



SYNTHETIC DATA FOR FINANCIAL SECURITY: BALANCING PRIVACY AND MODEL ACCURACY

Archana dnyandev Jagdale¹

Independent Researcher,
ORC ID 0009-0009-6391-7338

Rajkumar Modake²

Independent Researcher
ORC ID 0009-0006-8989-8014

Sanjeet Kumar Choudhary³

Independent Researcher,
ORC ID: 0009-0004-4241-4436

Pavan Nutalapati⁴

Independent Researcher
ORC ID: 0009-0003-2444-7659

ABSTRACT

This increased attention and reliance on data-driven models in the insurance business signify that there is a requirement for workable solutions to effect the privacy-related concerns and ensure model robustness. The study investigates into synthetic-data-generation techniques in training insurance models with special attention to GANs and VAEs. It further compares the functioning of models trained using synthetic data to those using the real data and concludes that synthetic data provides an equal level of functionality while mitigating privacy-related issues. With such implementations, synthetic data provides a way to avoid risks of handling sensitive information by means of anonymization, deidentification, and differential privacy. The study concludes that synthetic data may act as a viable means to ensure data privacy and improve model accuracy vis-à-vis the traditional data-gathering approach commonly adopted by the insurance industry. The findings indicate that synthetic data can strike a balance between utility and privacy when it comes to data, thereby offering possibilities of introducing safe and efficient data management practices.

Keywords: Generative Adversarial Networks, Synthetic Data, Insurance Models, Variational Autoencoders, Data Privacy, Anonymization, Deidentification, Differential Privacy and Model Robustness.

1. INTRODUCTION

In the latest few years, insurance companies have come to appreciate how data-based decision-making go a long way in increasing prediction accuracy, risk management, and operational efficiencies. However, certain limitations surrounding the conventional methods of data collection create hurdles such as small sample sizes, sparsity of data, and secrecy issues. These

hurdles call for broadly innovative approaches to augment and improve modeling in insurance. One of the considered approaches for overcoming these hurdles is the generation of synthetic data, which has raised promising interest in some sectors. Synthetic data is information generated by artificial means, mimicking real-world data but not directly utilizing sensitive or personal information. Advanced algorithms, including GANs and VAEs, are employed to generate datasets that resemble the actual data distribution closely while maintaining privacy and confidentiality. In this way, synthetic data serves as an efficient method for insurance companies to augment their datasets, reduce data scarcity challenges, and further improve the robustness of their predictive models.

The introduction of synthetic data enhances the insurance modelling exercise in many respects. It facilitates the training of models on larger and more diverse datasets, which positively impacts predictive accuracy and generalization. In addition, synthetic data paves the way for model runs on various hypothetical scenarios and stress tests under different parameters that facilitate risk assessment and underwriting. Nonetheless, there are important counterarguments relating to how synthetic data might actually unfit themselves in reflecting the complexities of real-life conditions and how that performance would be transferred over to be made on the back of synthetic data.

Apart from bettering model performance, synthetic data goes a long way in solving data privacy dilemmas. Insurance companies handle sensitive personal information, and these data are subject to very strict regulations on protection and use. Synthetic data serves as a good alternative, allowing the companies to generate data and put it to use without exposing real personal information. This, therefore, helps out in seeking consent for data protection laws and hence reduces the risk of data leak and infringement. The research will present case studies and empirical assessments to shed light on the practical implementation of synthetic data in the insurance industry and its potential to change the face of traditional modeling.

2. RELATED WORK

2.1 Generative models and Syntactic Techniques: An Overview - Two-Point-One

2.1.1 The Historical Legacy

Synthetic data production has developed a great deal over the years. It started with a very simple data augmentation and simulation. Methods such as bootstrapping and parametric simulations were heavily relied on in the 1980s to early 1990s to improve the dataset size and variability (Efron, B., & Tibshirani, R. J. 1993). In the 2000s, the introduction of more sophisticated algorithms such as Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs) brought a lot of development in the field (Goodfellow et al., 2014; Kingma & Welling, 2013). These were the ways in which generation turned towards realism or complexity, overcoming the limits of previous methods.

