

Synthetic Geomagnetic Field Data Creation for Geomagnetic Disturbance Studies using Time-series Generative Adversarial Networks

Anna Zhang, Pooria Dehghanian, Melvin Stevens, Jonathan Snodgrass, Thomas Overbye
Department of Electrical and Computer Engineering
Texas A&M University, College Station, TX
{a-zhang; pooria.dehghanian; mrstevens; snodgrass; overbye}@tamu.edu

Abstract—A key challenge to Geomagnetic Disturbance (GMD) studies is the scarcity of severe geomagnetic field data available to researchers due to its low event occurrence. This study aims to address this challenge by first creating realistic “synthetic” data that represents the geomagnetic field fluctuations caused by recent GMD events. This paper utilizes a machine-learning approach to generate synthetic geomagnetic field data. Specifically, the application and preliminary results of a modified form of the generative adversarial network (GAN) to create time-series synthetic geomagnetic field data of three different severities are described here. The purpose of this paper is to document the first step towards creating severe synthetic geomagnetic field data to advance power system research. Future studies beyond this paper will extend on this work to generate data representing severe GMD storms.

Index Terms—Geomagnetic Disturbances, Magnetometer, Geomagnetic Field, Geomagnetically Induced Current, Machine Learning, Neural Networks, Generative Adversarial Networks.

I. INTRODUCTION

G EOMAGNETIC DISTURBANCE (GMD) simulations and analysis on the power grid are important for furthering innovations for the protection of the bulk power system. Due to their rare occurrence and potential to cause long-term and widespread damage to the electrical infrastructure, GMDs are classified as High Impact, Low Frequency (HILF) events. A significant challenge to GMD studies on the power grid is posed by the scarcity of severe GMD geomagnetic data. The purpose of this work is to address this gap by first generating realistic “synthetic” data that represents the geomagnetic field fluctuations caused by a GMD event. Future studies will extend beyond this work to generate data representing severe GMDs. This work is a part of the larger goal of modeling severe GMD events.

Copyright ©2023 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubspermissions@ieee.org. Presented at the 2023 IEEE Texas Power and Energy Conference (TPEC), College Station, TX, February 2023.

A. Geomagnetic Disturbance Formulation and its Impact on the Power Grid

Geomagnetic disturbances are caused by an injection of charged particles from the sun. The interaction between these charged particles and the earth’s magnetosphere results in a disturbance in the earth’s magnetic field. As described through Faraday’s law of induction, changes in the earth’s magnetic field can induce an electric field across the earth’s surface [1].

In power systems modeling, frequency domain transformations are commonly utilized to compute geomagnetically induced electric fields [2]. The induced electric field is defined as the product of the earth’s surface impedance and magnetic field, as shown,

$$E_X(\omega) = Z(\omega) \frac{B_Y(\omega)}{\mu_0} \quad (1)$$

$$E_Y(\omega) = -Z(\omega) \frac{B_X(\omega)}{\mu_0} \quad (2)$$

where $E_X(\omega)$ corresponds to the northward electric field, $E_Y(\omega)$ corresponds to the eastward electric field, $B_X(\omega)$ corresponds to the northward magnetic field, $B_Y(\omega)$ corresponds to the eastward magnetic field, $Z(\omega)$ corresponds to the earth’s surface impedance, and μ_0 corresponds to the magnetic permeability of free space. By computing the inverse Fourier transform of $E(\omega)$, the time-series electric field, $E(t)$, is obtained as shown,

$$E(t) = \mathcal{F}^{-1}\{E(\omega)\} \quad (3)$$

The presence of this electric field on the electric power grid results in the formation of a quasi-dc voltage, V_{dc} , across transmission lines. The voltage induced on a transmission line can be calculated by integrating $E(t)$ along the incremental length, dl , of the transmission line as shown,

$$V_{dc} = \oint E \cdot dl \quad (4)$$

The induced voltages generate geomagnetically induced quasi-dc currents (GICs) on the power grid that, when uncontrolled, could lead to instability in system voltages, affect operations of relays, protection systems, and cause irreversible physical damage to high-voltage power transformers [3], [4]. To assess the risk of voltage instability, GIC values can be

included into the power flow analysis. Changes in reactive power and bus voltages can be used to determine the likelihood of voltage collapse [5].

