https://doi.org/10.31449/inf.v50i15.11996

Generative Adversarial Networks for Primary Healthcare Planning: A Systematic Review of Synthetic Data Generation and Budget Forecasting

I Komang Sugiartha, Adang Suhendra, and I Made Wiryana
Information Technology, Gunadarma University, Jl. Margonda Raya 100, Depok, Indonesia
E-mail: sugiartha@staff.gunadarma.ac.id, asuhendra@staff.gunadarma.ac.id, mwiryana@staff.gunadarma.ac.id

Keywords: Generative adversarial networks, GANs, healthcare, planning

Received: September 20, 2025

Primary healthcare systems, particularly in low and middle income countries, are caused by insufficient data, constraints on resources, and shifts in the demand for services. The focus of this paper is on a systematic literature review to understand the impact of Generative Adversarial Networks (GANs) in the area of budget and planning in primary healthcare systems. Peer-reviewed literature was used for the analysis, and 35 studies published in the period of 2017 to 2025 were selected and reviewed based on the inclusion criteria of relevance to the area of forecasting, simulation, and decision support in healthcare data. Some of the key variations of GANs that are discussed in the paper involve the use of Conditional GANs, Conditional Tabular GANs (CTGANs) and other recurrent or temporal based GANs, which are commonly used for imaging, and electronic health record and time series data in healthcare. The generated findings suggest that, in resource-limited healthcare contexts, the use of GAN- based synthetic data generation improves the flexibility and consistency of predictive models, facilitates resource planning through scenario-based simulation, and allows for the controlled sharing of data across collaborating institutions. However, the review findings point to several gaps that include the absence of operational use in primary healthcare budgeting, difficulties in interpretability and governance of models, and no agreed-upon criteria for evaluating the trade-off between the reliability of synthetic data and associated privacy concerns. Therefore, the review demonstrates the possibilities and constraints of GAN-based methods and delineates paths for subsequent research to foster the integration of these methods into primary healthcare budgeting processes.

Povzetek: Pregledni članek kaže, da GAN-i v primarnem zdravstvu obetajo pri generiranju sintetičnih podatkov za napovedovanje, simulacije in načrtovanje virov, vendar njihovo uporabo pri dejanskem proračunskem odločanju še omejujejo težave z razložljivostjo, upravljanjem in presojo razmerja med zasebnostjo ter zanesljivostjo podatkov.

1 Introduction

Budget planning in the healthcare sector, particularly in developing countries like Indonesia, faces several significant challenges. These challenges can be categorized into various aspects such as financial constraints, resource allocation, and systemic inefficiencies. Limited operational funds, as seen in regions such as Aceh, and the disproportionate allocation of resources favoring curative over preventive care weaken health outcomes and resilience [1, 2]. Misalignment between health policies and budget allocations, along with limited local capacity in decentralized systems, further reduce efficiency and community participation [3]. Complex allocation processes, fragmented funding sources, and inflexible budget structures exacerbate delays, mismanagement, and ineffective resource utilization [2, 4]. Researches show that artificial intelligence and machine learning offer solutions, but are constrained by data availability, budget constraints, and the increasingly complex demands of healthcare services. Artificial intelligence in the field of generative deep learning, such as Generative Adversarial Networks (GANs), has emerged as a promising technology to address these challenges.

GANs promote planning through the augmentation of limited datasets, strengthening model robustness for patient load forecasting and their requirements [5, 6, 7]. Synthetic data generation also allows secure data sharing across institutions, preserving privacy at the cost of not losing analytical value [8, 9, 10]. These capabilities facilitate collaborative research and evidence based policymaking. Through it all, GANs lead to healthcare service planning with fewer uncertainties and greater fairness.

Apart from data synthesis, GANs are also used to mimic varied planning situations, like funding fluctuations or disease pandemics [5, 11]. That ability benefits healthcare managers by maximizing the allocation of limited assets in uncertain situations. Personalized care planning and treatment of chronic disease are also facilitated by GANs through patient oriented simulation [5, 6, 11]. Diagnostic precision is further increased through image quality and segmentation [5, 7, 12].

Although promising, application of GANs is not without challenges, ranging from data quality concerns to model interpretability to training data bias [5, 6]. Moreover, high computational cost and limited technological capabilities can discourage uptake in low resource environments [5, 11]. Furthermore, safety and ethical issues pertaining to synthetic data must also be reconciled [5, 13, 14]. Despite their demonstrated potential, the application of GANs in healthcare remains challenged by issues related to data quality, model interpretability, computational complexity, and ethical considerations. Moreover, high computational cost and limited technological capabilities can discourage uptake in low resource environments [5, 11]. Furthermore, safety and ethical issues pertaining to synthetic data must also be reconciled [5, 13, 14]. These challenges must be addressed to support the safe integration of GANs into budget planning in healthcare systems. However, previous review studies have focused more on the application of GANs to clinical problems, medicine, patients, and biomedical imaging, such as in skin lesion analysis [15], tabular healthcare data generation with privacy concerns [16], ethical implications in medical practice [17], patient feedback optimization [18], high quality ultrasound imaging and reconstruction [19], and biomedical image and signal analysis [20]. To address this gap, this review differs from prior studies by focusing specifically on the role of GAN-based synthetic data generation, forecasting, and simulation in healthcare budget planning, rather than on clinical diagnosis or biomedical imaging applications that dominate existing reviews.

This review aims to provide a comprehensive overview of the potential applications, benefits, and challenges of using GANs in budget planning, specifically issues of data availability and quality, interpretability, and ethics. The goal is to provide insight into how this emerging technology can support more adaptive, efficient, and data driven budget planning in community health centers. A novelty of this study is that it provides an overview of how previous studies have generated high quality synthetic data, enhanced predictive analytics, and simulated various financial scenarios, while maintaining data privacy in the community health center. This overview has not been presented in previous GAN reviews.

