from IEEE Xplore (Institute of Electrical and Electronics Engineers, New York, NY, USA), ACM Digital Library (Association for Computing Machinery, New York, NY, USA), SpringerLink (Springer Nature, Berlin, Germany), and Elsevier ScienceDirect (Elsevier B.V., Amsterdam, The Netherlands), while Google Scholar (Google LLC, Mountain View, CA, USA) was consulted for broader access to preprints, theses, and cross-disciplinary work. Also, arXiv (Cornell University, Ithaca, NY, USA) was used to include the latest high-impact preprints. The search combined preliminary scoping with general terms, refined keyword combinations using Boolean operators, cited reference searching from key studies, and filtering by publication type, language, and year. The overall search covered the period 2014–2025, while the identification of solutions focused on publications from 2023 onwards to highlight the most recent advances. Earlier influential works are also included when foundational. Keywords included phrases such as "GAN effectiveness", "energy-conscious GANs", "privacy-protecting GANs", "differentially private GAN", and "GAN reduction".

During screening and eligibility assessment, the initial pool of publications was refined using predefined inclusion and exclusion criteria. We considered peer-reviewed journal articles, conference proceedings, and highly cited preprints that introduced innovative methodological approaches, addressed at least one of the survey's main themes (data scarcity, computational efficiency, or privacy and security), and provided

quantitative results, comparative analyses, or formal evaluation metrics. Duplicates, papers not in English, and studies focused solely on applications without methodological contributions were excluded. The initial search yielded 237 records, of which 225 were fully evaluated after abstract review and duplicate removal, ultimately incorporating 143 studies into the synthesis.

In the categorization stage, studies were organized into three primary themes: first, data scarcity, encompassing methods for limited training sets such as data augmentation, transfer learning, and few-shot learning; second, computational and energy efficiency, including approaches such as model compression, pruning, and hardware acceleration; and third, privacy and security, covering differential privacy, federated GANs, cryptographic frameworks, and adversarial defense mechanisms.

During synthesis and analysis, methods within each category were evaluated in terms of accuracy, scalability, energy footprint, and privacy guarantees. Methodological trade-offs were identified, such as privacy versus utility or efficiency versus stability, and research gaps were highlighted to pinpoint areas in need of further exploration.

This structured methodology, illustrated in Figure 5 ensures that the survey not only summarizes the state of the art but also provides a framework for developing efficient, privacy-preserving, and resource-aware GAN models.

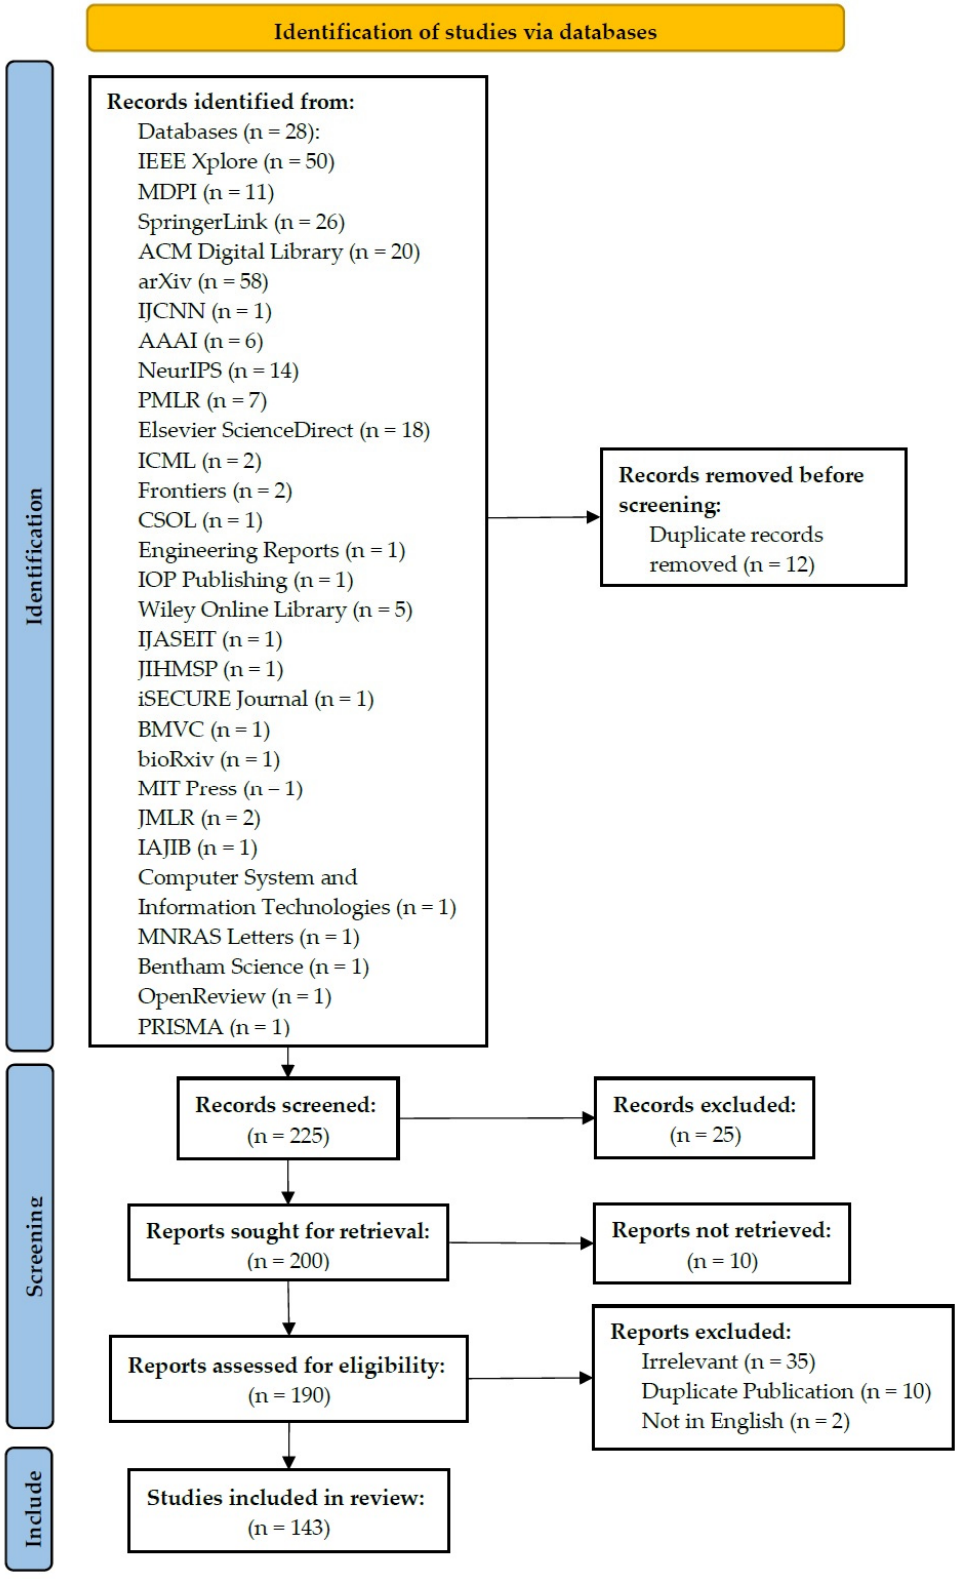

Figure 5. Prisma flow diagram.

5. Data Limitations

5.1. Solution Strategies

To address this issue, we identified and assembled the most impactful studies proposed in the literature. Adaptive Discriminator Augmentation (ADA) [97] applies augmentations to both real and generated data. Typical transformations of data include rotating, blurring, cropping, color jittering, and pitch shifting, ensuring that the discriminator

encounters only augmented data [98]. This method of training neural networks reduces training instabilities. However, data augmentation cannot replace real, high-quality data; thus, assembling a large and diverse training dataset remains essential. Augmentation is best used to address residual gaps in the data distribution [49]. Moreover, applying noise-based augmentation may lead to inherently noisy outputs [99].

Furthermore, employing pretrained GAN models with Adaptive Filter Modulation (AdaFM) [100] can lower the risk of overfitting in data-scarce scenarios [101,102]. Even with limited datasets, this technique removes the possibility of overfitting by enabling the GAN to make use of previously acquired information. Although, such approaches still face limitations when the target samples are restricted between 25 and 1000 samples [52]. In AdaFM, low-level filters capture the general features, while high-level filters encode domain-specific structures. Although low-level filters can be readily applied across various domains, high-level filters exhibit less transferability, potentially hindering performance when adjusting to new domains.

Ensuring semantic congruence between the generated and the original data is another challenge for GAN models. Consistency regularization addresses this issue by constraining GANs to produce outputs that preserve the essential meaning and recognizable structure of the data despite input variations [103]. Prior to the implementation of consistency regularization, semantic alignment was conducted on the modified data instead of the unprocessed versions [104]. After applying consistent regularization, the GAN model is constrained to preserve stable interpretations, ensuring that alterations do not affect the semantics of the data. By keeping the processed images close to their originals, the model learns to maintain the correct classification [105]. Figure 6 illustrates how consistency regularization aids the model to recognize that the essential meaning of an image remains unchanged despite minor modifications.

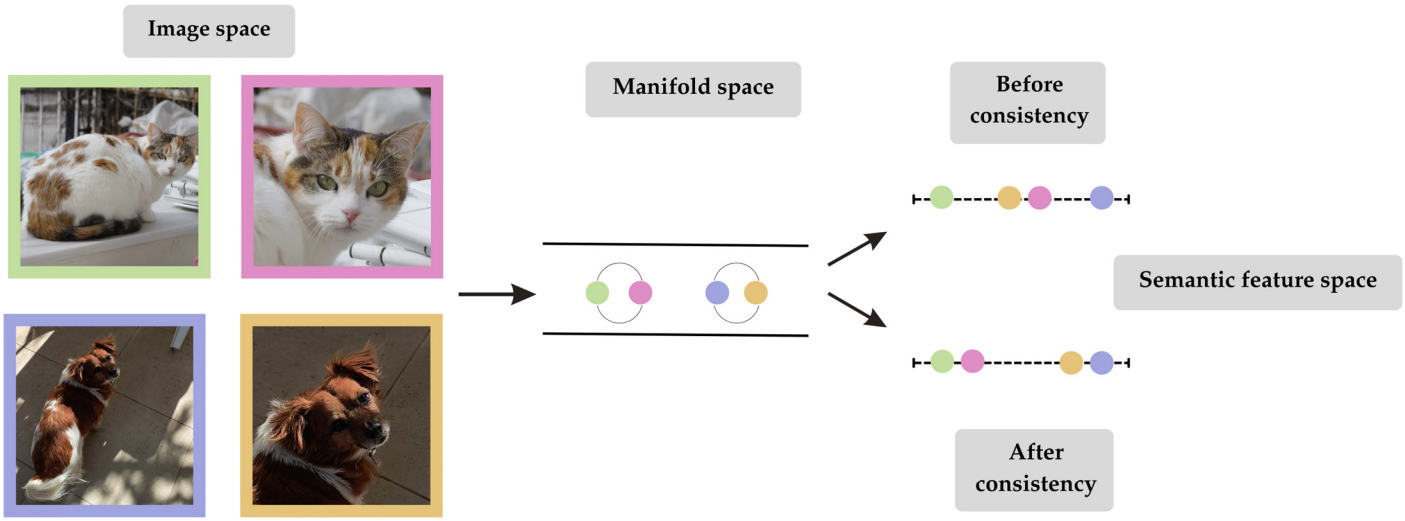

Figure 6. The improvement induced by consistency regularization is illustrated in this figure. Each picture in the model's image space is depicted by a colored dot. Images with similar semantic meaning are expected to be geometrically closer in the semantic feature space. Semantic similarity is not well captured before consistency regularization is applied, as the processed images of the dog and cat appear closer to each other than the two images of the same object. By bringing the augmented images (pink and yellow dots) closer to their originals, regularization ensures semantic consistency.