B. Historical Context of Geomagnetic Disturbances on the Power Grid

On March 1989, a severe GMD storm moving at a peak variation of 500 nT/min triggered a chain of power system disturbances that ultimately resulted in the collapse of the Hydro Quebec power grid [6]. Further examination of historic storm data suggests that disturbance levels of 5,000 nT/min have occurred as a result of a GMD storm in 1921 [6]. With today’s extensive network of high voltage transmission lines, it is reasonable to expect that a geomagnetic storm of similar intensity would cause unprecedented damage to the grid. These historic events demonstrate the extremity of GMD storms and threat that could be posed on the power grid.

Since the historic GMD storm of 1989, the power sectors have adopted operational procedures to protect against GMDs [6]. However, more research is necessary to consider the impacts more severe storms may have on the power grid so that practical mitigation procedures can be developed to harden the grid against more powerful GMD events.

C. Challenges to GMD Research in Power Systems

The immediate goal of this work is to synthesize GMD data that closely resembles different severity levels of magnetic field fluctuations. Much like hurricanes, tornadoes, and earthquakes, a standard rating of GMD severity has been established under the National Oceanic and Atmospheric Administration (NOAA) Space Weather Scales. Table I provides details on GMD severity and maps the severity scales to their potential impact on critical infrastructure, with G1 being the most minor and G5 being the most severe and thus, associated with the most major impacts [7].

One issue hindering research on GMD analysis on the power grid is the limited quantity of severe GMD data. Although geomagnetic field data can be acquired through the North American Electric Reliability Corporation (NERC) and other observatories supporting academic research, such as the INTERMAGNET network [8], data capturing fluctuations caused by GMD of severity G5 is very limited. To the authors’ best knowledge, the last known G5 storm to strike earth occurred in late-October 2003 (with a speed of 300 nT/min [6]).

This study is a part of the larger goal of synthesizing data that represents severe GMDs. A method to achieve this goal

is to first generate data that is statistically similar to presently available geomagnetic field data, and then scale up magnetic field fluctuations for severe GMD representations. Future studies will extend beyond this work to scale-up magnetic field changes of synthetic data sets to represent severe storms. This paper serves to document a machine-learning approach, or more specifically, the Time-series Generative Adversarial Network (TimeGAN), to generate realistic synthetic geomagnetic field data.

The rest of the paper is structured as follows. Synthetic data generation and validation metrics are fully explained in Section II. Experimental results and analysis are presented in Section III. Section IV discusses challenges faced during the study. Finally, Section V concludes with a summary of the work.

II. SYNTHETIC DATA GENERATION AND VALIDATION METHODS

The goal of this section is to provide an overview of the method and tools used for creating synthetic geomagnetic field data as well as explain the metrics used to guide evaluations of the quality and accuracy of synthetic data.

A. Generative Adversarial Networks (GAN)

GAN is a relatively new unsupervised machine learning algorithm first proposed by [9] in 2014 to generate very realistic high-resolution data. The basic structure of GAN includes two machine learning models, called the generator network and the discriminator network. The generator assumes an unstructured prior distribution p_z over a random noise vector z that would be fed into the generator function, $G(z)$. The goal for the generator is then to learn the function $G(z)$ that would transform unstructured samples z into realistic samples that can be drawn from the probability distribution of the real data set, p_r . The role of the discriminator is to learn a function $D(x)$ that can correctly distinguish between data synthesized by the generator and data used for training. During the learning process, the discriminator draws samples from the real data set, x , and is trained to classify the data as “real”. Similarly, the discriminator learns the statistics of data generated by the generator and is trained to classify this data as “fake”.