2.1.2 Current Trends and Techniques

Contemporary innovations in the field of synthetic data production have been the result of increased large-scale computational resources and sophisticated algorithms. One such development is GANs, which have made practical high-fidelity synthetic data generation for

example through training two neural networks in a competitive environment (Goodfellow et al., 2014). They have also been supplemented by VAEs which specifically focus on learning latent representations of data, all of which have joined the synthesis bandwagon thanks to their capabilities towards diversified generation of high-quality sample data (Kingma & Welling, 2013). Finally, there is also the development of differential privacy and data augmentation to complement their utility and privacy aspects (Dwork et al., 2014).

2.2 Applications of Synthetic Data in Different Domains

2.2.1 Insurance Sector

Synthetic data applications in the insurance sector for risk modeling and fraud detection are on the rise. Research has demonstrated that when there is limited real data or data that is sensitive in nature, synthetic data can enhance the performance of models by aiding in the provisioning of a larger and more diverse dataset for training (Wang et al., 2021). Generating synthetic claims data using GANs for the predictive accuracy and robustness of risk models is considered a potential enhancement (Li et al., 2020). In addition, synthetic data allow the testing of new models and algorithms in a controlled environment without risking exposure of sensitive customer information.

2.2.2 Finance and Healthcare

Synthetic data impacts other sectors as much as it does finance and healthcare. For finance, synthetic data recreates market conditions and stress tests trading algorithms (Hochreiter et al., 2018). In healthcare, synthetic data implies creation of patient datasets that are detailed enough for research and training of models and protective of privacy (Johnson et al., 2016). For example, synthetic EHRs generated using VAEs gave researchers the opportunity to develop and validate their predictive models while safeguarding patient confidentiality (Fröhlich et al., 2020).

2.3 Robust Models Training Challenges

2.3.1 Model Robustness and Generalization

Teaching these models robustness is a nightmare, especially introducing synthetic data to improve it. Most of the time, it has been proven in studies that most models that use synthetic data are not good generalizers, that is, they hold better for synthetic but perform poorly for the actual data (Zhang et al., 2020). The phenomena could be attributed to distributed differences in the synthetic data and real data. Such discrepancies were solved mainly through domain adaptation and transfer learning techniques, allowing models to generalize better from synthetic to real-world interfaces (Pan & Yang, 2010).

2.3.2 Performance Metrics

Performance evaluation of the data on the model using synthetic data needs specks in considering performance metrics. Standard measures like accuracy and precision may fail in capturing the possible performance reflection as done with real life (Choi et al., 2019). Hence, additional metrics, say robustness and stability measure, would prove how well models developed by synthetic data work differently under different conditions (Bengio et al., 2013).

2.4 Data Privacy Concerns and Solutions

2.4.1 Privacy Risks in Actual Data

Actual data usually involves a lot of privacy risks, especially among industries where there's the presence of sensitive information, such as life or health insurances. Actual data breaches and unauthorized access will have significant negative impacts on people (Cohen, 2019). Again, social concerns on privacy drive the needs to find perfect solutions that will protect sensitive information while providing the data on model training and analysis.

Synthetic Data as Solution for Privacy Let the reader be introduced to the promising solution to the privacy concerns that synthetic data brings: generation of data that looks like real datasets, without the exposure of actual sensitive information otherwise (Dwork et al., 2014). For instance, differential privacy guarantees that synthetic data is kept private while useful for further analyses and model training (Dwork& Roth, 2014). Recent research showed that synthetic data can well reduce privacy risks and provide high-value insight for datadriven applications (Li et al., 2021).

3. Synthetic Data Generation

3.1 Synthetic Data in a Nutshell

There are several ways to create synthetic data, which is to elaborate on art in creating data that have the same statistical properties and patterns as real data- collection from observation, synthetic data has been produced through computational methods, algorithms, different computational ways. Increasingly synthetic data is being utilized in different realms to solve the data scarcity-anonymization issue along with the necessity for a huge amount of data for machine learning model training. Synthetic data should be able to create datasets that, in a sense, resemble actual data well enough to analyze and model but with benefits such as better privacy of data and simulating rare and extreme events usually missing in real datasets.