2 Research method

This study adopts a systematic literature review approach to explore the potential applications and challenges of implementing Generative Adversarial Networks (GANs) in budget planning for primary healthcare centers. The methodology aims to identify existing research trends, practical use cases, and knowledge gaps related to the integration of GANs in the healthcare budget planning domain [21, 22].

2.1 Literature search strategy

A comprehensive literature search was conducted to ensure a robust foundation of existing studies and methodologies relevant to the intersection of Generative Adversarial Networks (GANs), healthcare, and public health budget planning. The search was performed across multiple reputable academic databases, namely IEEE Xplore, PubMed, Scopus, SpringerLink, and ScienceDirect, due to their respective strengths in covering interdisciplinary scientific literature.

IEEE Xplore was selected for its authoritative coverage of advanced computational methods, machine learning, and artificial intelligence, particularly in engineering and applied computer science. Since the core method in this study involves GANs, IEEE Xplore provides access to highly cited technical papers and proceedings on neural network architectures and AI model developments.

PubMed is a leading database for biomedical literature, maintained by the U.S. National Library of Medicine. It was included to gather health related applications of AI, including studies on synthetic data generation for medical records and healthcare decision making. This is particularly useful for integrating GANs into public health or clinical settings.

Scopus offers comprehensive indexing of multidisciplinary research, including computer science, health sciences, and social sciences. Its inclusion ensured wide coverage of peer reviewed journal articles, especially those discussing the applications of AI in public sector decision making, healthcare systems, and data generation.

SpringerLink was chosen for its extensive repository of journals and books related to applied sciences and engineering, with relevance to emerging technologies in healthcare systems and AI based modeling techniques like GANs.

ScienceDirect, by Elsevier, was included for its strong collection of peer reviewed literature in both computational sciences and health economics. It serves as a reliable source for studies involving AI based policy planning, synthetic data validation, and optimization in healthcare.

To enhance reproducibility, representative search strings were explicitly defined for each database. For example, in Scopus and IEEE Xplore, the query ("Generative Adversarial Networks" OR "GANs") AND ("synthetic data" OR "health data") AND ("budget planning") AND ("healthcare" OR "primary care" OR "public health") was applied to titles, abstracts, and keywords. Comparable adaptations were used for PubMed and ScienceDirect to align with database-specific indexing. The initial search across all databases yielded 330 records prior to duplicate removal, with the number of retrieved articles per database recorded and summarized during the screening process.

Figure 1 explains the stages of collecting and selecting included studies. The first stage begins with collecting studies from various databases and predetermined search strings. Then, the studies are filtered by applying inclusion and exclusion criteria, such as papers written in English and journal articles published between 2017 and 2025.

Figure 1 illustrates the systematic review process for exploring the application of Generative Adversarial Networks (GANs) in budget planning. The process begins with

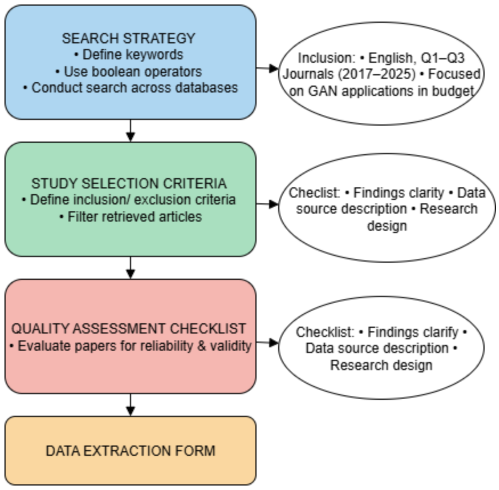

Figure 1: Stages of the systematic literature review process

the search strategy, where researchers define relevant keywords, use boolean operators, and conduct comprehensive searches across academic databases. This stage also includes setting inclusion criteria specifically, selecting English language articles published in Q1–Q3 journals between 2017 and 2025 that are focused on GAN applications in budgeting.

Next, the study selection criteria are applied. This involves defining clear inclusion and exclusion rules and filtering retrieved articles accordingly. Each selected article is evaluated using a checklist that includes the clarity of findings, a description of the data source, and the appropriateness of the research design.

Following selection, a quality assessment checklist is employed to ensure the reliability and validity of the studies. This stage uses predefined quality indicators such as transparency in findings, adequacy of data description, and robustness of the research methodology.

Once quality assessment is complete, relevant information is compiled using a data extraction form. This form is used to systematically collect and organize data from the selected studies to support the review objectives.

This structured approach ensures that the review is comprehensive, replicable, and yields meaningful insights into the role of GANs in budget planning.

To ensure relevance, a combination of Boolean operators and keyword groups was used:

- ”Generative Adversarial Networks” AND ”healthcare”
- ”GANs” AND ”primary care”
- ”Synthetic data” AND ”health data”
- ”GANs” AND ”budget planning”
- ”AI in public health”

This strategy aimed to capture both technical advancements in GAN models and their contextual applications within public health systems, particularly in resource constrained environments like primary healthcare units.

To reduce selection bias, several authors conducted the screening and review processes independently. Evaluations of titles, abstracts, and full texts were completed using established inclusion and exclusion criteria. Reviewer discussions and consensus established methodology and firm criteria to the selection process in case of differences in the study's eligibility or classification.

2.2 Inclusion criteria

The inclusion criteria were carefully defined to ensure that the selected literature was both relevant and of high quality, with a focus on the intersection of Generative Adversarial Networks (GANs) and healthcare applications, particularly in data synthesis, prediction, and budget planning.

- The search included peer reviewed publications from 2017 to 2025 to capture both foundational and emerging research in this area.
- Peer reviewed journal articles, conference proceedings, and technical reports.
- Studies discussing the application of GANs in healthcare data synthesis, prediction, planning, and decision making [23].
- Articles related to resource allocation, budget optimization, and data driven healthcare management using GANs [24].

2.3 Selection and review process

The selection and review process followed the PRISMA 2020 (Preferred Reporting Items for Systematic Reviews and Meta Analyses) guidelines, which offer a structured and standardized method for conducting and reporting systematic reviews [25]. This process ensured methodological transparency, minimized selection bias, and improved the reproducibility of the literature review. The complete process is illustrated in Figure 2.