Despite the effectiveness of consistency regularization in GAN training, several limitations have been recognized. As augmentations are applied only to the actual data, the training process becomes imbalanced, and the regularization only affects the discriminator, thereby limiting the overall performance. The generator may internalize artificial

augmentation features and embed them into the synthetic samples as undesired artifacts. Finally, consistency is restricted in the image space, ignoring the latent space where further improvements could be achieved [60].

In addition to regularization-based approaches, few-shot learning (FSL) is an extensively discussed technique to prevent imbalances in GAN training under limited data availability [106,107]. Although consistency regularization focuses on semantic stability, FSL approaches offer a broader framework for training GANs with limited data. FSL includes three main categories: (i) Meta-learning [108], where the model trains on small tasks and learns to generalize this "learning-to-learn" skill to new tasks [109,110]. (ii) Transfer learning, which modifies pretrained GANs on large datasets for new tasks through fine-tuning, facilitating the application of previously gained knowledge [111,112]. (iii) Metric learning is a strategy that enables the GAN to compare samples within a learnt feature space, promoting the identification of differences and preventing overfitting [113,114]. In FSL, extensive research has evaluated classifiers using the Omniglot dataset, which consists of 1623 classes of handwritten characters spanning 50 alphabets, with only 20 samples per class [115,116]. According to [117], the most powerful few-shot classifiers are Memory-Augmented Neural Network (MANN) [118], Convolutional Siamese Nets [119], Matching Nets [120], Siamese Nets with Memory [121], Neural Statistician [122], Meta Nets [123], Prototypical Nets [124], Model-Agnostic Meta-Learning (MAML) [125], and RELATION NET [117]. Their performance on 1-shot and 5-shot classification tasks (5-way and 20-way) is summarized in Table 1, which provides a comparative overview of few-shot classifiers employed in research to mitigate data scarcity. These results highlight performance differences among models.

Table 1. This table displays the classification accuracy of few-shot classifiers used to resolve the setback of data limitations. The column Fine Tune determines whether the GAN architecture is trained by solely the same data (N) or by using the new samples from few-shot task (Y). The two categories 5-way Acc. and 20-way Acc. refer to the quantity of classes used in each classification problem while '-' is not reported. 1-shot and 5-shot determine the number of samples every GAN model is trained on. The symbol ± indicates the standard deviation that was calculated. The outcome of this experiment indicates that Relation Net consistently achieves the highest performance among all four categories. Statistics were computed in study [117].

Models	Fine Tune	5-Way Acc.		20-Way Acc.	
		1-shot	5-shot	1-shot	5-shot
MANN	N	82.8%	94.9%	-	-
Convolutional Siamese Nets	N	96.7%	98.4%	88.1%	96.5%
Convolutional Siamese Nets	Y	97.3%	98.4%	88.1%	97.0%
Matching Nets	N	98.1%	98.9%	93.8%	98.5%
Matching Nets	Y	97.9%	98.7%	93.5%	98.7%
Siamese Nets with Memory	N	98.4%	99.6%	95.0%	98.6%
Neural Statistician	N	98.1%	99.5%	93.2%	98.1%
Meta Nets	N	99.0%	-	97.0%	-
Prototypical Nets	N	98.8%	99.7%	96.0%	98.9%
MAML	Y	98.7 ± 0.4%	99.9%	95.8 ± 0.3%	98.9 ± 0.2%
RELATION NET	N	99.6 ± 0.2%	99.8%	97.6 ± 0.2%	99.1± 0.1%

The results of this experiment reveal notable differences in classification accuracy among the evaluated modes. As shown in Table 1, Relation Net consistently outperforms all other models, achieving accuracies above 97% in both 1-shot and 5-shot tasks for 5-way and 20-way classification. Its main advantage is its stability, especially in the challenging 1-shot/20-way setting, where other few-shot learning models often show considerable

performance declines. In GAN frameworks, few-shot classifiers are usually integrated into GAN models by enhancing the discriminator with an auxiliary classification head that is trained on the support set. A static few-shot encoder is also used to provide consistency regularization. Therefore, the training becomes stable when data is limited. Such methods, however, continue to depend on pre-trained models and large-scale datasets (e.g., FFHQ with 70,000 images), leading to considerable computational cost and substantial resource requirements. The performance of these models significantly decreases under large domain shifts between auxiliary and target datasets, as the learned representations fail to transfer effectively. Several studies have attempted to address the limitations using simple augmentation techniques. The original data distribution is altered, misleading the generator and reducing the robustness and generalization of FSL approaches [106].

In addition to hybrid FSL-GAN approaches, several architectures have been developed to encounter the challenge of small datasets [126]. These consist of Residual Weight Masking Conditional GAN (RWM-CGAN) [127], Dynamic GAN (DYNAGAN) [128], Gaussian-Poisson GAN (GP-GAN) [52], Diverse and Limited data GAN (DeLiGAN) [21] and Frequency-aware GAN (FreGAN) [129].

5.2. Critical Analysis

In Section 5.1., the following solutions were examined in detail: (i) Adaptive Discriminator Augmentation, (ii) AdaFM/pre-trained GANs, (iii) Consistency regularization, and (iv) Few-Shot Learning. ADA offers a practical and lightweight yet limited solution that is effective primarily for moderately sized datasets. However, when applied to very small datasets the artifacts may prevail, reducing fidelity. Despite its cost, the adaptation of pre-trained GANs with AdaFM is generally an efficient and balanced solution, especially suitable for small to medium datasets, as it leverages prior knowledge. Consistency regularization is effective as a supportive technique but insufficient as a standalone solution for limited data training. Few-Shot Learning is the most discussed technique regarding small data training. Theoretically, it is the most powerful for very small datasets, but its practical applicability is limited by complexity and cost. Pre-trained GANs with AdaFM emerge as the most effective and practical solution overall, given their balance of generalization, robustness and applicability. ADA follows a lightweight approach but with artifact risks. FSL is theoretically powerful but demonstrates less potential in practice. Finally, Consistency Regularization is best used as an auxiliary technique rather than a primary method. The limitations of the aforementioned approaches are often associated with high operational costs, an issue which is discussed in the next section.

6. Energy Consumption—Computational Cost

6.1. Problem-Solving Approaches

A promising strategy utilizes hardware-accelerated GANs that employ memristors-based neuromorphic computing [130]. Unlike traditional processors, memristors combine storage and computation, reducing the need for frequent memory access, which decreases energy consumption while enhancing training speed [131,132].

The proposed architecture includes a DiffBlock, which computes the cost function directly within the GAN, improving accuracy and efficiency. The generator and discriminator work in parallel: during the forward phase, both networks process data at the same time; during the backward phase, the DiffBlock sends error signals to update both models. This adaptive scheduling enhances parallelism, minimizes memory traffic, and improves energy efficiency [71]. Figure 7 illustrates the operational flow of the proposed memristor-based GAN architecture, including both forward and backward computations. The flow

data is clarified, demonstrating how neuromorphic computing reduces memory traffic while enhancing computational efficiency.

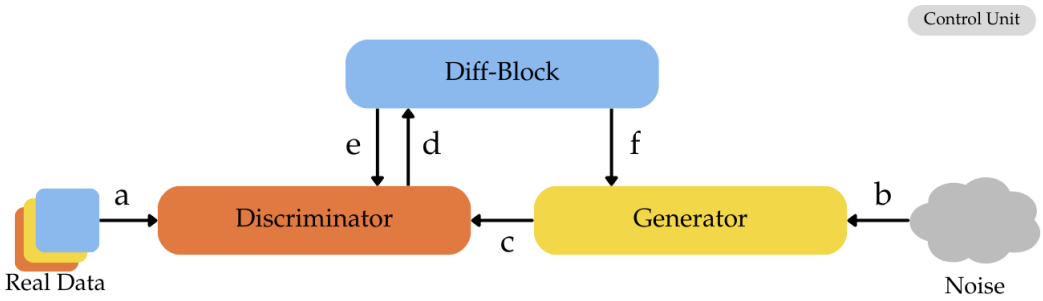

Figure 7. Architecture of the proposed memristor-based GAN computation. In this figure, "a" represents the real data input of the discriminator. The noise input provided to the generator producing synthetic data is denoted as "b". The generated data sent to the discriminator are referred to as "c". The discriminator's output, denoted by "d," is passed to the Diff-Block for loss calculation, and "e" utilizes this loss to update D's parameters. The Diff-Block updates the generator's weights in "f" by providing gradients. The path from "a" to "d" represents forward computations, while transitioning from "e" to "f", corresponds to the backward computations.

Table 2 illustrates that, according to the analysis conducted in [71], the memristor-based accelerator achieves considerable performance improvements over GPU- and FPGA-based counterparts. On the ImageNet dataset, the results yield 4.8× training speed, compared to FPGA, and 6.1× less energy demands than GPU. Large Scale Scene Understanding LSUN/bedroom dataset achieves the best performance as it delivers up to 5.5× faster training than FPGA and 6.1× lower energy consumption in comparison to GPU performance, while maintaining competitive output quality. These results demonstrate that neuromorphic hardware can significantly reduce the energy of GANs, allowing for more sustainable large-scale training and practical edge deployment.

Table 2. Performance and Energy Comparison of Hardware-Accelerated GAN compared to Traditional Accelerators on ImageNet and LSUN/bedroom datasets [71]. On ImageNet dataset and LSUN/bedroom dataset, the efficiency of the proposed model is estimated by comparing it with GPU-based and FPGA-based GAN accelerators in terms of time and energy. Thus, the speedup and the model's energy efficiency were computed.

	ImageNet				**LSUN/Bedroom**			
	Time (h)	**Speedup**	**Energy (KW/h)**	**Saving**	**Time (h)**	**Speedup**	**Energy (KW/h)**	**Saving**
Hardware-accelerated GAN	6.3	-	0.51	-	47.2	-	3.8	-
GPU	17	2.7×	3.1	6.1×	130	2.8×	23.4	6.1×
FPGA	30	4.8×	0.79	1.5×	255	5.5×	5.5	1.4×

Memristors inherently offer low precision in computation, data storage and transmission. Achieving higher precision usually results in slower speeds and higher design expenses. Moreover, utilizing 8-bit data precision memristors leads to minor reduction in system accuracy. Another limitation is the balance between time, area and parallelism. Computing efficiency requires high computing parallelism; however, this increases the area and design cost.