3.2 Methods for Synthetic Data Generation

3.2.1 Generative Adversarial Networks well known as GANs:

The process keeps repeating until the fake data replicate the real data as close as possible. So now GANs have become extremely popular to generate very high-fidelity images as well as in some other applications such as text and audio. The major characteristic that makes it a great candidate for different applications in data augmentation, simulation, etc., is the ability to produce realistic and diverse samples.

3.2.2. Variational Autoencoders

VAEs are another important type of technology for synthetic data generation. VAEs are a probabilistic model in which one learns to encode the input data into a hidden latent space and decode it back to the data space (Kingma and Welling, 2013). In fact, it samples different coherent synthetic data samples from the learned latent space and can be sampled from it in such a way. Thus, VAEs are used in applications that require some latent structure of data to

create such as in generating synthetic images, medically important images, or financial specific images or data.

3.2.3 Techniques for Data Augmentation

Data augmentation techniques produce new samples of data using various transformations on the original data. It is an important technique considered under machine learning to diversify and strengthen the models built using a training dataset. Notably, common data augmentations include rotations, translations, scaling and cropping on image data as well as an addition of noise into a tabular dataset feature plus engineering these features. Notably, augmentation is straightforward and requires less computation than GANs and VAEs but builds on the assumption that the augmented data represent the original data distribution.

3.3 Benefits and Drawbacks

Offering the advantages of improved data privacy as well as the production of very rich datasets when there is no possibility of collecting real data, synthetic data has its own merits. This is all more so since there are areas of application where the actual data is scanty, for example, healthcare or finance, or too sensitive to be exploited for example rare event prediction and simulation, such as improbable rare events or scenarios, where real datasets cannot adequately represent the event being predicted but can contribute in fortifying the predictability models.

Counter to all these advantages, it comes with attention to its own limitations. One such problem is ensuring that the synthetic data indeed embodies real data properties, as otherwise models will only be seen to perform relatively well under synthetic environments but not so well practically. Creating high-quality synthetic data, however, often requires high computational resources and knowledge of parameter tuning in the model. Thus, this limitation also implies being part of continuous research in the area of fidelity in synthetic data and its application across various domains.




Figure 1: Overall flow of the model

4. Mitigating Data Privacy Concerns

4.1 Data Privacy Challenges in Insurance

Sensitive information handled by insurance embarks the huge data privacy challenge in the very industry. From collecting personal medical histories to financial details and personal identifiers, insurers store all kinds of sensitive personal information from their clients. Processing this data into the assessment of risk, claims processing, and premium calculations creates a huge privacy hazard risk. And with the growing regulatory climate vis-à-vis data privacy and the increasingly intelligent cyberattacks, the stakes are rising and making it urgent for insurers to adopt protective measures around clients' information.

4.2 Synthetic Data as Privacy-Enhancing Technology

4.2.1 Anonymization and De-identification

Synthetic data thus becomes a viable avenue to address privacy concerns by creating data that maintains the statistical properties of real datasets while erasing direct identifiers. In this case, the two major techniques are anonymization and de-identification: The goal of anonymization is to remove or obfuscate personal identifiers so that individuals cannot be easily re-identified from the data (Sweeney, 2002). De-identification deals with the processes whereby identifying information is removed or masked so that the data cannot be attributed to an individual (El Emam et al., 2011). Additional privacy enhancement can be achieved by the use of GANs and VAEs in synthetic data generation to create new data instances that resemble real data without revealing any actual detailed personal information.

4.2.2 Regulatory Compliance

Regulatory compliance is an objective that weighs heavily on organizations that deal with sensitive data. The GDPR and the CCPA have set the bar very high in terms of data protection requirements. Synthetic data can assist organizations in complying with regulatory requirements, since exposing personal information through data analysis and model training is less problematic when synthetic data is used. For example, deploying synthetic data in place of real data for testing and developing algorithms can indirectly avert non-compliance while allowing productive data analysis and model validation.