From an initial pool of 330 records, 290 unique studies were screened after duplicate removal. Following title, abstract, and full text assessments, 35 studies met all inclusion criteria. These studies form the basis of the qualitative synthesis, providing key evidence on the role of GANs in healthcare data generation, prediction, planning, and resource optimization.

This structured and systematic approach, as visualized in Figure 2, ensured the inclusion of high quality, contextually relevant literature to support the research objectives and to inform the development of GAN based solutions in healthcare budget planning and decision making.

The selection process involved multiple stages:

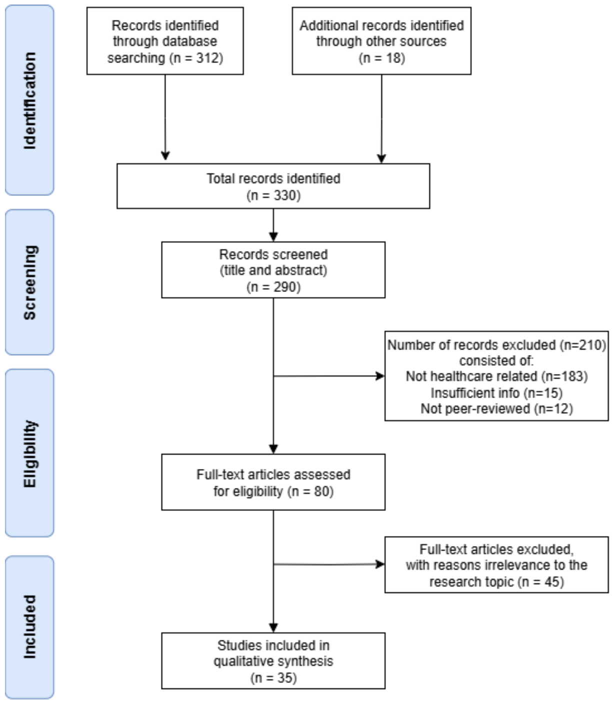

Figure 2: Overview of the study selection process following PRISMA 2020 guidelines

1. **Title and Abstract Screening**: All retrieved papers were screened based on relevance to the research topic [26].
2. **Full Text Review**: Relevant papers were read in full to extract detailed insights into methodologies, use cases, data types, GAN variants, and evaluation metrics [27].
3. **Categorization**: Selected papers were classified into four main thematic areas:
   – Synthetic Financial Data Generation for Budget Forecasting.
   – Scenario Based Simulation and Budget Stress Testing.
   – Cost Prediction Models with GAN Augmented Data.
   – Anomaly Detection in Financial Transactions.

Besides the selection process, an additional evaluative framework was established to assess the methodological strength of the included studies using a qualitative quality framework. Each study was evaluated across four criteria: (i) the study's data accessibility and transparency, (ii) the methodological rigor of the GAN architecture and training process, (iii) the strength of the validation, including the absence or presence of some quantitative performance metrics or downstream utility tests, and (iv) the reproducibility based on the clarity of the experimental setup and reporting. These criteria were constructed to provide a qualitative assessment of each study rather than providing a numerical score. This, in turn, helped to enhance the reliability of the review's findings, even though a particular strict threshold is avoided.

3 Results

This section gives an in depth examination of the findings and an exhaustive discourse. It starts with a broad overview of how the research articles are divided and continues with an extended examination of the findings on every research query.

3.1 Description of studies

As shown in Figure 3, we identified from the preliminary scrutiny that most articles, about 37.1% came from Scopus itself, thereby demonstrating its overarching position as a broad and multi disciplinary database in facilitating research through this review. IEEE Xplore was also the runner up and was responsible for 28.6% of the overall articles, and this reveals very strong applicability of engineering and computer science fields, especially on research involving artificial intelligence and healthcare tech.

PubMed, a widely used repository in the medical and health sciences domain, contributed 17.1%, showing a significant presence of medically focused literature. Meanwhile, ScienceDirect accounted for 11.4%, representing articles primarily from Elsevier journals, which span both health and technology fields. Lastly, SpringerLink contributed the smallest share, 5.7% yet still offered relevant insights across interdisciplinary research.

The review comprises scholarly articles between 2017 and 2025 and covers a broad time spectrum such that historical and recent work is included. The coverage is a balanced representation of technical and field specific literature, which is critical for an informative comprehension of the developing interface between public health and artificial intelligence.

3.2 Synthetic data generation for improved forecasting

In the context of Synthetic Data Generation for Improved Forecasting, Generative Adversarial Networks (GANs) have generally demonstrated the ability to expand sample size while preserving the underlying data distribution, thereby enabling more stable model training in data constrained settings [28, 29, 30, 31]. These benefits have been consistently observed across diverse domains, including medical imaging, electronic health record (EHR) tabular data, and time series forecasting.

Specifically, in medical imaging, synthetic abnormal brain MRI scans have been shown to improve tumor segmentation performance while simultaneously serving as an anonymization mechanism, thereby facilitating inter institutional data sharing without compromising patient identity [28]. Similarly, in hepatic CT imaging, augmenting

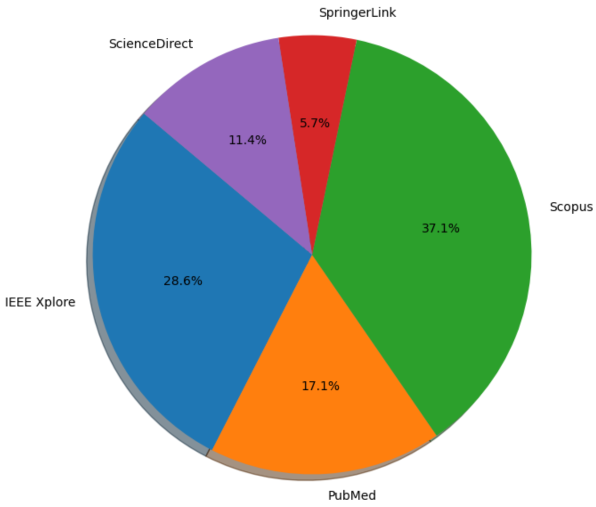

Figure 3: Articles per source

training datasets with GAN generated lesions substantially enhanced CNN sensitivity and specificity, highlighting the practical utility of synthetic data in scenarios where real samples are scarce [29].

