

VIRTUAL WEARABLE SENSOR DATA GENERATION WITH GENERATIVE ADVERSARIAL NETWORKS

Yining Huang

Operations Research Graduate Program
North Carolina State University
915 Partners Way
Raleigh, NC 27606, USA

Hong Wan

Department of Industrial and Systems Engineering
North Carolina State University
915 Partners Way
Raleigh, NC 27606, USA

Xi Chen

Grado Department of Industrial and Systems Engineering
Virginia Tech
1145 Perry Street
Blacksburg, VA 24061, USA

ABSTRACT

This study delves into the utilization of Generative Adversarial Networks (GANs) for generating subject-specific time series sensor data, offering an innovative alternative to traditional metamodel-based simulations. We undertake an in-depth analysis of DoppelGANger, a prominent GAN variant for time series data and metadata generation, evaluating its efficiency and efficacy. The sensor data for this investigation was sourced from the National Health and Nutrition Examination Survey, which served as the foundational training set. We scrutinized the synthesized sensor data corresponding to various physical attributes, focusing on the temporal and multi-dimensional statistical properties. Our empirical findings underscore the potential of GANs to adeptly capture the time-dependent correlations and the intricate statistical characteristics inherent in multi-dimensional data. This insight into GANs' capabilities is a crucial step towards more sophisticated synthetic data generation, with significant implications for future applications in wearable technology and personalized health monitoring systems.

1 INTRODUCTION

Generative models (Alex Lamb 2021), a class of statistical models, are widely used to generate new data that resembles the original dataset on which the model was trained. Typically, generative models are trained using large sets of existing data. Once trained, these models can generate new data by sampling from the learned distribution (Shalom et al. 2019; Bond-Taylor et al. 2022). Popular techniques like transformer technology (Vaswani et al. 2017) and Variation Autoencoder (VAE, Kingma and Welling 2022) have gained considerable attention in the field of Artificial Intelligence Generated Content (AIGC) due to their exceptional performance in producing images, videos, and text information. While generating tabular datasets has received relatively less attention, it holds immense potential for scientists who rely on a substantial amount of high-quality data for accurate research and exploration, especially when data are hard and/or expensive to get. The "Generating Virtual Data" approaches have naturally caught the attention of the simulation society. Cen et al. (2020) combined VAE with Long Short-Term Memory

(LSTM) components to capture intricate stochastic processes and efficiently generate sample paths. Cen and Haas (2022) enhanced simulation metamodeling through using Graph Neural Networks (GrNN) to generate more data. Chalé and Bastian (2021) compared the performance of Markov Chain Monte Carlo with Generative Adversarial Network (GAN) and VAE in estimating the joint probability distribution of network intrusion detection system data and demonstrated the advantage of generative methods.

In this paper, we focus on a specific category of generative models: GANs (Goodfellow et al. 2014), which consists of two Artificial Neural Networks (ANN): the Generator function $G(x)$ and the Discriminator function $D(x)$. The GANs excel at learning and replicating the underlying distribution of real data through an adversarial training process. In other words, the generator will continue to improve until the discriminator cannot distinguish the real and synthesized data. The synthetic data, therefore, will share the statistical properties and patterns present in the real data. By generating reliable synthetic samples, GANs can aid in increasing the diversity and size of available data and imputing missing data. Synthetic data also allows for the sharing of sensitive information without compromising privacy. From the simulation perspective, GANs' ability to generate synthetic data closely resembling real data makes them attractive for various simulation tasks. Barra Montevechi et al. (2021) used GANs for input modeling of discrete-event simulation (DES), demonstrating their performance over conventional approaches while requiring fewer assumptions and constraints. Montevechi et al. (2022) effectively leveraged the generator component for generating synthetic data and the discriminator component for differentiating between real and simulated outputs, serving as a validation mechanism. The performance is impressive.

While the benefits of using GANs-generated data in analysis are massive, it is important to acknowledge and address key challenges that arise. Due to limited training datasets, synthetic data generated by GANs may lack the complexity and diversity present in real-world data and fail to adequately learn the statistical properties of the entire population, potentially rendering biases and inaccuracies in subsequent analysis and difficulties in generalizing the model to unseen or different contexts. Lastly, evaluating certain types of data, particularly time series data, can be challenging when using GANs. There are few established metrics for assessing the quality of generated time series data, making it difficult to evaluate the effectiveness of the method accurately. There are works trying to address these challenges (Yadav et al. 2023; Bahrpeyma et al. 2021). In this paper, we implemented one variation of GANs model to generate synthetic wearable sensor data and virtual human subjects. We deliberately chose this example due to the inherent difficulties and high costs associated with acquiring such data. Our research provides valuable insights into the potential of GANs for generating realistic wearable sensor data and can be extended to other fields facing challenges related to data scarcity.