Model compression techniques, such as pruning and knowledge distillation, reduce the number of parameters and enable the deployment of smaller student models [133].

These approaches achieve computational efficiency by compressing GAN architectures. In [134], GAN Slimming (GS) is introduced, an All-in-One Compression framework that integrates channel pruning, quantization, and model distillation. Results indicate that GS achieves up to 47× compression with minimal quality loss, distinguishing it from other compression approaches. Training GANs is inherently unstable and insufficient. Multiple integrated compression strategies are embedded in a single framework often reducing instabilities, complicating the optimization process. Simple stacking of pruning, quantization, and distillation has led to degraded performance and low-quality outputs. Fixed discriminator–based methods commonly create inconsistencies regarding the compressed generator, reducing the quality of the generated data. The effectiveness of all-in-one compression methods relies on the architectural design of the student network. This dependency influences both the accuracy of the compressed model and its ability to generalize across different domains.

To further address the high energy consumption and memory traffic resulting from the dual-network design of GANs, the authors in [74] introduced the Fused Propagation (FusedProp) algorithm. Unlike conventional training, where generator and discriminator updates require separate forward-backward passes, FusedProp computes gradients for both networks simultaneously using one forward and backward propagation, scaling the propagation error by a constant factor λ.

$$\lambda = \frac{\partial L_G}{\partial L_D}, \tag{2}$$

In this scheme, depicted in Figure 8, the discriminator loss is multiplied by λ and the generator loss by $-\lambda$ variables, ensuring both networks receive their respective gradient signals within the same update cycle. FusedProp delivers 1.49× faster training compared to conventional GAN training. An inverted variant, Inverted Fused Propagation (InvFusedProp), was also proposed to address hinge loss when the parameter λ is not defined, by rescaling discriminator gradients to λ^{-1}.

$$\lambda^{-1} = \frac{\partial L_D}{\partial L_G}, \tag{3}$$

Both methods substantially reduce memory transfers and achieve higher training efficiency. The illustration directly contrasts conventional GAN training with the standard FusedProp method and with its inverted variant, offering a visual representation of how memory transfers and computational cost are reduced.

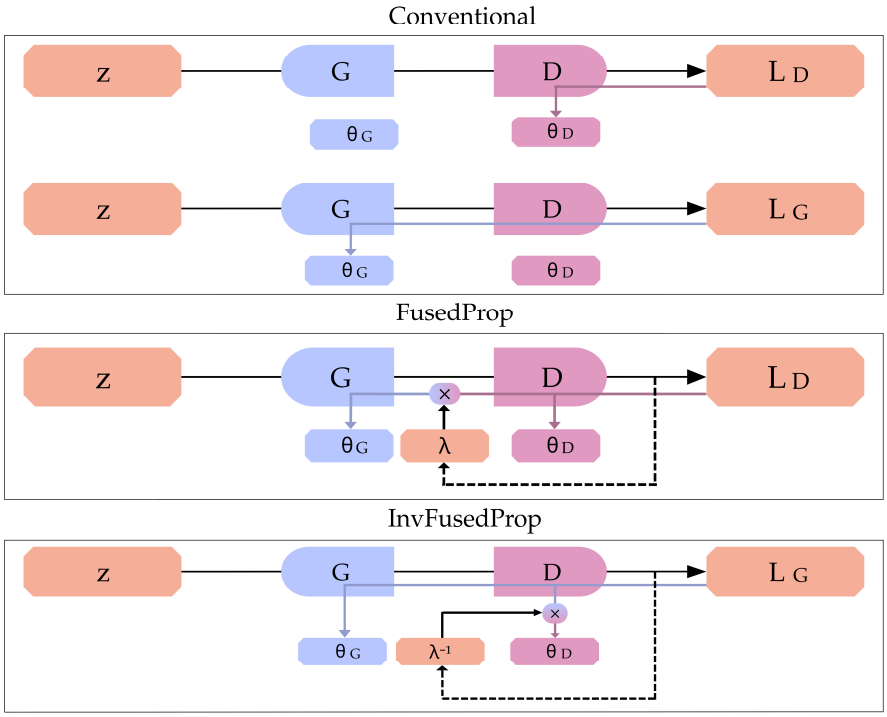

Figure 8. Illustration of FusedProp and InvFusedProp compared to conventional GANs. The gradients of the generator G and the discriminator D are denoted as "Θ_G" and "Θ_D", respectively. "λ" is the constant factor of the algorithm. The input noise provided to the generator is "z" while "Lg" and "Ld" are the loss functions of the generator and the discriminator, respectively.

Despite the computational efficiency of FusedProp algorithm several limitations have emerged. First, FusedProp results in restricted training speed when multiple discriminator updates are required per generator update. Second, gradient penalties like the R2 penalty are incompatible with FusedProp, because their second-order derivatives can destabilize the generator. However, the commonly used R1 penalty is fully supported. Third, architectures that combine auxiliary classification loss with adversarial loss are not currently incompatible, since balancing multiple objectives in a single backward pass remains unstable. In conclusion, FusedProp signifies an important advancement in enhancing GAN training by reducing memory use and computational cost. Future research should focus on enhancing stability and broadening compatibility with different penalties and GAN architectures. Among hardware-accelerated GAN, GAN slimming and FusedProp the most effective approach cannot be distinguished. A critical synthesis is provided in the next section, outlining strengths and weaknesses of every method.

6.2. Comparative Assessment

Section 6.2. presented three potential approaches to address the challenges associated with energy consumption and computational costs. These solutions comprise: (i) Memristor-based Neuromorphic GANs with DiffBlock, (ii) Model Compression, and (iii) Fused Propagation and Inverted Fused Propagation. Diffblock holds a strong potential for stable and large-scale training, but its dependence on specialized neuromorphic hardware restricts practical deployment. Model Compression including pruning, quantization and distillation frameworks such as GAN slimming, offers reduced computational costs, but its effectiveness relies on combined frameworks instead of standalone methods. FusedProp is a promising algorithmic optimization, but its application is limited by incompatibilities with advanced loss functions and complex training setups. When critically compared, the Memristor-based Neuromorphic method is the most balanced solution

offering long-term benefits for energy-efficient GANs. Model compression techniques are the most feasible option, providing computational savings without requiring specialized hardware. FusedPop provides notable efficiency but is best regarded as a complementary method due to its disadvantages.

Integrating control-based mechanisms like state-filtered disturbance rejection or multilayer neurocontrol with active disturbance rejection would connect traditional feedback theory with modern generative learning. This would create a path for energy-efficient, disturbance-resistant, and dynamically stable GAN architectures.

7. Privacy and Security

7.1. Addressing Approaches

To address these vulnerabilities, advanced GAN architectures have been proposed that balance privacy with model efficiency. One notable example is privacy-preserving GAN (privGAN), which employs a privacy discriminator (Dp) to distinguish memorized training samples and generated data [83]. When memorization is detected, the generator is penalized, reducing overfitting and the risk of information leakage [135]. In contrast to non-private GANs, PrivGAN successfully achieves its main objective of reducing MIAs. Nonetheless, this method is less resilient against other types of attacks. An additional limitation is that for specific datasets, increasing λ (regularization weight) or N (number of discriminators) beyond a certain point improves the privacy but lessens sample quality, reducing effectiveness against MIAs. It is also observed that as λ rises, certain classes become excessively represented. Moreover, privacy loss is increased with the number of epochs, indicating that extended training reduces the overall privacy guarantees. This is a limitation inherent to all GAN models and not exclusive to PrivGAN. However, this approach fails to provide a mechanism to address it. Finally, this approach demonstrates reduced efficiency on smaller datasets [136].

Two alternative models, Maximum Entropy GAN (MEGAN) and Mutual Information Minimization GAN (MIMGAN), have been developed to enhance privacy. MEGAN is based on Fano's inequality (Formula (4)) and Bayes error theory (Formula (5)). The overfitting of the discriminator is inversely related to classification error. Thus, by maximizing entropy, Fano's inequality limits overfitting and reduces the risk of MIAs [137].

$$P_e^d \geq \frac{H(r_e|x) - 1}{\log 2} \approx \frac{\mathrm{E}\left[H\left(D(x)\right)\right] - 1}{\log 2} \tag{4}$$

The variable P_e^d represents the error probability of an adversary attempting to decide whether a given sample belongs to the training set. $H(\cdot)$ denotes the entropy, i.e., the measure of uncertainty of a random variable, where r_e is the random variable representing membership information, and $|x$ is a candidate sample. The conditional entropy of the membership variable r_e given the sample x is expressed as $H(r_e|x)$. Furthermore, $\mathrm{E}\left[H\left(D(x)\right)\right]$ refers to the expected entropy of the discriminator's output distribution across all samples. Finally, the normalization factor $\log 2$ converts units into bits. Maximizing $\mathrm{E}\left[H\left(D(x)\right)\right]$ in Formula (4) the discriminator error is controlled from being too small and thus overfitting is prevented.

Bayes error [138] is the minimal possibility of a MIA to correctly distinguish training from non-training samples based on the score: $s = D(x)$

$$P_e^m = P(s \in R_1, \omega_0) + P(s \in R_0, \omega_1) \tag{5}$$

In Formula (5), the variable P_e^m denotes the probability of misclassification, s represents the sample being evaluated, while R_1 and R_0 correspond to the decisions "training" and "non-training", respectively. The true states are denoted as ω_1 for training and

ω_0 for non-training. The expression $P(s \in R_1, \omega_0)$ represents the probability of a false positive, whereas $P(s \in R_0, \omega_1)$ corresponds to the probability of false negative.

Together, Formulas (4) and (5) form the theoretical foundation of MEGAN, ensuring that score distributions remain indistinguishable between training and non-training samples. This helps mitigate MIAs [139]. Although MEGAN enhances privacy against membership inference attacks, two limitations must be considered. When the discriminator's outputs are near 0.5, it indicates that the model is at peak uncertainty and cannot distinguish real from generated samples. This limits the model's adaptability and reduces the diversity of its outputs. Consequently, its capacity to maintain privacy protection while maximizing data utility across various datasets is negatively affected [137].