For mixed type tabular data, Conditional Tabular GANs (CTGANs) effectively reproduced numerical categorical structures with low duplication and retained column wise statistics critical for downstream predictive tasks, such as length of stay forecasting [30]. In temporal applications, Recurrent/Conditional GANs successfully passed "train on synthetic, test on real" evaluations, demonstrating that synthetic sequences can serve as viable substitutes for limited observations in forecasting pipelines [31].

Moreover, methodological surveys have emphasized the importance of combining fidelity metrics (e.g., KS, MMD, JSD) with downstream utility and privacy assessments prior to deploying synthetic data within forecasting workflows [32].

Across the reviewed studies, the effectiveness of GAN-based synthetic data generation is commonly supported by quantitative evaluation metrics. Reported assessments include statistical fidelity measures such as Kolmogorov-Smirnov distance, Maximum Mean Discrepancy, and Jensen-Shannon divergence, which generally indicate closer alignment between synthetic and real data distributions compared to baseline augmentation methods. In forecasting and classification tasks, several studies further report improvements in downstream model performance, as measured by accuracy, precision-recall balance, or reduced generalization error, when training is augmented with GAN-generated data.

A comparative analysis across the reviewed studies indicates that the performance of GAN variants is highly dependent on data modality and evaluation objectives. Conditional GANs and CTGANs consistently demonstrate stronger performance on mixed-type tabular healthcare data by preserving conditional distributions and rare-event structures, while recurrent and temporal GAN variants are more effective for longitudinal forecasting tasks. In contrast, convolutional GAN architectures achieve higher spatial fidelity in imaging applications but often require larger datasets and greater computational resources. Evaluation practices commonly combine statistical fidelity metrics, such as Kolmogorov-Smirnov distance, Jensen-Shannon divergence, and Maximum Mean Discrepancy, with downstream utility assessments and privacy risk evaluations, forming a comprehensive framework for assessing synthetic data quality.

3.3 Scenario based simulation for resilience testing

Previous studies have extensively discussed Scenario Based Simulation for Resilience Testing, which essentially leverages conditional generation to enable credible "what if" simulations. By synthesizing trajectories and distributions under alternative operational assumptions, this approach allows systems to be stress tested against a variety of plausible shocks and extreme conditions [31].

In the field of neuroimaging, for instance, GAN based models have been applied to project glioma growth over time, achieving high levels of agreement with real observations as measured by overlap metrics such as Dice and Jaccard. These results demonstrate how generative models can effectively emulate evolving system states, offering valuable insights for resilience analysis [33]. Similarly, Recurrent Conditional GANs (RCGANs) have been employed as simulators of ICU dynamics, a framework that can be directly transferred to resource planning and budgeting contexts where sampling from realistic yet shifted distributions is required [31].

At a broader level, reviews have highlighted the ability of medical GANs to address domain shifts such as changes in imaging protocols or device configurations which further reinforces their suitability for resilience testing in environments governed by policy and privacy constraints [27]. Taken together, this body of work underscores not only the relevance of conditional generative approaches in medical research but also their potential as practical tools for operational planning and risk management across diverse domains.

3.4 Predictive models enhanced with augmented data

Growing evidence highlights that GAN based data augmentation offers tangible improvements for predictive modeling, particularly in situations where training data are limited or class distributions are skewed. Overall, synthetic augmentation has been shown to strengthen model robustness and generalization across diverse tasks, including classification, detection, and segmentation [27].

In liver lesion classification, for example, the addition of GAN generated regions of interest markedly improved sensitivity and specificity, surpassing the results of traditional augmentation techniques [29]. A similar effect was reported in brain tumor studies, where synthetic cohorts not only enhanced segmentation accuracy but, in some cases, achieved performance comparable to models trained exclusively on real datasets. These findings underscore the value of synthetic corpora in domains where access to authentic data remains restricted [28].

At the same time, field studies emphasize that the benefits of augmentation vary depending on dataset scale and heterogeneity. This variation calls for careful calibration of augmentation intensity to ensure effectiveness when these practices are transferred to new predictive contexts [33]. Collectively, these insights affirm that GAN based augmentation can be a powerful strategy for improving predictive models, provided that its application is thoughtfully adapted to the target domain.

3.5 Anomaly detection in complex datasets

Recent research has increasingly shown that GAN based approaches provide a promising framework for anomaly detection in complex and high dimensional datasets. By learning the manifold of normal data and flagging deviations, these methods enable principled anomaly identification without the need for exhaustive labeling, making them particularly valuable in domains where annotated data are scarce [33].

One line of work applies inpainting style reconstruction, where masked regions are replaced with healthy counterparts. This strategy has significantly improved unsupervised lesion localization, with Dice scores reported to increase from approximately 38% to 77% through adversarial reconstruction [33]. Beyond these experimental gains, contemporary reviews highlight both the opportunities and the challenges of GAN based anomaly detection. While such methods are especially promising when labels are limited, their performance varies across datasets, emphasizing the need for rigorous evaluation and careful calibration prior to operational deployment [34].

For anomaly detection, quantitative evidence is primarily reported through precision, recall, and reconstruction-based metrics. GAN-based approaches consistently demonstrate higher anomaly localization accuracy and improved recall compared to rule-based or autoencoder baselines, particularly in settings with limited labeled data. In addition, reconstruction error distributions and overlap-based metrics are frequently used to substantiate the ability of GANs to distinguish normal patterns from anomalous behavior, providing measurable support for their effectiveness in complex detection tasks.

Taken together, these findings illustrate that GAN based anomaly detection is not only conceptually sound but also practically effective, provided its application is tailored to the variability of real world data environments.