The rest of the paper is organized as follows. Section 2 provides a comprehensive literature review, examining the prior research in the field of GANs. In Section 3 describes the problem of interest and how the selected GANs work. The experimental setup is presented in Section 4, followed by the results and analysis in Section 5. Section 6 concludes the paper with discussions of insights, limitations, and potential directions for future research.

2 LITERATURE REVIEW

Machine learning techniques have been widely used in the healthcare sector for several years and made vast impacts (Sebastiani et al. 2022; Apell and Eriksson 2023). However, one of the primary challenges within the healthcare sector is the need for more data, due to the hardness to collect data and the strict privacy requirement. Innovative techniques such as GANs have been proposed to overcome this obstacle. Initially, GANs have mainly been applied to image synthesis (Cheng et al. 2020; Beji et al. 2023) and analysis in medical and other fields, for example, anomaly detection (Naidoo and Marivate 2020; Nho et al. 2021). However, in the past few years, its application has extended to non-imaging data augmentation. Researchers have made significant progress in generating data with complex correlation structures (Vaccari et al. 2021; Rashidian et al. 2020), which has the potential to improve the data richness and fair representation of various areas where data collection is challenging.

One important data type in healthcare is time series data, which are, for example, common in continuous monitoring of vital signs such as heart rate and blood pressure. The Recurrent Conditional GAN (RCGAN) (Esteban et al. 2017) is the first GAN-based models that can generate realistic labeled multi-dimensional medical time series data. It is capable of generating time series data conditional on variables such as oxygen saturation and heart rate monitoring, but it fails to adequately capture temporal dependencies within the data. Yoon et al. (2019) introduced TimeGAN. By incorporating both unsupervised and supervised training, this framework enables the generation of realistic samples while preserving the underlying dynamics of the training data during the sampling process. Additionally, TimeGAN introduces a stepwise supervised loss that utilizes the original data as supervision, explicitly promoting the model to capture the stepwise conditional distributions inherent in the data. In other studies, Wasserstein Generative Adversarial Network (WGAN) has been utilized as an alternative to traditional GANs to improve the stability and converging rate of the procedure; see, e.g., Smith and Smith (2019) and Bratu and Czibula (2021).

Besides the temporal dynamics, researchers must also generate time series data in context or time series with metadata (multi-labels). Metadata (label) here refers to descriptive information that provides context and additional details of the time series; it serves as a form of “data about data,” which can be thought of as a set of attributes or properties that describe the characteristics, origin, structure, or usage of the data. This specific data structure poses challenges for using traditional methods, such as linear regression (Faraway 2014) or autoregressive moving average models (Shumway et al. 2000). These challenges arise due to the disparity in dimensions or distributions between the contextual variables and the corresponding time series data. Traditional methods have difficulties effectively handling these two components simultaneously. However, GANs offer a promising solution to reconcile these differences and facilitate a more integrated analysis of the contextual variables and time series data. Lu et al. (2022) presented Multi-Label GAN (MTGAN). This method enhances the quality of uncommon disease generation by employing a gated recurrent unit (GRU) and a smooth conditional matrix for sequence and disease generation. The generator initially generates patient-level diagnosis probabilities recursively using the GRU and then incorporates the concept of a conditional vector to create a smooth conditional matrix across all visits in the sequences. Li et al. (2022) proposed transformer-based time-series conditional GAN (TTS-CGAN) to train on existing multi-class datasets and generated class-specific synthetic time-series sequences of variable lengths. Furthermore, Ehrhart et al. (2022) proposed an LSTM Fully Convolutional Networks conditional GAN (LSTM-FCN cGAN) architecture to generate conditional time series data with a single label. The aforementioned studies have primarily focused on incorporating a unique situation or a single-label context for time series data, but they have not yet fully explored the integration of time series data with metadata. The DoppelGANger (DGAN) introduced by Lin et al. (2020) effectively incorporates both time series data and subject-specific latent space information, making it applicable in various scenarios. Lin et al. (2020) demonstrated that DGAN holds immense promise as an expressive time series model capable of generating synthetic data with exceptional fidelity. A detailed explanation of their work is provided in the next section.