MIMGAN reduces privacy leakage by minimizing the mutual information between training and generated data, thereby lowering the similarity between real and synthetic samples. This model effectively reduces the accuracy of MIAs by increasing the overlap parameter λ, while the generalization gap remains largely unaffected. However, the additional variance introduced in the discriminator's output may negatively impact training stability. Achieving a balance between privacy and utility requires careful parameter tuning. As a consequence, complexity of the system is increased, presenting challenges for effective implementation [137]. Differentially Private GAN (DPGAN) adds carefully controlled noise during training. Although this may decrease accuracy, it provides solid privacy guarantees. DPGANs face a trade-off among data privacy and output quality. Increasing the amount of added noise improves the level of privacy protection, but reduces the quality of the generated data. Training also becomes unstable when excessive noise disrupts convergence and makes optimization less effective. Furthermore, the accuracy in tasks such as classification decreases when models are trained on noisy synthetic data. This leads to reduced effectiveness in comparison to training with real data. In datasets with high sparsity, such as electronic health records, noise amplifies sparsity, causing information loss and limiting the model's ability to identify important connections across variables [140]. In [137], a comparative analysis was conducted on the Modified National Institute of Standards and Technology (MNIST) dataset [116], a widely used benchmark for classification and generative modeling. The results, as shown in Table 3, demonstrate the trade-offs among these architectures. In particular, standard GANs achieve the highest performance in generating realistic data but offer no privacy protection. MEGAN and MIMGAN balance accuracy and privacy, while privGAN and DPGAN prioritize privacy at a cost to output quality.

Table 3. Comparative analysis of privacy-preserving GAN models conducted by [137]. The standard GAN yields the most realistic outputs, while the DPGAN provides the strongest privacy guarantees.

Model	MIA (10% Accuracy)	GAN- Test Accuracy	Comments
GAN	59.20%	96.88%	Strong output quality, no privacy protection.
MEGAN	12.08%	94.16%	Reduced privacy risk with minimal accuracy loss.
MIMGAN	13.01%	92.97%	Similar performance to MEGAN
PrivGAN	12.18%	77.51%	Moderate privacy, limited output quality.
DPGAN	10.07%	59.67%	Strongest MIA protection, lowest model effectiveness.

Further innovations include Compressive Privacy GAN (CPGAN), which prevents data reconstruction by privatizing input features. Unlike standard GANs, CPGAN replaces the discriminator with a Service Module (S) and an Attacker Module (A). The generator transforms inputs into privacy-preserving representations (Z), which are validated by S for utility and tested by A for vulnerability. This dual-module framework enhances resistance to reconstruction attacks [19,141]. Figure 9 illustrates the CPGAN framework.

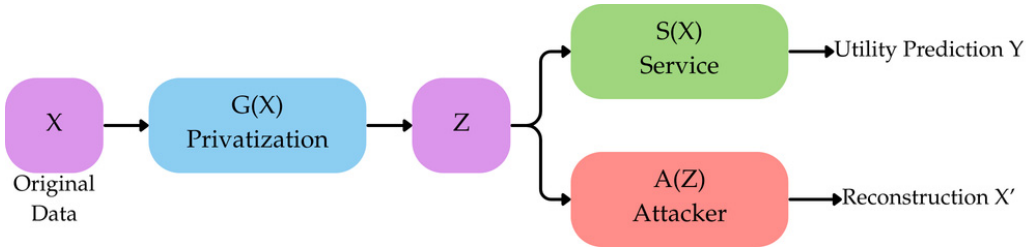

Figure 9. The CPGAN framework: The Service Module ensures usability of privatized data, and the Attacker Module assesses reconstruction resistance.

CPGAN preserves data privacy and, in comparison to alternative privacy-preserving mechanisms like DNN and RAN, achieves better balance between utility and privacy. Still, in smaller datasets the tradeoff is severe, leading to poor sample quality. Moreover, to ensure privacy, CPGAN diminishes classification accuracy, reducing utility accuracy by up to 2.58% for CIFAR-10 dataset compared to non-privacy preserving models. The privatizer resides in the device of the user and therefore is restricted to a lightweight design that remains undeveloped. Lastly, currently CPGAN is trained with raw sensitive data that are not protected from attacks during the training stage.

Another approach, Synthetic Adversarial GAN (SynGAN), generates synthetic adversarial traffic to improve robustness against DDoS attacks [142]. It comprises a generator, discriminator and evaluator, with the evaluator applying gradient boosting to measure realism through Root Mean Square Error (RMSE). While SynGAN enhances the training of IDS by generating realistic synthetic attack traffic, its capability to produce highly authentic attacks also raises concerns about potential misuse. Thus, further research is required to develop GAN architectures that balance privacy, effectiveness, and security. SynGAN generates synthetic adversarial traffic, but it also presents several limitations. Its current framework is designed only for producing DDoS attacks, which limits its efficiency. In addition, it relies on datasets such as NSL-KDD and CICIDS2017, which provide limited diversity and therefore reduce the reliability of the generated attacks. The evaluation process is constrained since the gradient boosting classifier used for validation had only limited success in differentiating real and synthetic traffic. Training with GP-WGANs is computationally demanding, as it requires powerful GPUs and large memory capacity. Finally, SynGAN has not yet been evaluated on commercial Network Intrusion Detection System (NIDS) leaving its effectiveness uncertain [143]. A critical analysis is provided in Section 7.2. focusing on two criteria for the classification of the methods: (i) solely privacy and (ii) the best balance between security and efficiency.

7.2. Analytical Evaluation

Various approaches have been proposed to achieve private-preserving GAN models. The methods presented in Section 7.1. are: (i) PrivGAN, (ii) Maximum Entropy GAN, (iii) Mutual Information Minimization GAN, (iv) Differentially Private GAN, (v) Compressive GAN, and (vi) SynGAN. PrivGAN balances between privacy and efficiency, remaining sensitive to hyperparameters and small data availability, restricting its robustness. ME-GAN offers privacy but is unable to maintain a variety of data, limiting its efficiency in application. Despite its training stability and degraded output quality, MIMGAN enhances privacy in a principled way. DPGAN provides the highest level of privacy protection, but its effectiveness is limited. CPGAN is a promising approach as it presents an overall private model but remains in process development, relying on lightweight designs. In network security, SynGAN is the most efficient architecture; however, its generalizability is limited, and it is constrained by significant computational costs. Considering that the classification is determined by privacy strength, DPGAN is the most efficient

method. However, an overall balanced approach is typically required. Therefore, PrivGAN provides the best balance between privacy and efficiency. MEMGAN and MIM-GAN are strong theoretical foundations in privacy. DPGAN, as stated earlier, is the most privacy-aware model but results in lower sample quality. CPGAN is still an innovative framework with promising privacy and effectiveness. Finally, SynGAN is highly specialized, useful primarily for DDoS attacks and various network attacks. This critical analysis highlights that the trade-offs between accuracy and privacy are the main issue regarding the scalability of research to larger datasets, while the sensitivity to limited data and the challenges of real-world deployment remain open challenges that should be addressed in future work.

8. Conclusions and Limitations

Although diffusion models and transformer-based generators have become strong alternatives, GANs continue to set the standard for measuring realism, efficiency, and reliability in creating synthetic data. This survey has examined critical yet under-researched challenges of generative adversarial networks, including data limitations, efficiency and security. While notable progress has been achieved, significant gaps remain, such as privacy-preserving models, pre-trained embeddings and lightweight architectures. The insufficiency of high-quality training data and environmental implications for energy-efficient and privacy-aware models continue to pose major setbacks. Current research provides solutions for models that are computationally and energetically expensive, insufficient and vulnerable models, without addressing the underlying problems. This review contributes to the relevant literature by exploring under-investigated challenges and offering further insights into the gravity and significance of these issues. By synthesizing findings from existing studies and related research, the survey assembles a series of distinct solutions aiming to mitigate these challenges. The greater part of the objective is achieved, as several approaches are proposed to address concerns relating to privacy, security, data limitation, cost and efficiency. Nevertheless, each solution resolves part of the problem while disregarding other aspects of the issue. As summarized in Section 7, most approaches show both advantages and disadvantages, and no method provides a comprehensive solution to address all challenges. Instead, existing approaches typically target specific domains, models or individual issues. Future research should emphasize developing a GAN architecture that integrates limited data training with energy-efficient computation and privacy-aware mechanisms, aiming to balance utility, robustness and security. Thus, this work offers insights that may guide future developments toward more reliable and sustainable generative models.

This literature review is informative in scope but has certain limitations that should be acknowledged. Our analysis primarily relies on widely used datasets such as MNIST, CIFAR-10 and LSUN, not fully capturing the complexity of real-world data. The scalability of larger diverse datasets is not examined, limiting the easy application of this work. While this survey emphasizes computational and energy efficiency, most methods often provide partial effectiveness, degrading the model's accuracy. Most architectures provide theoretical improvement, but their deployment is unstable due to real-world constraints, such as hardware availability, latency and regulatory requirements. Furthermore, as this work is a survey study, it synthesizes methods and models for already existing literature without providing a new experimental validation. Therefore, the direct comparison between architectures and approaches is limited under unified evaluation frameworks. These limitations suggest further efficient and large-scale research is required to strengthen the practical relevance of the findings presented here.

9. Future Work

Within the scope of this survey considering the limitations mentioned, several directions for future research emerge. Future work should concentrate on developing privacy-aware and efficiency-oriented GANs, able to achieve desirable performance in larger and heterogeneous datasets. More attention should be given to understanding how secure but computationally efficient models evolve across different domains and how to maintain high accuracy under realistic conditions such as limited data. Reducing the energy footprint of GAN training, avoiding stability loss and decreased accuracy, is a pressing issue that deserves further exploration. Furthermore, systematic evaluations under unified frameworks would provide a basis for the comparison and analysis of the models, enabling a fair assessment of private and efficient GANs.

Beyond the scope of this review, further research is warranted in several under-explored areas. Non-visual domains such as signal generation, alphanumeric character recognition, time-series analysis, and graph simulation, are neglected when compared to video and image domains. These areas introduce unique setbacks such as data scarcity, labeling costs, and temporal dynamics, not addressed by the current adversarial methods. Another important gap in research is the lack of information of the semantic and geometric properties of GANs, particularly in multi-modal and 3D GANs. Enhancing the understanding of the semantic features would support more secure and accurate data. Finally, evaluation metrics fail to align with human perception, are vulnerable to manipulation and are unstable across domains, limiting the reliability of comparisons between models.

Future research should also explore the interaction between control theory and adversarial learning. Concepts like state-filtered disturbance rejection and multilayer neurocontrol with active disturbance rejection can guide adaptive gradient filtering and robust optimization techniques. Embedding these controllers into GAN frameworks can reduce training oscillations, minimize unnecessary calculations, and improve stability in non-stationary or resource-limited conditions.

In summary, future work should concentrate in achieving efficient, computationally lightweight and private design while simultaneously addressing broader challenges relating to scalability, feasibility, non-visual domains, latent space understanding and evaluation metrics reliability. Resolving these issues will be crucial for ensuring that GANs evolve into a secure, efficient and trustworthy tool capable of operating across a wide range of scientific and real-world applications.