3.6 Summary of reviewed studies

To provide a structured overview of the reviewed literature and to clarify the research gap addressed by this study, Table 1 summarizes key characteristics of the selected studies. The table groups prior work by GAN architecture, application domain, dataset type, and main findings, highlighting that most existing research focuses on medical imaging and electronic health records, while direct applications of GANs in primary healthcare budget planning and forecasting remain limited.

Table 1: Summary of reviewed studies by GAN type, application domain, and key findings

GAN Type	Application Domain	Dataset Type	Key Findings	Reference
DCGAN / cGAN	Medical Imaging	MRI, CT scans	Improved image quality and segmentation performance; supports data augmentation under limited samples	[28, 29]
CTGAN	Electronic Health Records	Tabular EHR data	Preserves statistical distributions and rare events; improves downstream predictive robustness	[30]
RCGAN	Time-Series Simulation	ICU and clinical time series	Enables realistic temporal simulation and scenario analysis for resource utilization	[31]
cGAN	Anomaly Detection	Medical images, signals	Enhances unsupervised anomaly detection and reconstruction fidelity	[33]
GAN-based Models	Budget Forecasting / Planning	Financial and utilization data	Limited to exploratory or simulation-based studies; no large-scale primary healthcare deployment reported	[23, 31]

3.7 Quantitative performance metrics reported in prior studies

Through the reviewed literature, it can be seen that the effectiveness of the GAN-supported methods primarily come from the reported quantitative metric of the study. In budget forecasting and other predictive modeling tasks, GAN-supported datasets are noted to increase downstream models' accuracy, sensitivity, and specificity as compared to models trained on limited real world datasets. Several of the cited works noted an increase in recall for infrequent and/or critical events which suggests improvement of the models sustained robustness in a data scarce environment.

When assessing the fidelity of synthetic data, authors tend to use statistical measures of similarity for the real and synthetic data distributions which include Kullback-Leibler divergence, Wasserstein distance, the Kolmogorov-Smirnov statistic, and Maximum Mean Discrepancy. Typically, GAN-generated data has less divergence and distance than sampling from a baseline or other methods of noise augmentation. Authors of anomaly detection studies tend to use precision, recall, area under the ROC curve, and reconstruction error to describe the model performances, and in these studies closed-world GAN models are reported to perform better than other models in the case of limited labeled anomalies. Taken collectively, these models describe the effectiveness of GAN methods quantitative in a variety of application fields and a variety of evaluation conditions.

3.8 Research gaps and opportunities

Despite the growing interest in GANs, their use in budgeting and planning particularly for primary healthcare centers remains limited. Much of the work to date is still situated in experimental or simulation settings, with only a few attempts at real world implementation [23, 31].

This gap has led researchers to stress the importance of building stronger bridges between theory and practice. Two key directions often highlighted are the need for closer interdisciplinary collaboration and the design of domain specific GAN architectures that can address the unique challenges of healthcare planning [35].

Together, these insights suggest that while GANs hold strong potential in this area, moving forward will require both technical innovation and practical integration. Advancing from proof of concept studies toward operational, context sensitive applications could open meaningful opportunities for supporting resource allocation and decision making in primary healthcare.

Recent studies have begun to report real-world and pilot implementations of GAN-based systems in healthcare settings, including synthetic electronic health record generation for hospital analytics, GAN-augmented forecasting for resource utilization, and privacy-preserving data sharing across institutions. Although these implementations are primarily reported in hospital-level or national health system contexts, they provide empirical evidence that GAN-based approaches are technically feasible and operationally relevant, offering transferable insights for future adoption in primary healthcare budgeting and planning.

4 Discussion

4.1 Synthetic financial data generation for budget forecasting

In this research, we demonstrate that synthetic data generation can be a practical solution to various challenges in healthcare budget forecasting. Often, access to detailed financial data is limited due to privacy concerns, hampering analysis. With synthetic data, we can generate datasets that resemble real world data but remain anonymous, allowing for forecasting model testing without the risk of leaking sensitive information.

We use Generative Adversarial Networks (GANs) as the technical foundation. This technology mimics the statistical structure of real budgets and expenditures. We found several architectures suitable for different contexts: Conditional Tabular GANs (CTGANs) and stable variants like medBGANs work well for mixed numerical categorical data, while Recurrent Conditional GANs (RCGANs) are particularly useful for capturing temporal patterns in cash flows and healthcare claims. Conditional approaches are also important because they retain rare but crucial events, such as sudden spikes in claims or intensive care cases, which are highly relevant for testing the financial resilience of healthcare services.

To improve practical interpretability and replicability, a representative example from the reviewed literature can be highlighted. For instance, studies employing Conditional Tabular GANs (CTGANs) for healthcare-related financial or utilization data typically model input vectors composed of mixed numerical and categorical features, with latent dimensions ranging from tens to low hundreds. Training commonly follows an adversarial min–max objective using variants of the Wasserstein or cross-entropy loss with gradient penalty to stabilize convergence. Models are trained over multiple epochs with mini-batch updates, while conditional vectors are used to enforce category-level constraints and preserve rare event distributions. Although specific hyperparameter values vary across studies, this configuration reflects a widely adopted architectural pattern for tabular synthetic data generation in healthcare contexts.

Before using this synthetic data, we ensure its thorough validation. We tested statistical goodness of fit to see if the synthetic data distribution resembled real world data, and conducted usability tests to ensure it supported decision making. Privacy was also a key concern. To that end, we implemented techniques such as differential privacy and federated learning, along with additional testing such as membership inference testing to prevent data leakage. We documented all of these steps through data cards or model cards, which record the data's origin, validation results, and privacy settings, allowing for transparent auditing.

Despite the fact that many studies point to the fact that poorly constructed generative models still have the potential to leak sensitive information, while synthetic data may often be presented as a way to avoid using real data from the healthcare system that violates user privacy, DPP GANs and PATE-GANs reduce the potential for user data to be disclosed by incorporating noise and data from an aggregated teacher model, at the expense of a reduction data utility.

