



| RESEARCH ARTICLE

Article DOI: 10.21474/JNCS01/127

DOI URL: <http://dx.doi.org/10.21474/JNCS01/127>

PRIVACY-PRESERVING SYNTHETIC DATA GENERATION USING GENERATIVE ARTIFICIAL INTELLIGENCE

Aarav Mehta, Elena Petrova and Mohammed Tariq Hassan

Corresponding Author: Aarav Mehta

| ABSTRACT

The rapid advancement of Artificial Intelligence (AI) has transformed the way organizations collect, analyze, and utilize data. However, concerns regarding data privacy, legal regulations, and ethical considerations have limited access to high-quality datasets for research and industrial applications. Generative Artificial Intelligence (Generative AI) offers an innovative solution through synthetic data generation, enabling realistic datasets to be produced without exposing sensitive personal information. Synthetic data preserves statistical characteristics of real-world data while reducing the risk of privacy violations. This paper explores privacy-preserving synthetic data generation techniques based on Generative AI models, including Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs), Diffusion Models, and Large Language Models (LLMs).

| KEYWORDS

Generative AI, Synthetic Data, Privacy Preservation, Machine Learning, Data Security, GANs, Differential Privacy

| ARTICLE INFORMATION

RECEIVED: 24 February 2026

ACCEPTED: 30 March 2026

PUBLISHED: April 2026

Abstract:-

The rapid advancement of Artificial Intelligence (AI) has transformed the way organizations collect, analyze, and utilize data. However, concerns regarding data privacy, legal regulations, and ethical considerations have limited access to high-quality datasets for research and industrial applications. Generative Artificial Intelligence (Generative AI) offers an innovative solution through synthetic data generation, enabling realistic datasets to be produced without exposing sensitive personal information. Synthetic data preserves statistical characteristics of real-world data while reducing the risk of privacy violations. This paper explores privacy-preserving synthetic data generation techniques based on Generative AI models, including Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs), Diffusion Models, and Large Language Models (LLMs). The study discusses methodologies, privacy-preserving mechanisms, evaluation metrics, practical applications, challenges, and future research directions. The findings indicate that Generative AI has the potential to revolutionize secure data sharing and AI model development while supporting regulatory compliance.

Introduction:-

Data has become one of the most valuable assets in the digital economy. Healthcare institutions, financial organizations, government agencies, and technology companies rely heavily on large datasets for machine learning and predictive analytics. However, sharing sensitive data often violates privacy laws and exposes individuals to identity theft and data breaches. Traditional anonymization techniques frequently fail to protect privacy because attackers can reconstruct identities through auxiliary information. Consequently, researchers have explored synthetic data generation as an alternative. Synthetic datasets imitate the statistical properties of original datasets without directly revealing personal information. Recent developments in Generative Artificial Intelligence have significantly improved synthetic data quality. Deep learning

architectures can now generate realistic images, text, medical records, financial transactions, and sensor data. These capabilities make Generative AI an important tool for privacy-preserving machine learning. This paper reviews modern Generative AI techniques for synthetic data generation and evaluates their role in enhancing privacy, security, and data accessibility.

Literature Review:-

Several studies have investigated synthetic data generation using deep learning techniques. Goodfellow et al. introduced Generative Adversarial Networks (GANs), enabling realistic image synthesis through adversarial learning. Kingma and Welling proposed Variational Autoencoders (VAEs), providing probabilistic latent representations suitable for data generation. Recent diffusion models have achieved state-of-the-art image generation by gradually denoising random noise into realistic samples. Similarly, Large Language Models have demonstrated exceptional capabilities in generating synthetic textual datasets. Researchers have integrated Differential Privacy with generative models to mathematically guarantee privacy protection while maintaining data utility. The literature suggests that privacy-preserving synthetic data has become an active research area due to increasing regulatory requirements and the demand for trustworthy AI systems.

Synthetic Data Generation Using Generative AI:-

Synthetic data refers to artificially generated information that resembles real-world data. Instead of copying original records, generative models learn the probability distribution of the data and generate new samples.

The general workflow includes:-

1. Data Collection
2. Data Preprocessing
3. Model Training
4. Synthetic Data Generation
5. Privacy Evaluation
6. Utility Assessment
7. Deployment

The generated datasets can then be used for research, software testing, AI model training, and benchmarking.

Major Generative AI Techniques:-

Generative Adversarial Networks (GANs):-

GANs consist of two competing neural networks:-

- Generator
- Discriminator

The generator creates synthetic samples while the discriminator attempts to distinguish real from fake data. Through iterative competition, the generator improves its ability to create highly realistic data.

Advantages include:-

- High-quality synthetic images
- Realistic tabular data
- Strong performance for computer vision

Limitations include:-

- Training instability
- Mode collapse
- Large computational requirements

Variational Autoencoders (VAEs):-

VAEs encode data into latent variables and reconstruct samples from compressed representations.