3 METHODOLOGY AND DATA DESCRIPTION

3.1 Background and Motivation

According to NIH’s definition (NIH 2023), circadian rhythms (molecular/internal clock) refer to “physical, mental, and behavioral changes that follow a 24-hour cycle.” One of the most critical and well-known circadian rhythms is the sleep-wake cycle. There is overwhelming evidence of the vital link between circadian rhythms and various chronic diseases, and a person’s circadian rhythm, therefore, can serve as a strong indicator of his/her health. Wearable devices equipped with accelerometers and light and temperature sensors can provide continuous and objective measures of circadian rhythms, making it the most popular approach to measuring circadian rhymes nowadays. On the other hand, while the technology is maturing and the data collection becomes easier, the process can still be compromised by participant compliance,

privacy and ethics, cost and logistics, etc. Hence, virtual wearable sensor data hold significant potential in advancing circadian rhythm studies, which our paper focuses on (Hannay and Moreno 2020; Sarwar et al. 2023).

3.2 DoppelGANger Architecture

In this paper, we employ DGANs proposed by Lin et al. (2020) to tackle the challenge of attribute-conditioned time series data generation. We apply this method to train and generate individual patients' sensor data and a virtual patient cohort encompassing patient attributes and sensor data. The overall architecture is shown in Figure 1. We define Y as a collection of samples $Y_1, Y_2, \dots, Y_n$ (for example, the data record for n clients), and each sample Y_i comprise m metadata attributes $A_i = [A_1^i, A_2^i, \dots, A_m^i]$ (the i th participant's attributes like age, gender, and weights). The second component of each sample is a time series of records $R^i = [R_1^i, R_2^i, \dots, R_T^i]$, where R_j^i represents the j th time-unit record of the i th participant. One unique property of the DGANs is the decoupling of attributes and the time series data-generating process, each with its own generators and discriminators. There are two Multi-Layer Perceptron (MLP) generators to generate metadata for the synthetic time series data. The "metadata generator" generates real input time series data metadata. This MLP takes as input a vector of metadata about the input time series (such as its mean and standard deviation) and generates a new vector of metadata that represents the underlying structure of the input data. This generator is conditioned during training on the input time series data and its corresponding real metadata. The "Min/Max generator" generates "fake" metadata for each time series's minimum and maximum values. This MLP takes as input the real metadata generated by the metadata generator MLP and generates two new vectors representing the minimum and maximum values of each time series. These "fake" vectors are used during training to help alleviate mode collapse in the synthetic time series. The generator used for time series data is a bidirectional Long-Short-Term Memory (LSTM) network, which maps a window of real time series data into a lower-dimensional latent space representation that captures the underlying structure of the data. The bidirectional LSTM is a type of recurrent neural network (RNN) that addresses the issue of learning long-term dependencies. During training, the generator considers the conditional relationship between time series data and metadata: $P(A^i, R^i) = P(A^i) \times P(R^i|A^i)$ using the RNN with s samples generated at each pass. By incorporating information about the input data's structure and statistical properties, DGANs can generate synthetic time series data that preserves important characteristics of the original data.

There are two discriminators used in DGANs. The auxiliary discriminator only discriminates metadata, while the original discriminator discriminates between real and synthetic time series data. The auxiliary




Figure 1: The architecture of DGAN, adapted from Figure 7 in Lin et al. (2020).

discriminator inputs real and synthetic metadata (such as mean, standard deviation, etc.) associated with each window of time series data and outputs a scalar value representing the probability that the metadata is real. This discriminator uses a Wasserstein-distance-based loss function to ensure the generator network produces synthetic metadata statistically similar to the real metadata. This original discriminator is a standard Wasserstein GAN (WGAN) discriminator that takes a window of real or synthetic time series data as input and outputs a scalar value representing the probability that the input data is real. Both discriminators are trained simultaneously with a weight parameter α determining how much weight to give to each loss function. By combining these two discriminators, DGANs can generate synthetic time series data that preserve important statistical properties of their underlying structure through the generated metadata and their temporal dependencies through the generated time series data.

3.3 Wearable Sensor Data for Training Purpose

The National Health and Nutrition Examination Survey (NHANES) is a comprehensive research program to evaluate individuals' health and nutritional status in the United States. As part of the 2011-2012 NHANES study cycle, wrist accelerometers were incorporated, marking the first time 24-hour object actigraphy data were available on a nationally representative sample. This advancement provides valuable insights into the physical activity patterns of Americans, which is critical for informing public health policies and interventions.