Author Contributions: Conceptualization, E.T. and I.V.; methodology, N.E.A., E.V.B. and M.M.; validation, E.T., I.V., A.P. and M.S.; formal analysis, N.E.A., E.V.B. and M.M.; investigation, N.E.A., E.V.B. and M.M.; resources, N.E.A., E.V.B. and M.M.; data curation, E.T. and A.P.; writing—original draft preparation, N.E.A., E.V.B. and M.M.; writing—review and editing, N.E.A., E.V.B., M.M. and E.T.; visualization, N.E.A., E.V.B. and M.M.; supervision, M.S.; All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The data presented in this study are available in the cited references. These data were derived from publicly available resources, including IEEE Xplore (Institute of Electrical and Electronics Engineers, New York, NY, USA; https://ieeexplore.ieee.org/, accessed on 2 May 2025), ACM Digital Library (Association for Computing Machinery, New York, NY, USA; https://dl.acm.org/, accessed on 2 May 2025), SpringerLink (Springer Nature, Berlin, Germany; https://link.springer.com/, accessed on 2 May 2025), Elsevier ScienceDirect (Elsevier B.V., Amsterdam, The Netherlands; https://www.sciencedirect.com/, accessed on 2 May 2025), Google Scholar (Google LLC, Mountain View, CA, USA; https://scholar.google.com/, accessed on 2 May 2025), and arXiv (Cornell University, Ithaca, NY, USA; https://arxiv.org/, accessed on 2 May 2025).

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

ADRC	Active Disturbance Rejection Control
ADA	Adaptive Discriminator Augmentation
AdaFm	Adaptive Filter Modulation
AR	Augmented Reality
CGAN	Conditional GAN
CNN(s)	Convolutional Neural Network(s)
CPGAN	Compressive Privacy GAN
CRF	Conditional Random Field
DCGAN	Deep Convolutional GAN
DeLiGAN	Diverse and Limited data GAN
DL	Deep Learning
DDoS	Distributed Denial of Service
DPGAN	Differentially Private GAN
DynaGAN	Dynamic GAN
EEG	Electroencephalogram
FreGAN	Frequency-aware GAN
FPGA	Field-Programmable Gate Array
FSL	Few-Shot Learning
FusedProp	Fused Propagation
GAN(s)	Generative Adversarial Network(s)
GP-GAN	Gaussian-Poisson GAN
GS	GAN Slimming
(Inv)FusedProp	(Inverted) Fused Propagation
IS	Inception Score
LSUN	Large Scale Scene Understanding
MAML	Model-Agnostic Meta-Learning
MANN	Memory-Augmented Neural Network
MEGAN	Maximum Entropy GAN
MI	Model Inversion
MIA	Membership Inference Attacks
MIMGAN	Information Minimization GAN
ML	Machine Learning
MNIST	Modified National Institute of Standards and Technology
MUA	Models Under Attack
(N)IDS	(Network) Intrusion Detection System
PI-GAN	Plant Identification GAN
PII	Personally Identifiable Information
PrivGAN	Privacy-preserving GAN
RWM-CGAN	Residual Weight Masking Conditional GAN
RMSE	Root Mean Square Error
SLR	Systematic Literature Review
SFDRC	State Filtered Disturbance Rejection Control
SynGAN	Synthetic Adversarial GAN
TDP	Thermal Design Power
VAE	Variational Autoencoders
WGAN	Wasserstein GAN
WGAN-GP	Wasserstein GAN with Gradient Penalty

References

1. LeCun, Y.; Bengio, Y.; Hinton, G. Deep learning. *Nature* **2015**, *521*, 436–444. https://doi.org/10.1038/nature14539.
2. Goodfellow, I.; Bengio, Y.; Courville, A. *Deep Learning*; MIT Press: Cambridge, MA, USA, 2016. Available online: https://mit-press.mit.edu/9780262035613/deep-learning/ (accessed on 2 May 2025).

3. Creswell, A.; White, T.; Dumoulin, V.; Arulkumaran, K.; Sengupta, B.; Bharath, A.A. Generative adversarial networks: An overview. *IEEE Signal Process. Mag.* **2018**, *35*, 53–65. https://doi.org/10.1109/MSP.2017.2765202.
4. Karras, T.; Aila, T.; Laine, S.; Lehtinen, J. Progressive growing of GANs for improved quality, stability, and variation. *arXiv* **2017**, arXiv:1710.10196. https://doi.org/10.48550/arXiv.1710.10196.
5. Brock, A.; Donahue, J.; Simonyan, K. Large scale GAN training for high fidelity natural image synthesis. *arXiv* **2019**, arXiv:1809.11096. https://doi.org/10.48550/arXiv.1809.11096.
6. Radford, A.; Metz, L.; Chintala, S. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv* **2015**, arXiv:1511.06434. https://doi.org/10.48550/arXiv.1511.06434.
7. Karras, T.; Laine, S.; Aila, T. A style-based generator architecture for generative adversarial networks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 15–20 June 2019; pp. 4401–4410. https://doi.org/10.1109/TPAMI.2020.2970919.
8. Yu, L.; Zhang, W.; Wang, J.; Yu, Y. SeqGAN: Sequence generative adversarial nets with policy gradient. In Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, San Francisco, CA, USA, 4–9 February 2017. https://doi.org/10.1609/aaai.v31i1.10804.
9. Yi, X.; Walia, E.; Babyn, P. Generative adversarial network in medical imaging: A review. *Med. Image Anal.* **2019**, *58*, 101552. https://doi.org/10.1016/j.media.2019.101552.
10. Eckerli, F.; Osterrieder, J. Generative adversarial networks in finance: An overview. *arXiv* **2021**, arXiv:2106.06364. https://doi.org/10.48550/arXiv.2106.06364.
11. Luleci, F.; Catbas, F.N.; Avci, O. A literature review: Generative adversarial networks for civil structural health monitoring. *Front. Built Environ.* **2022**, *8*, 1027379. https://doi.org/10.3389/fbuil.2022.1027379.
12. Alqahtani, H.; Kavakli-Thorne, M.; Kumar, G. Applications of generative adversarial networks (GANs): An updated review. *Arch. Comput. Methods Eng.* **2021**, *28*, 525–552. https://doi.org/10.1007/s11831-019-09388-y.
13. Saxena, D.; Cao, J. Generative adversarial networks (GANs): Challenges, solutions, and future directions. *ACM Comput. Surv.* **2022**, *54*, 63. https://doi.org/10.1145/3446374.
14. Mujeeb, S.; Javaid, N. Deep Learning Based Carbon Emissions Forecasting and Renewable Energy's Impact Quantification. *IET Renew. Power Gener.* **2023**, *17*, 873–884. https://doi.org/10.1049/rpg2.12641.
15. Henderson, P.; Hu, J.; Romoff, J.; Brunskill, E.; Jurafsky, D.; Pineau, J. Towards the systematic reporting of the energy and carbon footprints of machine learning. *J. Mach. Learn. Res.* **2020**, *21*, 1–43. Available online: https://www.jmlr.org/papers/v21/20-312.html (accessed on 9 June 2025).
16. Strubell, E.; Ganesh, A.; McCallum, A. Energy and policy considerations for deep learning in NLP. In Proceedings of the AAAI Conference on Artificial Intelligence, New York, NY, USA, 7–12 February 2020; pp. 3645–3650. https://doi.org/10.1609/aaai.v34i09.7123.
17. Schwartz, R.; Dodge, J.; Smith, N.A.; Etzioni, O. Green AI. *Commun. ACM* **2020**, *63*, 54–63. https://doi.org/10.1145/3381831.
18. Li, T.; Clifton, C. Differentially private imaging via latent space manipulation. *arXiv* **2021**, arXiv:2103.05472. https://doi.org/10.48550/arXiv.2103.05472.
19. Cai, Z.; Xiong, Z.; Xu, H.; Wang, P.; Li, W.; Pan, Y. Generative adversarial networks: A survey toward private and secure applications. *ACM Comput. Surv.* **2022**, *54*, 132. https://doi.org/10.1145/3459992.
20. Padariya, D.; Wagner, I.; Taherkhani, A.; Boiten, E. Privacy-Preserving Generative Models: A Comprehensive Survey. *arXiv* **2025**, arXiv:2502.03668. https://doi.org/10.48550/arXiv.2502.03668.
21. Gurumurthy, S.; Sarvadevabhatla, R.K.; Babu, R.V. DeLiGAN: Generative adversarial networks for diverse and limited data. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Honolulu, HI, USA, 21–26 July 2017; pp. 4941–4949. https://doi.org/10.1109/CVPR.2017.525.
22. Yadav, P.; Sihag, G.; Vijay, V. Rebalancing the scales: A systematic mapping study of generative adversarial networks (GANs) in addressing data imbalance. *arXiv* **2025**, arXiv:2502.16535. https://doi.org/10.48550/arXiv.2502.16535.
23. Lin, Z.; Sekar, V.; Fanti, G. On the privacy properties of GAN-generated samples. In Proceedings of the International Conference on Artificial Intelligence and Statistics, Online, 13–15 April 2021; pp. 1522–1530. https://doi.org/10.48550/arXiv.2206.01349.
24. Alshantti, A.; Rasheed, A.; Westad, F. Privacy re-identification attacks on tabular GANs. *Secur. Priv.* **2025**, *8*, e469. https://doi.org/10.48550/arXiv.2404.00696.
25. Mirza, M.; Osindero, S. Conditional generative adversarial nets. *arXiv* **2014**, arXiv:1411.1784. https://doi.org/10.48550/arXiv.1411.1784.