In addition to technical safeguards, governance frameworks are vital for responsible deployment. The primary purpose of model and data cards is to increase transparency by recording the source of the data, the assumptions made during training, how the model was validated, and the limitations of the model. When these frameworks are not coupled with formal processes of decision making, auditing and accountability, their utility is greatly diminished. In the case of frameworks for governance in public health safety net budgeting, there is a need to go beyond the descriptive elements to set prescriptive elements that create standards in civil relations about the accountability of institutions and governance regarding the behavior of models and privacy risks.

From this entire process, we concluded that GAN based synthetic data has significant potential to strengthen healthcare budget forecasting. When properly created and maintained, this data not only improves the quality of analysis but also enables collaboration between institutions without compromising the confidentiality of real financial records.

4.2 Scenario based simulation and budget stress testing

Scenario based simulations can be viewed as a more realistic approach to assessing health budget resilience, especially when compared to deterministic spreadsheets that produce only a single predicted value. Their primary advantage is their ability to generate a full distribution of possible outcomes, allowing risks such as deficits, cash shortages, or accelerated spending to be assessed probabilistically, rather than simply through a single point estimate.

To maintain compliance with the policy framework, these simulations are equipped with explicit constraints, such as spending ceilings, benefit schedules, and applicable accounting rules. Robustness tests are also conducted to ensure the model's robustness to structural changes, such as post pandemic service utilization patterns or tariff reforms. The results are then integrated into the Medium Term Expenditure Framework (MTEF) through various visual tools, such as fan charts and stress test tables. An explanatory layer is added to link the variation in results between scenarios to key cost factors healthcare personnel, pharmaceuticals, logistics, and referrals so that the analysis can be directly linked to concrete fiscal response options.

The primary value of this approach lies in its ability to bridge technical analysis with the decision making process. Scenario based simulations not only offer insights into possible futures but also foster cross agency conversations about how to anticipate fiscal risks more transparently and collaboratively. Thus, simulations serve a dual purpose: as analytical tools that generate quantitative evidence, and as communication tools that clarify policy options, consequences, and opportunities for adjustment in health budget planning.

4.3 Cost prediction models with GAN augmented data

Cost modeling often faces challenges due to data limitations or skewed distribution, which can bias estimates toward facilities or services with more data. GAN based data augmentation offers a systematic solution by synthesizing underrepresented segments such as remote clinics, low volume services, or high cost cases so that models can generate more stable estimates of unit costs, claim burdens, and program needs.

To maintain validity, synthetic data must maintain conditional frequencies across policy relevant categories, such as combinations of service types, drug classes, and suppliers. In time series models, seasonal patterns and calendar effects must also be maintained to ensure projections do not lose real world dynamics. In addition to technical quality, rigorous governance is crucial. Champion challenger evaluation, rolling origin backtesting, and utility testing prior to implementation help prevent overfitting and the emergence of synthetic artifacts. Equivalence also needs to be explicitly addressed, for example, by monitoring differences in prediction errors in underserved areas.

Cost prediction results should not only be presented as raw numbers but also triangulated with policy benchmarks. To support transparency, easy to read summaries, such as cost cards, can be used to summarize the range of credible estimates, key drivers, and levels of uncertainty. With this approach, GAN based models not only improve technical accuracy but also strengthen communication with parties involved in fiscal reviews, such as the ministry of finance or budget committees.

4.4 Anomaly detection in financial transactions

Anomaly detection is a crucial aspect of financial transaction monitoring, particularly because anomalous patterns are often difficult to identify using conventional rules. Within an adversarial framework, a discriminator can function as an unsupervised detector that learns to distinguish normal financial behavior across accounts, programs, suppliers, and facilities. This approach allows the system to identify potential anomalies more adaptively than fixed rule based methods.

However, to avoid excessive false positives, the model needs to account for temporal signals such as seasonality and weekday patterns, as well as policy factors like regulatory changes or tariff updates. Furthermore, monitoring concept drift and rescheduling model training are crucial steps to ensure the system remains relevant to the latest dynamics. Sensitivity to collusion patterns can also be enhanced by utilizing graph based features that capture the relationships between accounts and suppliers within a network.

Cross entity collaboration in anomaly detection is essential, but must maintain privacy. Approaches such as federated learning, differential privacy, or the use of non invertible embeddings enable systemic monitoring without the need to expose the raw transaction ledger. At the operational stage, successful detection relies heavily on clear investigative guidelines from the triage process and document verification to targeted field audits. Feedback from adjudicated cases then needs to be fed back into the model to continuously improve precision and recall.

In this way, GAN based anomaly detection not only helps identify irregularities but also builds a continuous learning system that supports fiscal accountability and strengthens public financial governance.

4.5 Comparison with state-of-the-art models

When compared to conventional machine learning techniques regularly applied in budget modeling like linear regression, tree based methods, autoregressive time-series models, etc, methods based on Generative Adversarial Networks (GAN) have considerable adaptability to capture intricate high dimensional data distributions. Conventioanl models use stringent assumptions and could have challenges with sparsity and shifts in distributions. GANs on the other hand, are better equipped to produce plausible synthetic samples which preserve joint dependencies with regards to the multiple financial and service utilization variables. GANs come with training instability and high computational demands which may pose challenges for certain low resource environments.

In certain circumstances GANs may be more effective than other generative models like Variational Autoencoders (VAEs) to owing to the fact that GANs produce more realistic and sharper samples than VAEs which tend to oversmooth outputs and in doing so underestimate budget distribution tail risks. Despite the fact that autoregressive models are suitable for sequential forecasting, they tend to lack the capacity to generate a wide range of synthetic datasets. The trade off in data fidelity and modeling flexibility when Generative Adversarial Networks are used is a feature in synthetic data generation and scenario analyses during primary health care budget planning. The use of other generative models in conjunction with GANs may be a valid path for upcoming research in this area.