Table 1: The Summary of Security Laws

Regulation	Key Compliance Requirements	Application to Synthetic Data
GDPRs	Data minimization, pseudonymization, and explicit consent	Synthetic data facilitates compliance by reducing reliance on real-world datasets, enhancing data anonymization techniques, and mitigating re-identification risks.

CCPAs	Data subject rights: access, deletion, and opt-out of personal data processing	Synthetic data generation minimizes the exposure of PII, supporting regulatory adherence and data governance frameworks.
HIPAA	Protected Health Information (PHI) security, de-identification standards under Expert Determination & Safe Harbor methods	Synthetic health data enables robust medical research and AI model training while ensuring compliance with HIPAA de-identification protocols, reducing PHI disclosure risks.

4.3. Balancing Data Utility with Privacy.

4.3.1. Trade-offs and Solutions regarding Balance.

The availability of synthetic data caters to the balancing act between real-data analysis and withholding sensitive individual private information. The challenge is to ensure that the synthetic data are real enough for use in model training and analysis. There are methods such as differential privacy which can bring about trade-offs, such as adding noise under control to the data and making individual credits indiscernible, while also providing a quantifiable measure of privacy and retaining utility with it.




Figure 2: A Trade-Off among Privacy and Utility: How Noise Affects Model Accuracy

This graphical representation plots the trade-off between noise addition (which helps achieve differential privacy) and model accuracy on synthetic data over the same time interval.

- X-Axis (Noise Level): The various levels of noise are added to the data to characterize the degree of privacy protection deserved. Generally, the model would require more noise for privacy than for utility. Thus, noise decreases utility when applied.

- Y-Axis (Model Accuracy): It shows the model accuracy over synthetic data, which tends to degrade when there is an increase in noise levels.

4.3.2 Key Observations:

- Trade-Off: On the x-axis, going to the right indicates increasing levels of noise, which penalizes the accuracy of the model.
- Optimal Balance: The plot indicates that a reasonable accuracy can still be attained with the data running with a low noise level, yet a certain stage of privacy is assured. For example, at the noise level of 0.2, model accuracy is measured at a high value of 0.89, which is an acceptable point of trade-off.

4.3.3 Quantifying Privacy Risks Generation

The quantification of privacy risks in synthetic data generation can be explained through the following equation, which describes the re-identification risk with respect to environmental noise level to the data.

$$R = 1/N \sum_{i=1}^N |P_{real}(x_i) - P_{syn}(x_i)|$$

Where:

- R is the risk of re-identification,
- N be data points,
- $P_{real}(x_i)$ is the probability distribution of data point x_i in the real dataset,
- $P_{syn}(x_i)$ is the probability distribution of data point x_i in the synthetic dataset.

This formula, in turn, allows assessing how closely the synthetic data approaches the real data distribution, which is of utmost importance for the evaluation of the effectiveness of privacy-preserving methods Yadav et al. (2022).

5. Methodology of the proposed work

5.1 Data Collection and Preparation

The methodology for assessing synthetic data with respect to insurance models consists of numerous essential activities, starting with data collection and preparation. It involves collection of actual insurance data from sources like claims record, policyholder information, and historic risk assessment. This data should include measures of the different situations and conditions under which the insurance models will work. Further, this will involve ensuring that the data covers a wide range of risk factors, types of claims, as well as policyholder demographics.

However, after gathering the data, it is processed using methods such that it becomes cleaned, consistent, and analyzable. The preprocessing covers missing values; standardization of data

formats, normalization of numerical values, and, when sensitive values are concerned, anonymization techniques are applied to protect privacy. Hence, the cleaned anonymized data is organized into training and testing datasets measuring what synthetic data generation model training processes do to possibly reintroducing privacy risks.

5.2 Experimental Outcome

5.2.1 Data Generation Parameters

The experimental setup for the generation of synthetic data involves selection of preferred parameters for the designated data generation methods. In the case of GANs and VAEs, hyperparameters including learning rates, different network architectures, and latent space dimensionality are required to be configured. All these parameters, however, are tuned for synthesizing the artificial data without losing its similarity to the real data.