In the 2011-2012 cycle of NHANES, individuals aged six years and older wear an accelerometer (ActiGraph Model GT3X+, ActiGraph, Pensacola, Florida) on their non-dominant wrist for seven consecutive days, both during the day and at night. Raw signals were collected every 1/80 of a second (80 Hz) on the x , y , and z axes and subsequently processed, flagged, and summarized at the minute level. NHANES released the summary measures in November 2020 as a minute summary file known as PAXMIN. PAXMIN stands for "Physical Activity Monitor Minutes," which employs monitor-independent movement summary (MIMS) units. These MIMS units were developed by researchers at Northeastern University as a non-proprietary, open-source, and device-independent universal summary metric (John et al. 2019).

Our study uses the MIMS triaxial value (variable name: PAXMTSM) at the minute level as our input. According to NHANES, there was a total of 9,756 persons who participated and completed the NHANES interview, and 9,338 underwent an examination in the 2011-2012 cycle. However, MIMS triaxial values were changed to missing (i.e., a value of 0) if certain conditions were met. To analyze the circadian rhythm pattern from this dataset, we selected only participants who had completed sensor data, meaning they had 11,530 recordings in the dataset, and there are 2,854 individuals in our dataset.

3.4 Data Preprocess

Each participant in our dataset has 11,530 recordings spanning nine days. Data are measured in minutes and each day has 1,440 recordings except the first and last days which have fewer recordings. We removed recordings on those two days for each participant, resulting in 10,080 recordings per participant. Additionally, we obtained demographic information for each participant from NHANES, which we associated with their Physical Activity Monitor (PAX) sensor data. In this study, we only retain gender, age, and race information from the demographic dataset. The resulting customized dataset is comprised of two parts: the demographic information and the time series sensor data for each participant. In other words, each record contains the subject's gender, age, and race information, and seven days of PAX sensor data.

4 EXPERIMENTAL SETTINGS

4.1 Settings of Experiments

In this section, we explore the power of DGANs and delve into their capability to generate synthetic time series data, both with and without attributes. In this study, a cohort of 50 participants were randomly

selected. Each record contains one participant's attributes and the corresponding sensor data. We conducted a two-phase experiment. The first phase involves training individual DGAN models for each participant utilizing their respective data, thereby enabling the generation of a 7-day information forecast tailored to each individual. Consequently, we had a dataset containing time-series values at the minute level, 1,440 records per day for seven days, for each set of attributes.

During the second phase, we trained a collective DGAN model by integrating the information from all participants. The model will generate synthetic data containing both attributes and time series features. Since the generated attribute sets can be a random combination of all possible values, the synthetic attribute sets may contain new attribute sets that do not exist in our training data. For example, given two attribute sets of a male aged 35 and a 23-year-old female, we can generate 35-year-old males and 23-year-old females with in-training set attributes and 23-year-old males and 35-year-old females with out-of-training set attributes.

We employ a benchmark known as conditional tabular GAN (CTGAN, Xu et al. 2019) to assess the attributes generated in Phase II. CTGAN comprises three essential components: the conditional vector, the generator loss, and the training-by-sampling method. By utilizing these elements, CTGAN guarantees that its generator produces synthetic data that closely aligns with real data distribution, preserving its structural properties and ensuring accurate training even in the presence of data imbalance. We implemented the experiments in Python, and used the gretel-synthetics package (Gretel Labs, Inc. 2023). The code and data used in this study are available at <https://github.com/dreamprompt/dgantest>.

4.2 Evaluation Metrics

Since the output consists of two distinct components—the attributes generated in the second phase and the time series data for each individual generated in both phases—evaluating these diverse datasets using the same metrics presents a challenge. Therefore, to ensure an accurate assessment, it is essential to separate and evaluate these datasets individually. This approach will enable comprehensive analysis and evaluation of each dataset's performance using metrics specifically tailored to their unique characteristics.

To assess the quality of the generated attributes, we will employ the Kolmogorov-Smirnov (KS) distance (Massey 1951) as a quantitative measure. This distance metric will give us a numerical tool to gauge the similarity between the generated attributes and the original dataset. Additionally, we will utilize histograms as a visualization tool to understand better the distribution patterns and potential discrepancies between the generated attributes and the ground truth data. By combining the KS distance and histograms, we can effectively evaluate the accuracy and fidelity of the generated attributes.

The KS distance measures the distance between the empirical distributions implied by two different datasets. The empirical distribution function F_n for a sample of n independent and identically distributed ordered observations X_i is defined as $F_n(x) = n^{-1} \# \{\text{elements in the sample} \leq x\} = n^{-1} \sum_{i=1}^n \mathbf{1}_{(-\infty, x]}(X_i)$, where $\mathbf{1}_{(-\infty, x]}(X_i)$ takes value 1 if $X_i \leq x$ and 0 otherwise. The KS statistic for the two given empirical distribution functions $F_n^1(x)$ and $F_n^2(x)$ is given by $D_n = \sup_{x \in \mathbb{R}} |F_n^1(x) - F_n^2(x)|$.