26. Denton, E.L.; Chintala, S.; Fergus, R. Deep Generative Image Models Using a Laplacian Pyramid of Adversarial Networks. *arXiv* **2015**, arXiv:1506.05751. https://doi.org/10.48550/arXiv.1506.05751.
27. Arjovsky, M.; Chintala, S.; Bottou, L. Wasserstein generative adversarial networks. In Proceedings of International Conference on Machine Learning (ICML), Sydney, NSW, Australia, 6–11 August 2017. Available online: https://proceedings.mlr.press/v70/arjovsky17a.html (accessed on 30 May 2025).
28. Salimans, T.; Goodfellow, I.; Zaremba, W.; Cheung, V.; Radford, A.; Chen, X. Improved techniques for training GANs. In *Advances in Neural Information Processing Systems*; American Institute of Physics: College Park, MD, USA, 2016; Volume 29. Available online: https://proceedings.neurips.cc/paper/2016/hash/8a3363abe792db2d8761d6403605aeb7-Abstract.html (accessed on 21 May 2025).
29. Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; Hochreiter, S. GANs trained by a two time-scale update rule converge to a local Nash equilibrium. In *Advances in Neural Information Processing Systems*; American Institute of Physics: College Park, MD, USA, 2017; Volume 30. Available online: https://proceedings.neurips.cc/paper/2017/hash/8a1d694707eb0fefe65871369074926d-Abstract.html (accessed on 16 July 2025).
30. Isola, P.; Zhu, J.Y.; Zhou, T.; Efros, A.A. Image-to-image translation with conditional adversarial networks. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Honolulu, HI, USA, 21–26 July 2017; pp. 1125–1134. https://doi.org/10.1109/CVPR.2017.632.
31. Zhu, J.Y.; Park, T.; Isola, P.; Efros, A.A. Unpaired image-to-image translation using cycle-consistent adversarial networks. In Proceedings of the IEEE International Conference on Computer Vision, Venice, Italy, 22–29 October 2017. https://doi.org/10.1109/ICCV.2017.244.
32. Karras, T.; Laine, S.; Aittala, M.; Hellsten, J.; Lehtinen, J.; Aila, T. Analyzing and improving the image quality of StyleGAN. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 14–19 June 2020. https://doi.org/10.1109/CVPR42600.2020.00813.
33. Karras, T.; Aittala, M.; Laine, S.; Härkönen, E.; Hellsten, J.; Lehtinen, J.; Aila, T. Alias-free generative adversarial networks. *arXiv* **2021**, arXiv:2106.12423. https://doi.org/10.48550/arXiv.2106.12423.
34. Yang, G. State Filtered Disturbance Rejection Control. *Nonlinear Dyn.* **2025**, *113*, 6739–6755. https://doi.org/10.1007/s11071-024-10449-6.
35. Yang, G.; Yao, J. Multilayer Neurocontrol of High-Order Uncertain Nonlinear Systems with Active Disturbance Rejection. *Int. J. Robust Nonlinear Control* **2024**, *34*, 2972–2987. https://doi.org/10.1002/rnc.7118.
36. Shiri, M.; Bruno, A.; Loiacono, D. Memory-Efficient 3D High-Resolution Medical Image Synthesis Using CRF-Guided GANs. In Proceedings of the International Conference on Pattern Recognition (ICPR), Kolkata, India, 2–5 December 2024; Springer Nature: Cham, Switzerland, 2025; pp. 184–194. https://doi.org/10.1007/978-3-031-87660-8_14.
37. Rahman, Z.U.; Asaari, M.S.M.; Ibrahim, H.; Abidin, I.S.Z.; Ishak, M.K. Generative adversarial networks (GANs) for image augmentation in farming: A review. *IEEE Access* **2024**, *12*, 179912–179943. https://doi.org/10.1109/ACCESS.2024.3505989.
38. Olaniyi, E.; Chen, D.; Lu, Y.; Huang, Y. Generative adversarial networks (GANs) for image augmentation in agriculture: A systematic review. *Comput. Electron. Agric.* **2022**, *200*, 107208. https://doi.org/10.1016/j.compag.2022.107208.
39. Donahue, C.; McAuley, J.; Puckette, M. Adversarial audio synthesis. *arXiv* **2018**, arXiv:1802.04208. https://doi.org/10.48550/arXiv.1802.04208.
40. Fahimi, F.; Zhang, Z.; Goh, W.B.; Ang, K.K.; Guan, C. Towards EEG generation using GANs for BCI applications. In Proceedings of the IEEE EMBS International Conference on Biomedical & Health Informatics, Chicago, IL, USA, 19–22 May 2019; pp. 1–4. https://doi.org/10.1109/BHI.2019.8834503.
41. Amin, S.S.; Ragha, L. Alphanumeric character recognition on tiny dataset. In Proceedings of the International Conference on Computer Networks, Big data and IoT (ICCBI), Madurai, India, 15–16 December 2020; Springer: Cham, Switzerland, 2020; Volume 49. https://doi.org/10.1007/978-3-030-43192-1_73.
42. Iyer, S.; Hou, T.T. GAT-GAN: A graph-attention-based time-series generative adversarial network. *arXiv* **2023**, arXiv:2306.01999. https://doi.org/10.48550/arXiv.2306.01999.
43. Yoon, J.; Jarrett, D.; Van der Schaar, M. Time-series generative adversarial networks. In *Advances in Neural Information Processing Systems*; American Institute of Physics: College Park, MD, USA, 2019; Volume 32. Available online: https://papers.nips.cc/paper/8789-time-series-generative-adversarial-networks.pdf (accessed on 24 July 2025).
44. Wang, J.; Ju, B.; Qian, X.; Ye, M.; Huo, W. Graph-based generative adversarial networks for molecular generation with noise diffusion. In Proceedings of the 14th International Conference on Bioscience, Biochemistry and Bioinformatics, Kyoto, Japan, 12–15 January 2024. https://doi.org/10.1145/3640900.3640911.

45. Choi, J.; Lee, J.; Yoon, C.; Park, J.H.; Hwang, G.; Kang, M. Do not escape from the manifold: Discovering the local coordinates on the latent space of GANs. *arXiv* **2021**, arXiv:2106.06959. https://doi.org/10.48550/arXiv.2106.06959.
46. Bhagoji, A.N.; He, W.; Li, B.; Song, D. Exploring the space of black-box attacks on deep neural networks. *arXiv* **2017**, arXiv:1712.09491. https://doi.org/10.48550/arXiv.1712.09491.
47. Zhang, K. On mode collapse in generative adversarial networks. In *Artificial Neural Networks and Machine Learning—ICANN 2021, Proceedings of the International Conference on Artificial Neural Networks, Online, 14–17 September 2021*; Farkaš, I., Ed.; Springer: Cham, Switzerland, 2021; Volume 12892. https://doi.org/10.1007/978-3-030-86340-1_45.
48. Figueira, A.; Vaz, B. Survey on synthetic data generation, evaluation methods and GANs. *Mathematics* **2022**, *10*, 2733. https://doi.org/10.3390/math10152733.
49. Karras, T.; Aittala, M.; Hellsten, J.; Laine, S.; Lehtinen, J.; Aila, T. Training generative adversarial networks with limited data. *arXiv* **2020**, arXiv:2006.06676. https://doi.org/10.48550/arXiv.2006.06676.
50. Wang, Y.; Yao, Q.; Kwok, J.T.; Ni, L.M. Generalizing from a few examples: A survey on few-shot learning. ACM Comput. *ACM Comput. Surv.* **2021**, *53*, 63. https://doi.org/10.1145/3386252.
51. Tanaka, F.H.K.; Aranha, C. Data augmentation using GANs. *arXiv* **2019**, arXiv:1904.09135. https://doi.org/10.48550/arXiv.1904.09135.
52. Zhao, M.; Cong, Y.; Carin, L. On leveraging pretrained GANs for generation with limited data. In Proceedings of the International Conference on Machine Learning, Vienna, Austria, 12–18 July 2020; pp. 11340–11351. Available online: https://proceedings.mlr.press/v119/zhao20a.html (accessed on 12 July 2025).
53. Bai, C.-Y.; Lin, H.-T.; Raffel, C.; Kan, W.C.-W. On Training Sample Memorization: Lessons from Benchmarking Generative Modeling with a Large-scale Competition. In Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, Online, 14–18 August 2021. https://doi.org/10.1145/3447548.3467198.
54. Feng, Q.; Guo, C.; Benitez-Quiroz, F.; Martinez, A.M. When do GANs replicate? On the choice of dataset size. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Montreal, BC, Canada, 11–17 October 2021. https://doi.org/10.1109/ICCV48922.2021.00663.
55. Li, Y.; Chao, X. Semi-supervised few-shot learning approach for plant diseases recognition. *Plant Methods* **2021**, *17*, 68, https://doi.org/10.1186/s13007-021-00770-1.
56. Middel, L.; Palm, C.; Erdt, M. Synthesis of medical images using GANs. In *Uncertainty for Safe Utilization of Machine Learning in Medical Imaging and Clinical Image-Based Procedures, Proceedings of the First International Workshop, UNSURE 2019, and 8th International Workshop, CLIP 2019, Held in Conjunction with MICCAI 2019, Shenzhen, China, 17 October 2019*; Springer International Publishing: Cham, Switzerland, 2019, pp. 125–134. https://doi.org/10.1007/978-3-030-32689-0_13.
57. Xu, F.; Dan, Y.; Yan, K.; Ma, Y.; Wang, M. Low-resource language discrimination toward Chinese dialects with transfer learning and data augmentation. *ACM Trans. Asian Low-Resour. Lang. Inf. Process.* **2022**, *21*, 27. https://doi.org/10.1145/3473499.
58. Carrasco, X.A.; Elnagar, A.; Lataifeh, M. A generative adversarial network for data augmentation: The case of Arabic regional dialects. *Procedia Comput. Sci.* **2021**, *189*, 92–99. https://doi.org/10.1016/j.procs.2021.05.072.
59. Chen, Z.; Ramachandra, B.; Vatsavai, R.R. Consistency regularization with generative adversarial networks for semi-supervised learning. *arXiv* **2020**, arXiv:2007.03844. https://doi.org/10.48550/arXiv.2007.03844.
60. Zhao, Z.; Singh, S.; Lee, H.; Zhang, Z.; Odena, A.; Zhang, H. Improved consistency regularization for GANs. In Proceedings of the AAAI Conference on Artificial Intelligence, Online, 2–9 February 2021. https://doi.org/10.1609/aaai.v35i12.17317.
61. Zhang, H.; Zhang, Z.; Odena, A.; Lee, H. Consistency regularization for generative adversarial networks. *arXiv* **2019**, arXiv:1910.12027. https://doi.org/10.48550/arXiv.1910.12027.
62. Bansal, A.; Sharma, R.; Kathuria, M. A Systematic Review on Data Scarcity Problem in Deep Learning: Solution and Applications. *ACM Comput. Surv.* **2022**, *54*, 208. https://doi.org/10.1145/3502287.
63. Abdalla, H.; Kumar, Y.; Marchena, J.; Guzman, S.; Awlla, A.; Gheisari, M.; Cheraghy, M. The Future of Artificial Intelligence in the Face of Data Scarcity. *Comput. Mater. Contin.* **2025**, *84*, 1073–1099. https://doi.org/10.32604/cmc.2025.063551.
64. Narayanan, H.; Ghanta, S. A Study of Data Augmentation Techniques to Overcome Data Scarcity in Wound Classification Using Deep Learning. *arXiv* **2024**, arXiv:2411.02456. https://doi.org/10.48550/arXiv.2411.02456.
65. Li, Z.; Usman, M.; Tao, R.; Xia, P.; Wang, C.; Chen, H.; Li, B. A Systematic Survey of Regularization and Normalization in GANs. *ACM Comput. Surv.* **2023**, *55*, 232. https://doi.org/10.1145/3569928.
66. Fayaz, S.; Alasmary, W.; Javed, A.R.; Band, S.S.; Abdel-Basset, M.; Baz, M.; Ghoneim, A. Advancements in Data Augmentation and Transfer Learning: A Comprehensive Survey to Address Data Scarcity Challenges. *Recent Adv. Comput. Sci. Commun.* **2024**, *17*, 14–35. https://doi.org/10.2174/0126662558286875231215054324.