5 Conclusion and future research

Generative Adversarial Networks (GANs) hold significant potential to revolutionize primary healthcare planning through their advanced capabilities in data generation, predictive modeling, and personalized treatment planning. In settings where data is scarce or incomplete such as in many community healthcare centers GANs offer promising solutions by synthesizing realistic datasets, simulating complex scenarios, and enhancing diagnostic models. These capabilities can directly support more informed decision making in resource allocation, activity planning, and long term health service strategies.

However, the adoption of GANs in real world primary healthcare environments is still in its early stages. Several challenges remain, including concerns about data privacy and ethics, limited computational resources, and the technical complexity involved in developing and deploying GAN models. In addition to these challenges, the deployment of GANs in public health budgeting should be guided by clear governance frameworks that emphasize accountability, transparency, and auditability in data use and model outputs. Such governance integration is critical to ensure regulatory compliance, institutional trust, and responsible decision-making within decentralized primary healthcare systems. Without addressing these barriers, the full potential of GANs in healthcare planning cannot be realized.

Future efforts should prioritize the development of context-aware and customized GAN solutions tailored to the specific needs of primary healthcare systems, supported by interdisciplinary collaboration among AI developers, clinicians, policymakers, and public health experts. In particular, future research should emphasize the integration of GANs with federated learning and differential privacy to enable collaborative model training across multiple healthcare institutions without requiring centralized data sharing. This combination allows synthetic data generation and predictive modeling to benefit from distributed data sources while preserving patient confidentiality and regulatory compliance. With appropriate technical and governance frameworks in place, such integrated approaches can enhance data availability, support ethical and accountable AI deployment, and ultimately strengthen equitable, efficient, and data-driven primary healthcare planning in decentralized and resource-constrained environments.

References

[1] A. Abdullah, K. Hort, A. Z. Abidin, and F. M. Amin, "How much does it cost to achieve coverage targets for primary healthcare services? A costing model from Aceh, Indonesia," *Int J Health Plann Manage*, vol. 27, no. 3, pp. 226–245, 2012. https://doi.org/10.1002/hpm.2099

[2] A. Fuady et al., "Bridging the gap: financing health promotion and disease prevention in Indonesia," *Health Res Policy Syst*, vol. 22, no. 1, p. 146, 2024. https://doi.org/10.1186/s12961-024-01206-7

[3] B. Tsofa, S. Molyneux, L. Gilson, and C. Goodman, "How does decentralisation affect health sector planning and financial management? a case study of early effects of devolution in Kilifi County, Kenya," *Int J Equity Health*, vol. 16, no. 1, p. 151, 2017. https://doi.org/10.1186/s12939-017-0649-0

[4] T. Moradi, M. J. Kabir, H. Pourasghari, S. J. Ehsanzadeh, and A. Aryankhesal, "Challenges of budgeting and Public Financial Management in Iran's health system: a qualitative study," *Med J Islam Repub Iran*, vol. 37, p. 80, 2023. https://doi.org/10.47176/mjiri.37.80

[5] K. Deeba, D. Vathana, D. Vanusha, and J. Ramaprabha, "Applications of generative adversarial networks (GANs) in healthcare," in *Artificial Intelligence Revolutionizing Cancer Care*, CRC Press, 2025, pp. 177–211.

[6] Y. S. S. M. Saleh, H. Mokayed, K. Nikolaidou, L. Alkhaled, and Y. C. Hum, "How GANs assist in Covid-19 pandemic era: a review," *Multimed Tools Appl*, vol. 83, no. 10, pp. 29915–29944, 2024. https://doi.org/10.1007/s11042-023-16597-y

[7] F. Asadi and J. A. O'Reilly, "Artificial computed tomography images with progressively growing generative adversarial network," in *2021 13th Biomedical Engineering International Conference (BMEiCON)*,

2021, pp. 1–5. https://doi.org/10.1109/bmeicon53485.2021.9745251

[8] G. O. Ghosheh, J. Li, and T. Zhu, "A survey of generative adversarial networks for synthesizing structured electronic health records," *ACM Comput Surv*, vol. 56, no. 6, pp. 1–34, 2024. https://doi.org/10.1145/3636424

[9] S. Rashidian et al., "SMOOTH-GAN: towards sharp and smooth synthetic EHR data generation," in *Artificial Intelligence in Medicine: 18th Int. Conf. on Artificial Intelligence in Medicine (AIME 2020)*, Minneapolis, MN, USA, Aug. 25–28, 2020, pp. 37–48. https://doi.org/10.1007/978-3-030-59137-3_4

[10] P. Jackson and M. Lussetti, "Extending a generative adversarial network to produce medical records with demographic characteristics and health system use," in *2019 IEEE 10th Annual Information Technology, Electronics and Mobile Communication Conf. (IEMCON)*, 2019, pp. 515–518. https://doi.org/10.1109/iemcon.2019.8936168

[11] V. Mukesh, H. Joshi, S. Saraf, and G. Singh, "Review on Biomedical Informatics Through the Versatility of Generative Adversarial Networks," in *Int. Conf. on Computational Intelligence in Data Science*, 2024, pp. 461–474. https://doi.org/10.1007/978-3-031-69986-3_35

[12] A. Boukhamla et al., "GANs investigation for multimodal medical data interpretation: basic architectures and overview," in *2023 Int. Conf. on Control, Automation and Diagnosis (ICCAD)*, 2023, pp. 1–6. https://doi.org/10.1109/iccad57653.2023.10152386

[13] K. K. Dixit, U. S. Aswal, V. Saravanan, M. Sararswat, N. Shalini, and A. Srivastava, "Data augmentation with generative adversarial networks for deep learning in healthcare," in *2023 Int. Conf. on Artificial Intelligence for Innovations in Healthcare Industries (ICAIIHI)*, 2023, pp. 1–6. https://doi.org/10.1109/icaiihi57871.2023.10489462

[14] D. Helen and N. V. Suresh, "Generative AI in Healthcare: Opportunities, Challenges, and Future Perspectives," in *Revolutionizing the Healthcare Sector with AI*, pp. 79–90, 2024. https://doi.org/10.4018/979-8-3693-3731-8.ch004