To gauge the authenticity and precision of the generated time series data, we will utilize various evaluation measures. Initially, we will compare the basic statistical characteristics of the generated data, including the mean, minimum, maximum, and standard deviation, with the corresponding statistics of the training set. Furthermore, we will utilize the Wasserstein distance (Vallender 1974) metric to quantify the dissimilarity between the generated and real data distributions. The Wasserstein distance, denoted as $W(P, Q)$, measures the dissimilarity between the empirical distribution P built on the observations $\{X_i\}_{i=1}^n$ and the empirical distribution Q based on the observations $\{Y_i\}_{i=1}^n$. It can be defined as

$$W(P, Q) = \inf_{\Pi} \left(\sum_{i=1}^n |X_i - Y_{\pi(i)}| \right),$$

where Π is a set of all permutations of n elements, $\pi(i)$ is the corresponding index of Y in a permutation, and X_i and Y_i represent the i th observations from distributions P and Q , respectively.

Lastly, we employ Principal Component Analysis (PCA) plots to visually depict the relationship between generated samples' distribution and the original data's distribution in a two-dimensional space. By doing so, we can effectively discern patterns, clusters, and relationships within the data.

5 EXPERIMENT RESULTS

5.1 Synthetic Attributes

As mentioned in Section 3.4, the model incorporates information regarding the gender, age, and race of each participant. In the second phase, after training the collective DGAN model, we can generate data of any desired sample size. To facilitate comparison, we generated 50 samples for 100 independent macro-replications and recorded the results.

To begin with, we compare the attributes generated in Phase II with those generated using CTGAN, as depicted in Figure 2. We see that the range of KS distances for all three attributes corresponding to DGAN (Figure 2(a)) spans from 0 to 0.26, while that corresponding to CTGAN (Figure 2(c)) is from 0.06 to 0.34. This disparity suggests that the distributions of the synthetic data generated by DGAN are closer to those of the original data than those generated by CTGAN. Additionally, the p -values calculated indicate whether there is sufficient evidence to reject the null hypothesis, which assumes that the two distributions being compared are identical. At a 0.95 confidence level, we will reject the null hypothesis for the attribute "Gender" in certain cases of CTGAN (Figure 2(d)), whereas for most instances of DGAN (Figure 2(b)), we will not reject the null.

Finally, let us examine the histograms of the synthetic data generated in Phase I and compare it with those of the original data shown in Figure 3. Although the outcome is random, it is evident that the synthetic data generated by DGAN effectively captures the distribution information for gender and race, which have fewer categories. However, when it comes to generating age, which entails distinct values, the results are not as ideal as with gender and race but still reasonable.

5.2 Synthetic Time Series

As outlined in Section 4.1, DGAN will produce time series data whose attributes may or may not be in our training dataset. To facilitate comparison, we focus only on solely on presenting the results for the generated data with attribute sets belonging to the training dataset.

In Figure 4, each subplot illustrates the results for a specific participant. Within each subplot, the different datasets are depicted using boxplots. The first boxplot corresponds to the original data, the second boxplot represents the Phase I synthetic data, and the third one represents the Phase II synthetic data. Upon comparing the Phase I and Phase II boxplots with that of the original data, it becomes evident that they closely resemble the original data in most subplots. Conversely, the Phase II boxplot exhibits more significant variability when compared to the those of the preceding two datasets. During Phase II, the collective model is trained to learn the characteristics that are influenced by other participants, rather than focusing solely on the individual patterns exhibited by each participant.

After examining the characteristics of the generated time series data, we compute the Wasserstein distance for the data generated in Phase I and Phase II compared to the original dataset. The results are presented in Figure 5. Overall, the data generated in Phase I exhibits a smaller Wasserstein distance when compared to the original dataset.

Finally, the PCA plots presented in Figure 6 offer valuable insights into the characteristics of the synthetic data generated in Phase I and II compared to the original data. In Figure 6(a), where the synthetic data are produced using an individual model trained on its corresponding dataset, we observe a more dispersed distribution than that of the original data. This broader dispersion captures a more comprehensive range of variability in the data, effectively representing the original data. Conversely, in Figure 6(b), the




Figure 2: Scatter plots of the KS statistics D_n and the corresponding p -values for DGAN and CTGAN.