67. Kim, S.; Kang, S.; Han, D.; Kim, S.; Kim, S.; Yoo, H.-J. An Energy-Efficient GAN Accelerator with On-Chip Training for Domain-Specific Optimization. *IEEE J. Solid-State Circuits* **2021**, *56*, 2968–2980. https://doi.org/10.1109/JSSC.2021.3094469.
68. Bhattacharya, G. From DNNs to GANs: Review of Efficient Hardware Architectures for Deep Learning. *arXiv* **2021**, arXiv:2107.00092. https://doi.org/10.48550/arXiv.2107.00092.
69. Douwes, C.; Esling, P.; Briot, J.-P. Energy consumption of deep generative audio models. *arXiv* **2021**, arXiv:2107.02621. https://doi.org/10.48550/arXiv.2107.02621.
70. Chi, P.; Li, S.; Xu, C.; Zhang, T.; Zhao, J.; Liu, Y.; Wang, Y.; Xie, Y. PRIME: A novel processing-in-memory architecture for neural network computation in ReRAM-based main memory. *SIGARCH Comput. Archit. News* **2016**, *44*, 27–39. https://doi.org/10.1145/3007787.3001140.
71. Liu, F.; Liu, C.; Bi, F. A memristor based unsupervised neuromorphic system towards fast and energy-efficient GAN. *arXiv* **2018**, arXiv:1806.01775. https://doi.org/10.48550/arXiv.1806.01775.
72. Li, M.; Lin, J.; Ding, Y.; Liu, Z.; Zhu, J.Y.; Han, S. GAN compression: Efficient architectures for interactive conditional GANs. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 14–19 June 2020; pp. 5284–5294. https://doi.org/10.1109/CVPR42600.2020.00533.
73. Tantawy, D.; Zahran, M.; Wassal, A. A survey on GAN acceleration using memory compression techniques. *J. Eng. Appl. Sci.* **2021**, *68*, 47. https://doi.org/10.1186/s44147-021-00045-5.
74. Polizzi, Z.; Tsai, C.-Y. FusedProp: Towards efficient training of generative adversarial networks. *arXiv* **2020**, arXiv:2004.03335. https://doi.org/10.48550/arXiv.2004.03335.
75. Satyam, S.; Nikam, H.; Sahay, S. Energy-efficient implementation of generative adversarial networks on passive RRAM crossbar arrays. *arXiv* **2021**, arXiv:2111.14484. https://doi.org/10.48550/arXiv.2111.14484.
76. Corda, S.; Veenboer, B.; Tolley, E. PMT: Power measurement toolkit. In Proceedings of the 2022 IEEE/ACM International Workshop on HPC User Support Tools (HUST), Dallas, TX, USA, 13–18 November 2022; pp. 44–47. https://doi.org/10.1109/HUST56722.2022.00011.
77. Lacoste, A.; Luccioni, A.; Schmidt, V.; Dandres, T. Quantifying the carbon emissions of machine learning. *arXiv* **2019**, arXiv:1910.09700. https://doi.org/10.48550/arXiv.1910.09700.
78. De Vries, A. The growing energy footprint of artificial intelligence. *Joule* **2023**, *7*, 2191–2194. https://doi.org/10.1016/j.joule.2023.09.004.
79. Patterson, D.; Gonzalez, J.; Hölzle, U.; Le, Q.; Liang, C.; Munguia, L.M.; Rothchild, D.; So, D.; Texier, M.; Dean, J. The carbon footprint of machine learning training will plateau, then shrink. *Computer* **2022**, *55*, 18–28. https://doi.org/10.1109/MC.2022.3148714.
80. Jayakodi, N.K.; Doppa, J.R.; Pande, P.P. SETGAN: Scale and energy trade-off GANs for image applications on mobile platforms. In Proceedings of the 2020 IEEE/ACM International Conference on Computer-Aided Design (ICCAD), San Diego, CA, USA, 2–5 November 2020; pp. 1–9. https://doi.org/10.1145/3400302.3415675.
81. Kumar, A.; Anand, K.; Mandloi, S.; Mishra, A.; Thakur, A.; Kasera, N.; AP, P. CoroNetGAN: Controlled pruning of GANs via hypernetworks. In Proceedings of the IEEE/CVF International Conference on Computer Vision, Paris, France, 2–3 October 2023; pp. 1262–1271. https://doi.org/10.1109/ICCVW60793.2023.00136.
82. Zhou, J.; Chen, Z.; Li, J.; Gu, Z.; Liang, Z. Property inference attacks against GANs. arXiv **2021**, arXiv:2111.07608. https://doi.org/10.48550/arXiv.2111.07608
83. Hu, H.; Salcic, Z.; Sun, L.; Dobbie, G.; Yu, P.S.; Zhang, X. Membership Inference Attacks on Machine Learning: A Survey. *ACM Comput. Surv.* **2022**, *54*, 235. https://doi.org/10.1145/3523273.
84. Ekramifard, A.; Amintoosi, H.; Hosseini Seno, S.A. Detection of membership inference attacks on GAN models. *ISeCure* **2025**, *17*, 43. https://doi.org/10.22042/isecure.2024.461639.1131.
85. Shafik, W. Generative adversarial networks: Security, privacy, and ethical considerations. In *Generative Artificial Intelligence (AI) Approaches for Industrial Applications*; Springer: Cham, Switzerland, 2025; pp. 93–117. https://doi.org/10.1007/978-3-031-76710-4_5.
86. Liu, Y.; Acharya, U.R.; Tan, J.H. Preserving Privacy in Healthcare: A Systematic Review of Deep Learning Approaches for Synthetic Data Generation. *Comput. Methods Programs Biomed.* **2025**, *260*, 108571. https://doi.org/10.1016/j.cmpb.2024.108571.
87. Fredrikson, M.; Jha, S.; Ristenpart, T. Model inversion attacks that exploit confidence information and basic countermeasures. In Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, Denver, CO, USA, 12–16 October 2015; pp. 1322–1333. https://doi.org/10.1145/2810103.2813677.

88. Yang, W.; Wang, S.; Wu, D.; Cai, T.; Zhu, Y.; Wei, S.; Zhang, Y.; Yang, X.; Tang, Z.; Li, Y. Deep learning model inversion attacks and defenses: A comprehensive survey. *Artif. Intell. Rev.* **2025**, *58*, 242. https://doi.org/10.1007/s10462-025-11248-0.
89. Li, Z.; Yang, M.; Liu, Y.; Wang, J.; Hu, H.; Yi, W.; Xu, X. GAN you see me? Enhanced data reconstruction attacks against split inference. In *Advances in Neural Information Processing Systems*; American Institute of Physics: College Park, MD, USA, 2023; Volume 36, pp. 54554–54566. Available online: https://proceedings.neurips.cc/paper_files/paper/2023/file/ab003a4f85ecb1b7b1514ff539dc7395-Paper-Conference.pdf (accessed on 18 August 2025).
90. Shieh, C.-S.; Nguyen, T.-T.; Lin, W.-W.; Huang, Y.-L.; Horng, M.-F.; Lee, T.-F.; Miu, D. Detection of adversarial DDoS attacks using generative adversarial networks with dual discriminators. *Symmetry* **2022**, *14*, 66. https://doi.org/10.3390/sym14010066.
91. Dunmore, A.; Jang-Jaccard, J.; Sabrina, F.; Kwak, J. A Comprehensive Survey of Generative Adversarial Networks (GANs) in Cybersecurity Intrusion Detection. *IEEE Access* **2023**, *11*, 76071–76094. https://doi.org/10.1109/ACCESS.2023.3296707.
92. Cui, F.; Ye, Q.; Kibenge-MacLeod, P. A Wasserstein GAN-based framework for adversarial attacks against intrusion detection systems. In Proceedings of the ICC 2023-IEEE International Conference on Communications, Rome, Italy, 28 May–1 June 2023; pp. 3187–3192. https://doi.org/10.1109/ICC45041.2023.10279233.
93. Xu, C.; Ren, J.; Zhang, D.; Zhang, Y.; Qin, Z.; Ren, K. GANobfuscator: Mitigating Information Leakage Under GAN via Differential Privacy. *IEEE Trans. Inf. Forensics Secur.* **2019**, *14*, 2358–2371. https://doi.org/10.1109/TIFS.2019.2897874.
94. Hu, R.; Li, X.; Zhang, Y.; Chen, W.; Wang, L. CB-GAN: Generate Sensitive Data with a Convolutional Bidirectional Generative Adversarial Networks. In Proceedings of the International Conference on Database Systems for Advanced Applications (DASFAA), Tianjin, China, 17–20 April 2023; Springer Nature: Cham, Switzerland, 2023. https://doi.org/10.1007/978-3-031-30678-5_13.
95. PRISMA. *Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) Statement*. Available online: https://www.prisma-statement.org/ (accessed on 5 October 2025)
96. Brereton, P.; Kitchenham, B.A.; Budgen, D.; Turner, M.; Khalil, M. Lessons from applying the systematic literature review process within the software engineering domain. *J. Syst. Softw.* **2007**, *80*, 571–583. https://doi.org/10.1016/j.jss.2006.07.009.
97. Tseng, H.Y.; Jiang, L.; Liu, C.; Yang, M.H.; Yang, W. Regularizing generative adversarial networks under limited data. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Online, 19–25 June 2021. https://doi.org/10.1109/CVPR46437.2021.00783.
98. Tsalera, E.; Papadakis, A.; Pagiatakis, G.; Samarakou, M. Impact evaluation of sound dataset augmentation and synthetic generation upon classification accuracy. *J. Sens. Actuator Netw.* **2025**, *14*, 91. https://doi.org/10.3390/jsan14050091.
99. Zhang, Z.; Hua, Y.; Sun, G.; Wang, H.; McLoone, S. Improving the Leaking of Augmentations in Data-Efficient GANs via Adaptive Negative Data Augmentation. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), Waikoloa, HI, USA, 2–7 January 2024; pp. 5412–5421.
100. Abdollahzadeh, A.; Xiao, Y.; Wang, Y.; Huang, T.; Patel, V.M.; Metaxas, D. A Survey on Generative Modeling with Limited Data, Few Shots, and Zero Shot. *arXiv* **2023**, arXiv:2307.14397. https://doi.org/10.48550/arXiv.2307.14397.
101. Yamaguchi, S.; Kanai, S.; Kumagai, A.; Chijiwa, D.; Kashima, H. Transfer learning with pre-trained conditional generative models. *arXiv* **2022**, arXiv:2204.12833. https://doi.org/10.48550/arXiv.2204.12833.
102. Song, Y.; Wang, T.; Cai, P.; Mondal, S.K.; Sahoo, J.P. A comprehensive survey of few-shot learning: Evolution, applications, challenges, and opportunities. *ACM Comput. Surv.* **2023**, *55*, 271. https://doi.org/10.1145/3582688.
103. Cobbinah, M.; Nunoo-Mensah, H.; Ebenezer Adjei, P.; Adoma Acheampong, F.; Acquah, I.; Tutu Tchao, E.; Selasi Agbemenu, A.; John Kponyo, J.; Abaidoo, E. Diversity in stable GANs: A systematic review of mode collapse mitigation strategies. *Eng. Rep.* **2025**, *7*, e70209. https://doi.org/10.1002/eng2.70209.
104. Bora, A.; Price, E.; Dimakis, A.G. AmbientGAN: Generative models from lossy measurements. In Proceedings of the International Conference on Learning Representations (ICLR), Vancouver, BC, Canada, 30 April–3 May 2018. Available online: https://openreview.net/forum?id=Hy7fDog0b (accessed on 17 September 2025).
105. Ni, Y.; Koniusz, P. NICE: Noise-modulated consistency regularization for data-efficient GANs. In *Advances in Neural Information Processing Systems*; American Institute of Physics: College Park, MD, USA, 2023; Volume 36, pp. 13773–13801. Available online: https://proceedings.neurips.cc/paper_files/paper/2023/file/2c8047bf3ed8ef6905351608d641f02f-Paper-Conference.pdf (accessed on 27 July 2025).
106. Yang, M.; Wang, Z.; Chi, Z.; Li, D.; Du, W. Adversarial semantic augmentation for training generative adversarial networks under limited data. *arXiv* **2025**, arXiv:2502.00800. https://doi.org/10.48550/arXiv.2502.00800.
107. Liang, W.; Liu, Z.; Liu, C. Dawson: A domain adaptive few-shot generation framework. *arXiv* **2020**, arXiv:2001.00576. https://doi.org/10.48550/arXiv.2001.00576.