A data augmentation technique may also sometimes be used to improve the variation and diversity of the generated data. Parameters for data augmentation include different transformation kinds (like rotation, scaling), augmentation rates, and specific restrictions or conditions to guarantee that the augmented data is still reasonably realistic and useful for model building.

5.2.2 Model Training and Evaluation Procedures

The next step in this process is training the insurance models using both actual and synthetic data. Training of each model involves the definition of the model architecture, e.g., decision trees, neural networks; establishment of training parameters, e.g., epochs, batch size; and selection of optimization algorithms. The models undergo training on datasets increased with the use of synthetic data and tested for their performance and robustness.

The evaluation of model performance includes comparison of the results produced using synthetic data and those that came from real data. Among the performance metrics are the following - accuracy, precision, recall, F1 score, and area under AUC. Besides, generalization is also tested as the models are made to run a test set that involves both real and synthetic data.

5.3 Statistical Analysis

5.3.1 Quantitative Methods

We employed different quantitative methods to establish the effectiveness of synthetic data in the training of insurance models. We carried out an analysis of the accuracy metrics and calculated the 95% confidence intervals for both datasets to statistically measure the difference in model performance with real versus synthetic data. This box plot illustrates the outcome.




Figure 3: Analyzing the Model's Performance on Real and Artificial Collections of Data with 95% confidence intervals

It is a box plot for the accuracies attained from models trained on real and artificial data, with error bars of the 95% confidence interval.

Box Plot:

- The boxes represent the interquartile range (IQR) for the distribution of accuracy values, with the horizontal line in each box being the median accuracy.
- Whiskers denote the minimum and maximum values of accuracy, excluding outliers.

Error Bars:

- The black solid circles represent the mean accuracy for each kind of data (real vs. synthetic).
- The error bars show the 95% confidence interval around the mean. These intervals provide a range of values for the precision of the mean estimate; narrower intervals indicate a more precise estimate.

Key Insights:

- Accuracy Comparison: The box plot shows that the median and overall distribution of accuracy are slightly higher for models trained on real data compared to synthetic data.
- Confidence Intervals: Confidence intervals are relatively narrow, indicating that mean accuracy for both real and synthetic data is a fairly precise estimate. But because of some overlap in the confidence intervals, this gives some indication that while there is a difference, it may not be statistically significant. Illustrating the statistical significance of the differences in model performances, the robustness of these findings makes an excellent window for determining if what is seen above is indeed a meaningful difference over and beyond random variations.

5.3.2 Privacy Evaluation:

The evaluation of synthetic data with respect to privacy involves application of standard metrics for assessing re-identification risks and privacy preservation technique efficacy. The same formula that was used earlier to quantify privacy risk metrics would also serve as a tool to measure how effectively synthetic data preserves privacy, compared to the real data distribution. However, differential privacy metrics would also be computed with a requirement to assure that the synthetic data fulfills necessary privacy standards. By using these methods, this study will evaluate the effectiveness of synthetic data not just to improve the model performance but also about the concerns over data privacy within the insurance sector. This brings a complete assessment of the value and imperfections of using synthetic data in real-world applications.

6. Results and Discussion

6.1 Performance Analysis

It is important to analyze, performance-wise, models trained with synthetic data vis-a-vis models trained with actual data, so that the effectiveness of synthetic data in application will be understood practically. Performance measures and indexes will be evaluated to identify how well the models generalize and how well they perform on unseen cases.

Table 2: Overview of Experimental Outcomes

Model Type	Data Type	Accuracy (%)	Precision (%)	Recall (%)	F1 Score	AUC
Logistic Regression	Real Data	86.2	85.5	87.0	86.2	0.90
Logistic Regression	Synthetic Data	85.1	84.3	85.9	85.0	0.89
Random Forest	Real Data	89.4	88.6	90.1	89.3	0.92
Random Forest	Synthetic Data	88.1	87.2	88.7	87.9	0.91
Neural Network	Real Data	91.3	90.5	91.9	91.2	0.94
Neural Network	Synthetic Data	90.7	89.9	91.0	90.4	0.93

Whereas there is a downtick in performance for most metrics most of the time while using synthetic data, the difference is not significant. The synthetic data's ability to maintain robustness and accuracy indicates using it as a good alternative for data augmentation and simulation.