Figure 3: Histograms of the DGAN synthetic data and the original data of gender, age, and race attributes.

synthetic data are generated collectively, utilizing all the time series data from the 50 participants. We observe a tighter clustering of the synthetic samples, indicating a higher level of similarity among the generated samples. This tighter distribution reflects the model's tendency to reproduce patterns in the training data. The quality of the results in Figure 6(b) is less satisfactory than that of those in Figure 6(a).




Figure 4: Comparison of boxplots: training data (O) vs. Phase I synthetic data generated by DGAN (I) vs. Phase II synthetic data generated by DGAN (II).




Figure 5: Comparison of the Wasserstein distances calculated in Phase I and Phase II.




Figure 6: Comparison of the PCA results.

6 DISCUSSION AND FUTURE RESEARCH

In this study, we conducted a small-scale experiment with a dataset of only three attributes and 50 samples. The purpose was to assess the suitability of DGANs for generating time series data with accompanying metadata after appropriate training. Our findings indicate that DGANs demonstrate capability in generating time series data with metadata following adequate training. However, while DGANs exhibit promising results when evaluating individual attributes, their performance when assessing jointly all three attributes remains to be determined. In contrast, CTGAN performs well in generating attributes jointly and we will explore it further.

We selected three discrete attributes for comparison purposes in this study. The model, however, can also include continuous data. The time series data generated by individually trained DGANs are of higher quality than those generated by collectively trained DGANs. The observation is reasonable since an individualized model eliminates the variability introduced by time series data belonging to other participants, as opposed to collective model training.

One notable limitation of DGANs, in contrast to CTGANs, is their inability to generate data with specific attributes. Addressing and improving this limitation will be a focal point for future direction of our model. During our experiment, we encountered a notable issue highlighted by Choi et al. (2020). Our experiment involved 50 participants, but when we utilized DGANs to generate synthetic data, we observed that two participants out of 50 could not be replicated (Figure 4). In other words, DGANs lacked population control, thus failing to address the bias problem. This issue was also discussed by Xu et al. (2019), who highlighted the challenge of generating underrepresented data. To tackle this challenge, we are actively pursuing the integration of DGANs and CTGANs, aiming to leverage the strengths of both methods.

REFERENCES

- Apell, P., and H. Eriksson. 2023. "Artificial Intelligence (AI) Healthcare Technology Innovations: The Current State and Challenges from a Life Science Industry Perspective". *Technology Analysis & Strategic Management* 35(2):179–193.
- Bahrpeyma, F., M. Roantree, P. Cappellari, M. Scriney, and A. McCarren. 2021. "A Methodology for Validating Diversity in Synthetic Time Series Generation". *MethodsX* 8:101459.
- Barra Montevechi, J. A., A. T. Campos, G. T. Gabriel, and C. H. dos Santos. 2021. "Input Data Modeling: An Approach Using Generative Adversarial Networks". In *Proceedings of the 2021 Winter Simulation Conference*, edited by S. Kim, B. Feng, K. Smith, S. Masoud, Z. Zheng, C. Szabo, and M. Lope, 1–12. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