Advantages:-

- Stable training
- Probabilistic framework
- Efficient latent representation

Limitations:-

- Lower image quality than GANs
- Blurry outputs in some cases

Diffusion Models:-

Diffusion models generate data by reversing a gradual noise process.

Benefits include:-

- High-quality image synthesis
- Better diversity
- Stable optimization

Applications:-

- Medical imaging
- Satellite imagery
- Autonomous driving

Large Language Models (LLMs):-**Modern LLMs generate synthetic textual datasets for:-**

- Chatbots
- Healthcare documentation
- Educational content
- Code generation
- Question-answer datasets

They significantly reduce the cost of creating annotated training datasets.

Privacy-Preserving Mechanisms:-

Several techniques improve privacy protection during synthetic data generation.

Differential Privacy:-

Differential Privacy introduces carefully calibrated noise during training.

Benefits:-

- Mathematical privacy guarantee
- Reduced information leakage
- Compliance with privacy regulations

Federated Learning:-

Instead of collecting centralized datasets, Federated Learning trains AI models locally on distributed devices.

Advantages include:-

- Local data remains private
- Lower transmission risks
- Better scalability

Data Masking:-**Sensitive attributes are transformed before model training through:-**

- Tokenization
- Encryption
- Randomization

Secure Multi-Party Computation

Multiple organizations collaboratively train AI models without sharing raw data.

Applications include:-

- Banking
- Healthcare
- Government analytics

Applications:-**Healthcare:-**

Synthetic patient records enable researchers to develop diagnostic AI models without exposing confidential medical information.

Examples include:-

- Disease prediction
- Drug discovery
- Medical imaging

Finance:-**Banks use synthetic financial transactions for:-**

- Fraud detection
- Credit scoring
- Risk analysis without exposing customer identities.

Autonomous Vehicles:-

Self-driving systems require millions of driving scenarios.

Synthetic environments generate:-

- Traffic
- Pedestrians
- Weather
- Road conditions to improve autonomous navigation.

Cybersecurity:-**Synthetic network traffic supports:-**

- Intrusion detection
- Malware classification
- Attack simulation without revealing confidential network logs.

Education:-**Educational institutions create synthetic student datasets for:-**

- Learning analytics
- Performance prediction
- Academic research while maintaining student privacy.

Benefits:-**Privacy-preserving synthetic data provides numerous advantages:-**

- Enhanced privacy protection
- Regulatory compliance
- Increased data accessibility
- Faster AI development
- Reduced legal risks
- Lower data collection costs
- Improved collaboration between organizations
- Better reproducibility of scientific research

Challenges:-

Despite significant progress, several challenges remain.

Data Utility:-

Synthetic datasets may lose important statistical relationships.

Privacy Leakage:-

Poorly trained models may memorize original records.

Computational Cost:-

Training advanced diffusion models and LLMs requires expensive GPU infrastructure.

Bias:-

Synthetic data may inherit societal biases from original datasets.

Evaluation:-

Measuring both utility and privacy simultaneously remains difficult.

Future Research Directions:-**Future work should focus on:-**

- Privacy-aware diffusion models
- Explainable synthetic data generation
- Federated generative AI
- Quantum-enhanced privacy mechanisms
- Automated privacy auditing
- Energy-efficient generative models
- Domain-specific synthetic datasets
- Regulatory compliance frameworks

Conclusion:-

Generative Artificial Intelligence has emerged as a transformative technology for privacy-preserving synthetic data generation. By leveraging advanced deep learning models such as GANs, VAEs, Diffusion Models, and Large Language Models, organizations can generate realistic datasets while minimizing privacy risks. These techniques facilitate secure research collaboration, regulatory compliance, and scalable AI development across diverse sectors, including healthcare, finance, cybersecurity, education, and autonomous systems. Although challenges such as privacy leakage, computational complexity, and bias remain, ongoing research continues to improve model robustness and data quality. As privacy regulations become increasingly stringent, synthetic data generation will play a crucial role in enabling responsible, ethical, and secure artificial intelligence.

References:-

1. Goodfellow, I., et al. (2014). Generative Adversarial Networks. NeurIPS.
2. Kingma, D. P., & Welling, M. (2014). Auto-Encoding Variational Bayes. ICLR.
3. Dwork, C. (2006). Differential Privacy. ICALP.
4. Kairouz, P., et al. (2021). Advances and Open Problems in Federated Learning. Foundations and Trends in Machine Learning.
5. Ho, J., Jain, A., & Abbeel, P. (2020). Denoising Diffusion Probabilistic Models. NeurIPS.
6. Vaswani, A., et al. (2017). Attention Is All You Need. NeurIPS.
7. Devlin, J., et al. (2019). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. NAACL.
8. Brown, T., et al. (2020). Language Models are Few-Shot Learners. NeurIPS.
9. Abadi, M., et al. (2016). Deep Learning with Differential Privacy. ACM CCS.
10. Salimans, T., et al. (2016). Improved Techniques for Training GANs. NeurIPS.