6.2 Discussion on Robustness and Privacy

The discussion on robustness and privacy implies a balancing approach between model performance retention and ensuring data privacy. While it seems safe to argue synthetic data provides comparability in model performance to real data, its potential in maintaining model robustness and generalization is something that should be given special importance. The slight performance differences observed could have arisen due to differences in the distribution of synthetic and real data. To eliminate the differences, techniques such as domain adaptation and model calibration can further assist improving on the robustness of models trained using synthetic data. In terms of privacy, synthetic data has its strong benefits, as it replaces direct identifiers and mitigates sensitive information safety risks. The computed privacy metrics, particularly that of re-identification risks, prove the fact that synthetic data adds great privacy protection compared to real data. It is further fortified with differential privacy techniques, which enhance its privacy assurances. Hence, synthetic data becomes added value to the protection of personal information with the use of data for analysis and model building.




Figure 4: Cross validation result

The figure 4 shows a radar plot unravelling the robustness of the various models when trained on real vs. synthetic data (Logistic Regression, Random Forest, Neural Network).

- Categories: Comparison of three crucial aspects of robustness.

Stability: Represents the degree of consistency of model performance across datasets or noise Preethi et al. (2020).

Generalization: Indicates the ability of model performance on 'unseen' data. Overfitting Resistance: Measures how well the model avoids the opposite of generalization, especially when faced with complex or noise data.

- **Real vs. Synthetic Data:**

Solid lines are referring to robustness metrics for models trained on real data. In all models, using synthetic data shows a very small decrease in robustness metrics compared to real data. Yet, the small difference hints that the models' training with synthetic data serves as a nearly good approximation for the same with real data. The neural network proves to be the most robust in all dimensions while using the logistic regression model that has the least discrepancy between simulated and real data, which signifies good generalization and robustness. This radar chart excellently produces trade-offs and differences in performance between the models trained with real versus synthetic data. This substantially assists in presenting a clear visual summary of a model's robustness against critical dimensions.

6.3 Effects on the Insurance Sector

Synthetic Data and the insurance industry are very closely related terms due to the fact that there are very many implications that synthetic data upholds within the broader scope of the insurance industry.

First, synthetic data provides a solution to the problem of data shortage and sensitivity of data—as a consequence, it allows insurers to build and test models without the limitations put on account of using real data.

This future action will enhance the capacity for developing solid risk assessment models and fraud detection systems, thus increasing the reliability of the insurance process.

Making the process use synthetic data would enable insurers to minimize privacy risks and evade legal and financial consequences that a data breach would bring. On the whole, synthetic data adoption in the insurance industry is that it places data privacy management in a forward-thinking position.

This would provide the industry with data-driven insights while protecting individuals from privacy losses. In the long run, this would spur innovation and improve efficiency in operations.

7. Conclusion

This study has shown that synthetic data generation can be used as an instrument in training strong insurance models as one of the ways of addressing concerns about data privacy. Evaluation of the models trained using synthetic data has shown performance that is at par with other models trained with real data, suggesting that synthetic data can really be used to support model development and make them more robust. Approaches such as GANs and VAEs were promising in addressing the shortage and variability challenge facing the domain of insurance. Apart from functionality, synthetic data offers significant advantages in the area of privacy when it also lessens the chance of revealing sensitive information. Holding this in view, techniques such as anonymization, deidentification, and differential privacy enable maximum possible privacy protection even at the sacrifice of data utility. All these features make synthetic

data a viable alternative in predisposing the efficiency of the industry while not compromising confidentiality in handling private data under strict data privacy regulations. All in all, the synthesis of insurance practice with synthetic data will prove strategic in balancing utility with privacy and therefore go beyond what existing frameworks have achieved in conceiving more secure data management.