108. Phaphuangwittayakul, A.; Guo, Y.; Ying, F. Fast adaptive meta-learning for few-shot image generation. *IEEE Trans. Multimed.* **2022**, *24*, 2205–2217. https://doi.org/10.1109/TMM.2021.3077729.
109. Gu, B.; Zhai, J. Few-shot image generation based on meta-learning and generative adversarial network. *Signal Process. Image Commun.* **2025**, *137*, 117307. https://doi.org/10.1016/j.image.2025.117307.
110. Hospedales, T.; Antoniou, A.; Micaelli, P.; Storkey, A. Meta-learning in neural networks: A survey. *IEEE Trans. Pattern Anal. Mach. Intell.* **2021**, *44*, 5149–5169. https://doi.org/10.1109/TPAMI.2021.3079209.
111. Elaraby, N.; Barakat, S.; Rezk, A. A conditional GAN-based approach for enhancing transfer learning performance in few-shot HCR tasks. *Sci. Rep.* **2022**, *12*, 16271. https://doi.org/10.1038/s41598-022-20654-1.
112. Parnami, A.; Lee, M. Learning from few examples: A summary of approaches to few-shot learning. *arXiv* **2022**, arXiv:2203.04291. https://doi.org/10.48550/arXiv.2203.04291.
113. Li, X.; Yang, X.; Ma, Z.; Xue, J.H. Deep metric learning for few-shot image classification: A review of recent developments. *Pattern Recognit.* **2023**, *138*, 109381. https://doi.org/10.1016/j.patcog.2023.109381.
114. Jiang, W.; Huang, K.; Geng, J.; Deng, X. Multi-scale metric learning for few-shot learning. *IEEE Trans. Circuits Syst. Video Technol.* **2021**, *31*, 1091–1102. https://doi.org/10.1109/TCSVT.2020.2995754.
115. Zhang, Y.; Yuan, X.; Luo, L.; Yang, Y.; Zhang, S.; Xu, C. Local contrast learning for one-shot learning. *Appl. Sci.* **2024**, *14*, 5217. https://doi.org/10.3390/app14125217.
116. Düzyel, O. A comparative study of GAN-generated handwriting images and MNIST images using t-SNE visualization. *arXiv* **2023**, arXiv:2305.09786. https://doi.org/10.48550/arXiv.2305.09786.
117. Sung, F.; Yang, Y.; Zhang, L.; Xiang, T.; Torr, P.H.; Hospedales, T.M. Learning to compare: Relation network for few-shot learning. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 18–22 June 2018. https://doi.org/10.1109/CVPR.2018.00131.
118. Santoro, A.; Bartunov, S.; Botvinick, M.; Wierstra, D.; Lillicrap, T. Meta-learning with memory-augmented neural networks. In Proceedings of the International Conference on Machine Learning PMLR, New York, NY, USA, 19–24 June 2016. Available online: https://proceedings.mlr.press/v48/santoro16.html (accessed on 18 May 2025).
119. Koch, G.; Zemel, R.; Salakhutdinov, R. Siamese neural networks for one-shot image recognition. In Proceedings of the 32nd International Conference on Machine Learning, Lille, France, 6–11 July 2015; pp. 1–6. Available online: https://api.semanticscholar.org/CorpusID:13874643 (accessed on 21 May 2025).
120. Vinyals, O.; Blundell, C.; Lillicrap, T.; Wierstra, D. Matching networks for one-shot learning. *arXiv*, **2016**, arXiv:1606.04080. https://doi.org/10.48550/arXiv.1606.04080.
121. Kaiser, O.; Nachum, O.; Roy, A.; Bengio, S. Learning to remember rare events. *arXiv* **2017**, arXiv:1703.03129. https://doi.org/10.48550/arXiv.1703.03129.
122. Edwards, H.; Storkey, A. Towards a neural statistician. *arXiv* **2016**, arXiv:1606.02185. https://doi.org/10.48550/arXiv.1606.02185.
123. Munkhdalai, T.; Yu, H. Meta networks. In Proceedings of the 34th International Conference on Machine Learning (ICML), Sydney, NSW, Australia, 6–11 August 2017; PMLR: Cambridge, MA, USA, Volume 70, pp. 2554–2563. Available online: https://proceedings.mlr.press/v70/munkhdalai17a.html (accessed on 30 May 2025).
124. Snell, J.; Swersky, K.; Zemel, R.S. Prototypical networks for few-shot learning. *arXiv* **2017**, arXiv:1703.05175. https://doi.org/10.48550/arXiv.1703.05175.
125. Finn, C.; Abbeel, P.; Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. *arXiv* **2017**, arXiv:1703.03400. https://doi.org/10.48550/arXiv.1703.03400.
126. Ali-Gombe, A.; Elyan, E.; Savoye, Y.; Jayne, C. Few-shot classifier GAN. In Proceedings of the 2018 International Joint Conference on Neural Networks (IJCNN), Rio de Janeiro, Brazil, 8–13 July 2018; pp. 1–8. https://doi.org/10.1109/IJCNN.2018.8489387.
127. Hu, J.; Qi, Z.; Wei, J.; Chen, J.; Bao, R.; Qiu, X. Few-shot learning with adaptive weight masking in conditional GANs. In Proceedings of the International Conference on Electronics and Devices, Computational Science (ICEDCS), Marseille, France, 23–25 September 2024; pp. 435–439. https://doi.org/10.1109/ICEDCS64328.2024.00083.
128. Kim, S.; Kang, K.; Kim, G.; Baek, S.H.; Cho, S. DynaGAN: Dynamic few-shot adaptation of GANs to multiple domains. In Proceedings of the SIGGRAPH Asia 2022 Conference Papers, Daegu, Republic of Korea, 6–9 December 2022; pp. 1–8. https://doi.org/10.1145/3550469.3555416.
129. Yang, M.; Wang, Z.; Chi, Z.; Zhang, Y. FreGAN: Exploiting frequency components for training GANs under limited data. *arXiv* **2022**, arXiv:2210.05461. https://doi.org/10.48550/arXiv.2210.05461.

130. Shrivastava, N.; Hanif, M.A.; Mittal, S.; Sarangi, S.R.; Shafique, M. A survey of hardware architectures for generative adversarial networks. *J. Syst. Archit.* **2021**, *118*, 102227. https://doi.org/10.1016/j.sysarc.2021.102227.
131. Thomas, A. Memristor-based neural networks. *J. Phys. D Appl. Phys.* **2013**, *46*, 093001. https://doi.org/10.1088/0022-3727/46/9/093001.
132. Sun, K.; Chen, J.; Yan, X. The future of memristors: Materials engineering and neural networks. *Adv. Funct. Mater.* **2021**, *31*, 2006773. https://doi.org/10.1002/adfm.202006773.
133. Aguinaldo, A.; Chiang, P.Y.; Gain, A.; Patil, A.; Pearson, K.; Feizi, S. Compressing GANs using knowledge distillation. *arXiv* **2019**, arXiv:1902.00159. https://doi.org/10.48550/arXiv.1902.00159.
134. Wang, H.; Gui, S.; Yang, H.; Liu, J.; Wang, Z. GAN slimming: All-in-one GAN compression by a unified optimization framework. In Proceedings of the European Conference on Computer Vision, Glasgow, Scotland, 24–27 August 2020; Springer: Cham, Switzerland, 2020; pp. 54–73. https://doi.org/10.1007/978-3-030-58548-8_4.
135. Sun, H.; Zhu, T.; Zhang, Z.; Jin, D.; Xiong, P.; Zhou, W. Adversarial Attacks Against Deep Generative Models on Data: A Survey. *IEEE Trans. Knowl. Data Eng.* **2023**, *35*, 3367–3388. https://doi.org/10.1109/TKDE.2021.3130903.
136. Mukherjee, S.; Xu, Y.; Trivedi, A.; Ferres, J.L. privGAN: Protecting GANs from membership inference attacks at low cost. *arXiv* **2019**, arXiv:2001.00071. https://doi.org/10.48550/arXiv.2001.00071.
137. Shateri, M.; Messina, F.; Labeau, F.; Piantanida, P. Preserving privacy in GANs against membership inference attack. *IEEE Trans. Inf. Forensics Secur.* **2024**, *19*, 1728–1743. https://doi.org/10.1109/TIFS.2023.3342654.
138. Theodoridis, S.; Koutroumbas, K. *Pattern Recognition*; Elsevier: Amsterdam, The Netherlands, 2006.
139. Sangeetha, S.; Shriaarthy, E.; Rithvik Pranao, N.; Priya, J.S.; Yogeswari, L.; Gokul, S.; Nandha, S.S. Elevating privacy: A differential privacy infused approach to GAN for robust data synthesis for deep learning models. In *Advanced Computing and Communications: Responsible AI, Proceedings of the 29th International Conference ADCOM, Bangalore, India, 18–20 December 2024*; Springer: Cham, Switzerland, 2025; Volume 2461. https://doi.org/10.1007/978-3-031-96473-2_19.
140. Liu, E.; Chu, Z.; Zhang, X. Wasserstein GAN for moving differential privacy protection. *Sci. Rep.* **2025**, *15*, 1–13. https://doi.org/10.1038/s41598-025-03178-2.
141. Tseng, B.-W.; Wu, P.-Y. Compressive privacy generative adversarial network. *IEEE Trans. Inf. Forensics Secur.* **2020**, *15*, 2499–2513. https://doi.org/10.1109/TIFS.2020.2968188.
142. Aldhaheri, S.; Alhuzali, A. SGAN-IDS: Self-attention-based generative adversarial network against intrusion detection systems. *Sensors* **2023**, *23*, 7796 https://doi.org/10.3390/s23187796.
143. Charlier, J.; Lagrue, S.; Francois, J.; Engel, T. SynGAN: Towards Generating Synthetic Network Attacks Using GANs. *arXiv* **2019**, arXiv:1908.09899. https://doi.org/10.48550/arXiv.1908.09899.