[15] A. Ravikumar, H. Sriraman, C. Chadha, and V. K. Chattu, "Alleviation of health data poverty for skin lesions using ACGAN: systematic review," *IEEE Access*, vol. 12, pp. 122702–122723, 2024. https://doi.org/10.1109/access.2024.3417176

[16] J. Coutinho-Almeida, P. P. Rodrigues, and R. J. Cruz-Correia, "GANs for tabular healthcare data generation: A review on utility and privacy," in *Int. Conf. on Discovery Science*, 2021, pp. 282–291. https://doi.org/10.1007/978-3-030-88942-5_22

[17] A. Gunawan and R. Wiputra, "Ethical Implications of Generative AI in Medicine: A Systematic Review of Current Practices and Future Directions," in *2024 IEEE 12th Conf. on Systems, Process & Control (ICSPC)*, 2024, pp. 239–244. https://doi.org/10.1109/icspc63060.2024.10862569

[18] M. Fahim-Ul-Islam, A. Chakrabarty, M. Hasan, A. H. Efaz, X. Wang, and M. J. Piran, "Optimizing Patient Feedback with Generative Adversarial Network Leveraging Knowledge Distillation to Improve Healthcare," *IEEE J Biomed Health Inform*, 2025. https://doi.org/10.1109/jbhi.2025.3584240

[19] Y. Li et al., "Review of high-quality ultrasound imaging and reconstruction," Editorial and Publishing Board of JIG, Jun. 1, 2024. [doi:10.11834/jig.240006](doi:10.11834/jig.240006)

[20] F. De Marco et al., "AI-based solutions for the analysis of biomedical images and signals," *Ital-IA*, pp. 171–176, 2023.

[21] I. J. Goodfellow et al., "Generative adversarial nets," *Adv Neural Inf Process Syst*, vol. 27, 2014. https://doi.org/10.48550/arXiv.1406.2661

[22] G. A. Kaissis, M. R. Makowski, D. Rückert, and R. F. Braren, "Secure, privacy-preserving and federated machine learning in medical imaging," *Nat Mach Intell*, vol. 2, no. 6, pp. 305–311, 2020. https://doi.org/10.1038/s42256-020-0186-1

[23] E. Choi, S. Biswal, B. Malin, J. Duke, W. F. Stewart, and J. Sun, "Generating multi-label discrete patient records using generative adversarial networks," in *Machine Learning for Healthcare Conference*, 2017, pp. 286–305. https://doi.org/10.48550/arXiv.1703.06490

[24] M. K. Baowaly, C.-C. Lin, C.-L. Liu, and K.-T. Chen, "Synthesizing electronic health records using improved generative adversarial networks," *J Am Med Inform Assoc*, vol. 26, no. 3, pp. 228–241, 2019. https://doi.org/10.1093/jamia/ocy142

[25] M. J. Page et al., "The PRISMA 2020 statement: an updated guideline for reporting systematic reviews," *BMJ*, vol. 372, 2021. https://doi.org/10.31222/osf.io/v7gm2

[26] D. Moher, A. Liberati, J. Tetzlaff, and D. G. Altman, "Preferred reporting items for systematic reviews and meta-analyses: the PRISMA statement," *BMJ*, vol. 339, 2009. https://doi.org/10.1371/journal.pmed.1000097

[27] X. Yi, E. Walia, and P. Babyn, "Generative adversarial network in medical imaging: A review," *Med Image Anal*, vol. 58, p. 101552, 2019. https://doi.org/10.1016/j.media.2019.101552

[28] H.-C. Shin et al., "Medical image synthesis for data augmentation and anonymization using generative adversarial networks," in *Simulation and Synthesis in Medical Imaging (SASHIMI 2018)*, Granada, Spain, Sept. 16, 2018, pp. 1–11. https://doi.org/10.1007/978-3-030-00536-8_1

[29] M. Frid-Adar, I. Diamant, E. Klang, M. Amitai, J. Goldberger, and H. Greenspan, "GAN-based synthetic medical image augmentation for increased CNN performance in liver lesion classification," *Neurocomputing*, vol. 321, pp. 321–331, 2018. https://doi.org/10.1016/j.neucom.2018.09.013

[30] D. Bietsch, R. Stahlbock, and S. Voß, "Electronic Health Data in the Context of Patient Length-of-Stay Prediction: Using Generative Adversarial Nets for Synthetic Data Creation," in *2023 Congress in Computer Science, Computer Engineering, & Applied Computing (CSCE)*, 2023, pp. 1597–1604. https://doi.org/10.1109/csce60160.2023.00262

[31] C. Esteban, S. L. Hyland, and G. Rätsch, "Real-valued (medical) time series generation with recurrent conditional GANs," *arXiv preprint arXiv:1706.02633*, 2017. https://doi.org/10.48550/arXiv.1706.02633

[32] V. C. Pezoulas et al., "Synthetic data generation methods in healthcare: A review on open-source tools and methods," *Comput Struct Biotechnol J*, vol. 23, pp. 2892–2910, 2024. https://doi.org/10.1016/j.csbj.2024.07.005

[33] R. Wang et al., "Applications of generative adversarial networks in neuroimaging and clinical neuroscience," *Neuroimage*, vol. 269, p. 119898, 2023. https://doi.org/10.1016/j.neuroimage.2023.119898

[34] S. Kazeminia et al., "GANs for medical image analysis," *Artificial Intelligence in Medicine*, vol. 110, p. 101938, 2020. doi:10.1016/j.artmed.2020.101938. https://doi.org/10.1016/j.artmed.2020.101938

[35] A. Goncalves, P. Ray, B. Soper, J. Stevens, L. Coyle, and A. P. Sales, "Generation and evaluation of synthetic patient data," *BMC Med Res Methodol*, vol. 20, pp. 1–40, 2020. https://doi.org/10.1186/s12874-020-00977-1