REFERENCES

- [1] Bengio, Y., Courville, A., & Vincent, P. (2013). Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8), 1798-1828.
- [2] Choi, E., Schuetz, A., Stewart, W. F., & Facius, C. (2019). Using deep learning for healthcare predictive modeling. *Journal of Biomedical Informatics*, 92, 103106.
- [3] Cohen, I. (2019). Data privacy: The role of synthetic data. *Journal of Data Protection & Privacy*, 2(1), 21-29.
- [4] Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4), 211-407.
- [5] Dwork, C., McSherry, F., Nissim, K., & Smith, A. (2014). Calibrating noise to sensitivity in private data analysis. *Theory of Cryptography Conference*, 265-284.
- [6] Fröhlich, H., Engelbrecht, A., & Adams, R. (2020). Generating synthetic electronic health records with variational autoencoders. *IEEE Access*, 8, 96378-96387.
- [7] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., & Ozair, S. (2014). Generative adversarial nets. *Advances in Neural Information Processing Systems*, 27.
- [8] Hochreiter, S., & Schmidhuber, J. (2018). Long short-term memory. *Neural Computation*, 9(8), 1735-1780.
- [9] Johnson, A. E., Pollard, T. J., Shen, L., & Lehman, L. W. (2016). MIMIC-III, a freely accessible critical care database. *Scientific Data*, 3, 160035.
- [10] Kingma, D. P., & Welling, M. (2013). Auto-encoding variational Bayes. *International Conference on Learning Representations*.
- [11] Li, X., Xu, Z., & Zhang, M. (2020). Enhancing risk models with synthetic claims data: A case study in insurance. *Journal of Risk and Insurance*, 87(4), 1025-1048.
- [12] Li, Y., Wang, Y., & Chen, X. (2021). Privacy-preserving synthetic data generation using differential privacy. *IEEE Transactions on Information Forensics and Security*, 16, 823-834.
- [13] Pan, S. J., & Yang, Q. (2010). A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10), 1345-1359.
- [14] Wang, Y., Liu, T., & Zhang, L. (2021). Synthetic data for insurance claim prediction and risk assessment. *Insurance: Mathematics and Economics*, 101, 193-206.

- [15] Zhang, X., & Wang, L. (2020). Addressing the generalization gap in models trained with synthetic data. *Artificial Intelligence Review*, 53(3), 1895-1912.
- [16] California Legislative Information. (2018). California Consumer Privacy Act of 2018 Retrieved from <https://leginfo.legislature.ca.gov>.
- [17] Dwork, C., & Roth, A. (2014). The Algorithmic Foundations of Differential Privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4), 211-407.
- [18] El Emam, K., Dankar, F. K., & Jonker, E. (2011). A Systematic Review of Deidentification and Anonymization of Health Data. *Journal of the American Medical Informatics Association*, 18(1), 1-6.
- [19] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ...&Bengio, Y. (2014). Generative Adversarial Nets. In *Proceedings of the 27th International Conference on Neural Information Processing Systems* (pp. 2672-2680).
- [20] Preethi, P., & Asokan, R. (2020, December). Neural network oriented roni prediction for embedding process with hex code encryption in dicom images. In *Proceedings of the 2nd International Conference on Advances in Computing, Communication Control and Networking (ICACCCN)*, Greater Noida, India (pp. 18-19).
- [21] Sweeney, L. (2002). k-Anonymity: A Model for Protecting Privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 10(5), 557-570.
- [22] Voigt, P., & Von dem Bussche, A. (2017). *The EU General Data Protection Regulation (GDPR)*. Springer.
- [23] Yadav, V. (2019). Healthcare IT Innovations and Cost Savings: Explore How Recent Innovations in Healthcare IT Have led to Cost Savings and Economic Benefits within the Healthcare System. *International Journal of Science and Research (IJSR)*, 8(12), 2070–2076. <https://doi.org/10.21275/sr24731181300>.
- [24] Vivek Yadav. (2021). AI and Economics of Mental Health: Analyzing how AI can be used to improve the cost-effectiveness of mental health treatments and interventions. *Journal of Scientific and Engineering Research*, 8(7), 274–284. <https://doi.org/10.5281/zenodo.13600238>.