- Beji, A., A. G. Blaiech, M. Said, A. B. Abdallah, and M. H. Bedoui. 2023. "An Innovative Medical Image Synthesis Approach Using Dual GAN Deep Neural Networks for Improved Segmentation Quality". *Applied Intelligence* 53(3):3381–3397.
- Bond-Taylor, S., A. Leach, Y. Long, and C. G. Willcocks. 2022. "Deep Generative Modelling: A Comparative Review of VAEs, GANs, Normalizing Flows, Energy-based and Autoregressive Models". *IEEE Transactions on Pattern Analysis & Machine Intelligence* 44(11):7327–7347.
- Bratu, A., and G. Czibula. 2021. "DAuGAN: An Approach for Augmenting Time Series Imbalanced Datasets via Latent Space Sampling Using Adversarial Techniques". *Scientific Programming* 2021:7877590.
- Cen, W., and P. J. Haas. 2022. "Enhanced Simulation Metamodeling via Graph and Generative Neural Networks". In *Proceedings of the 2022 Winter Simulation Conference*, edited by B. Feng, G. Pedrielli, Y. Peng, S. Shashaani, E. Song, C. G. Corlu, L. H. Lee, E. P. Chew, T. Roeder, and P. Lendermann, 2748–2759. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Cen, W., E. A. Herbert, and P. J. Haas. 2020. "NIM: Modeling and Generation of Simulation Inputs Via Generative Neural Networks". In *Proceedings of the 2020 Winter Simulation Conference*, edited by K.-H. G. Bae, B. Feng, S. Kim, S. Lazarova-Molnar, Z. Zheng, T. Roeder, and R. Thiesing, 584–595. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Chalé, M., and N. D. Bastian. 2021. "Challenges and Opportunities for Generative Methods in the Cyber Domain". In *Proceedings of the 2021 Winter Simulation Conference*, edited by S. Kim, B. Feng, K. Smith, S. Masoud, Z. Zheng, C. Szabo, and M. Lope, 1–12. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Cheng, K., R. Tahir, L. K. Eric, and M. Li. 2020. "An Analysis of Generative Adversarial Networks and Variants for Image Synthesis on the MNIST Dataset". *Multimedia Tools and Applications* 79:13725–13752.
- Choi, K., A. Grover, T. Singh, R. Shu, and S. Ermon. 2020. "Fair Generative Modeling via Weak Supervision". In *Proceedings of the 37th International Conference on Machine Learning*, edited by H. D. III and A. Singh, 1887–1898: JMLR.org.
- Ehrhart, M., B. Resch, C. Havas, and D. Niederseer. 2022. "A Conditional GAN for Generating Time Series Data for Stress Detection in Wearable Physiological Sensor Data". *Sensors* 22(16):5969.
- Esteban, C., S. L. Hyland, and G. Rätsch. 2017. "Real-valued (Medical) Time Series Generation with Recurrent Conditional GANs". <https://arxiv.org/abs/1706.02633>, accessed 01th May 2023.
- Faraway, J. J. 2014. *Linear Models with R*. CRC press.
- Goodfellow, I. J., J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. 2014. "Generative Adversarial Nets". In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2*, edited by Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, 2672–2680. Cambridge, MA, USA: MIT Press.
- Gretel Labs, Inc. 2023. "Gretel Synthetics". <https://github.com/gretelai/gretel-synthetics>, accessed 01th May 2023.
- Hannay, K. M., and J. P. Moreno. 2020. "Integrating Wearable Data into Circadian Models". *Current Opinion in Systems Biology* 22:32–38.
- John, D., Q. Tang, F. Albinali, and S. Intille. 2019. "An Open-Source Monitor-Independent Movement Summary for Accelerometer Data Processing". *J Meas Phys Behav* 2(4):268–281.
- Kingma, D. P., and M. Welling. 2022. "Auto-Encoding Variational Bayes". <https://arxiv.org/abs/1312.6114>, accessed 01th May 2023.
- Alex Lamb 2021. "A Brief Introduction to Generative Models". <https://arxiv.org/abs/2103.00265>, accessed 01th May 2023.
- Li, X., A. H. H. Ngu, and V. Metsis. 2022. "TTS-CGAN: A Transformer Time-series Conditional GAN for Biosignal Data Augmentation". <https://arxiv.org/abs/2206.13676>, accessed 01th May 2023.
- Lin, Z., A. Jain, C. Wang, G. Fanti, and V. Sekar. 2020. "Using GANs for Sharing Networked Time Series Data: Challenges, Initial Promise, and Open Questions". In *Proceedings of the ACM Internet Measurement Conference*, edited by D. Levin, A. E. Mislove, J. Amann, and M. J. Luckie, 464–483. New York, NY, USA: Association for Computing Machinery.
- Lu, C., C. K. Reddy, P. Wang, D. Nie, and Y. Ning. 2022. "Multi-Label Clinical Time-series Generation via Conditional GAN". <https://arxiv.org/abs/2204.04797>, accessed 01th May 2023.
- Massey, F. J. 1951. "The Kolmogorov-Smirnov Test for Goodness of Fit". *Journal of the American Statistical Association* 46(253):68–78.
- Montevecchi, J. A. B., G. T. Gabriel, A. T. Campos, C. H. dos Santos, F. Leal, and M. E. F. H. S. Machado. 2022. "Using Generative Adversarial Networks to Validate Discrete Event Simulation Models". In *Proceedings of the 2022 Winter Simulation Conference*, edited by B. Feng, G. Pedrielli, Y. Peng, S. Shashaani, E. Song, C. G. Corlu, L. H. Lee, E. P. Chew, T. Roeder, and P. Lendermann, 2748–2759. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Naidoo, K., and V. Marivate. 2020. "Unsupervised Anomaly Detection of Healthcare Providers Using Generative Adversarial Networks". In *Responsible Design, Implementation and Use of Information and Communication Technology*, edited by M. Hattingh, M. Matthee, H. Smuts, I. Pappas, Y. K. Dwivedi, and M. Mäntymäki, 419–430. Cham: Springer International Publishing.

- Nho, Y.-H., S. Ryu, and D.-S. Kwon. 2021. "UI-GAN: Generative Adversarial Network-based Anomaly Detection Using User Initial Information for Wearable Devices". *IEEE Sensors Journal* 21(8):9949–9958.
- NIH 2023. "Circadian Rhythms". <https://www.nigms.nih.gov/education/fact-sheets/Pages/circadian-rhythms.aspx>, accessed 01th May 2023.
- Rashidian, S., F. Wang, R. Moffitt, V. Garcia, A. Dutt, W. Chang, V. Pandya, J. Hajagos, M. Saltz, and J. Saltz. 2020. "SMOOTH-GAN: Towards Sharp and Smooth Synthetic EHR Data Generation". In *Artificial Intelligence in Medicine: 18th International Conference on Artificial Intelligence in Medicine*, edited by M. Michalowski and R. Moskovitch, 37–48. Berlin, Heidelberg: Springer-Verlag.
- Sarwar, A., E. O. Agu, and A. Almadani. 2023. "CovidRhythm: A Deep Learning Model for Passive Prediction of Covid-19 Using Biobehavioral Rhythms Derived From Wearable Physiological Data". *IEEE Open Journal of Engineering in Medicine and Biology* 4:21–30.
- Sebastiani, M., C. Vacchi, A. Manfredi, and G. Cassone. 2022. "Personalized Medicine and Machine Learning: A Roadmap for the Future". *J Clin Med* 11(14):4110.
- Shalom, O. S., G. Uziel, and A. Kantor. 2019. "A Generative Model for Review-based Recommendations". In *Proceedings of the 13th ACM Conference on Recommender Systems*, edited by T. Bogers, A. Said, P. Brusilovsky, and D. Tikk, 353–357. New York, NY, USA: Association for Computing Machinery.
- Shumway, R. H., D. S. Stoffer, and D. S. Stoffer. 2000. *Time Series Analysis and its Applications*, Volume 3. Springer.
- Smith, K. E., and A. Smith. 2019. *Time Series Generation using a One Dimensional Wasserstein GAN*. Spain: Universidad de Granada. Mathematical Methods and Computing.
- Vaccari, I., V. Orani, A. Paglialonga, E. Cambiaso, and M. Mongelli. 2021. "A Generative Adversarial Network (GAN) Technique for Internet of Medical Things Data". *Sensors* 21(11):3726.
- Vallender, S. 1974. "Calculation of the Wasserstein Distance Between Probability Distributions on the Line". *Theory of Probability & Its Applications* 18(4):784–786.
- Vaswani, A., N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. 2017. "Attention is All You Need". In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, edited by I. Murray and C. K. I. Williams, 6000–6010. Red Hook, NY, USA: Curran Associates Inc.
- Xu, L., M. Skoularidou, A. Cuesta-Infante, and K. Veeramachaneni. 2019. "Modeling Tabular Data Using Conditional GAN". In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, edited by H. M. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché Buc, E. B. Fox, and R. Garnett, 7335–7345. Red Hook, NY, USA: Curran Associates Inc.
- Yadav, P., M. Gaur, N. Fatima, and S. Sarwar. 2023. "Qualitative and Quantitative Evaluation of Multivariate Time-series Synthetic Data Generated Using MTS-TGAN: A Novel Approach". *Applied Sciences* 13(7):4136.
- Yoon, J., D. Jarrett, and M. van der Schaar. 2019. "Time-series Generative Adversarial Networks". In *Advances in Neural Information Processing Systems*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, Volume 32, 5508–5518. Curran Associates, Inc.

AUTHOR BIOGRAPHIES

YINING HUANG is a Ph.D. student of the Operations Research Graduate Program at North Carolina State University. Her research interests are simulation and blockchain technologies. Her email address is yhuang43@ncsu.edu.

HONG WAN is an associate professor in the Department of Industrial and Systems Engineering at North Carolina State University. Her research focuses on the areas of experimental design and data analysis of complex simulation models and blockchain. She is the director of the ISE blockchain lab, and serves as the editor in chief of Journal of Blockchain Research and the associate editor of ACM TOMACS. Her email address is hwan4@ncsu.edu and her website is <https://www.ise.ncsu.edu/people/hwan4/>.

XI CHEN is an associate professor in the Grado Department of Industrial and Systems Engineering at Virginia Tech. Her research interests include simulation modeling and analysis, applied probability and statistics, computer experiment design and analysis, and simulation optimization. Her email address is xchen6@vt.edu and her web page is <https://sites.google.com/vt.edu/xi-chen-ise/home>.