The mechanics of GAN can be viewed as a competition between the generator and discriminator. The generator’s goal is to generate realistic data that would fool the discriminator while the discriminator’s goal is to correctly classify data from the generator as counterfeit. If the discriminator can distinguish between real and generated content, the generator is penalized, otherwise, the discriminator itself is penalized while the generator is rewarded. This competition drives both networks to improve their strategies until the synthetic data becomes indistinguishable from the training data. The main objective of this game can be described by the following objective function [9]–[11],

$$\underset{G}{\text{Min}} \underset{D}{\text{Max}} V(G, D) = \mathbb{E}_{x \sim P_r} [D(x)] - \mathbb{E}_{z \sim P_z} [D(G(z))] \quad (5)$$

TABLE I
NOAA SPACE WEATHER SCALES FOR GEOMAGNETIC STORMS

Scale	Description	Impacts on Power Systems	Ave. Frequency (1 cycle=11 years)
G5	Extreme	Widespread voltage control issue; System collapse and blackouts	4 days per cycle
G4	Severe	Possible voltage issue; Miss operation of relay tripping	60 days per cycle
G3	Strong	Voltage correction needed; False alarms triggering on protection	130 days per cycle
G2	Moderate	Damage to high-latitude power system assets	360 days per cycle
G1	Minor	Weak power grid fluctuations	900 days per cycle

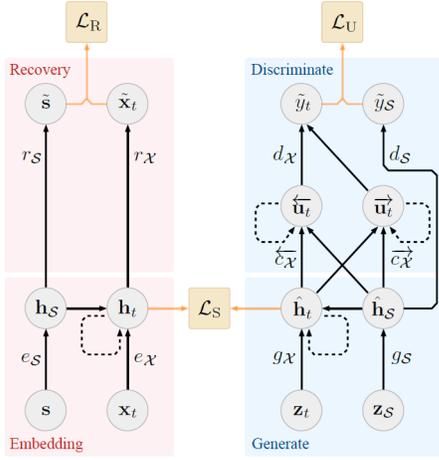


Fig. 1: Blocks to the TimeGAN model [12].

Although traditional GAN can be extremely successful in producing highly reliable data, its performance on learning temporal correlations unique to time-series data is insufficient. Fortunately, the popularity of GANs among the research and academic realm has yielded rapid improvements in GAN algorithms to accomplish unique objectives. This study seeks to test and evaluate TimeGAN proposed by [12] in generating synthetic geomagnetic field data. In this framework, the temporal dependencies of the real data are preserved in the model by merging the flexibility of the unsupervised GAN with the control offered by supervised training in autoregressive networks. TimeGAN introduces two new functions, called the embedding and recovery functions, whose purpose is to provide mappings between feature and latent representations. This allows the adversarial network to learn conditional temporal distributions of the data via a lower-dimensional space. A basic block diagram of TimeGAN is illustrated in Fig. 1. Further details on the TimeGAN architecture can be viewed here [12].

For this study, YData Synthetic [13], a publicly available python-based package based on [12], was implemented to generate and validate time-series synthetic data.

B. Data Set

This study utilized magnetometer data collected by an INTERMAGNET [8] observatory stationed at Boulder, Colorado, USA. Data from one magnetometer station is used in this study to avoid introducing spatially-dependent field variations. The magnetometer samples the geomagnetic field at a cadence of one sample per minute.

Nineteen hours of magnetic field, \mathbf{B} , data from November 3, 2021 to November 4, 2021 to capture fluctuations from G1 to G3 category GMD events, and five days worth of data from September 6, 2017 to September 10, 2017 to capture a G4 category GMD event, were used as inputs to train separate TimeGAN models representing GMD events of different severity levels.

One of the questions this study intended to answer was whether it is more difficult for TimeGAN to train from a data

set containing more severe magnetic field fluctuations. Hence, the time series data sets were partitioned based on their respective GMD severity category before using them as input to train the TimeGAN model. Data sets from the G1 and G2 categories were combined as a single input to the TimeGAN model; this set contained 150 samples and represented 2.5 hours of GMD duration. The data set from the G3 category contained 960 samples, representing 16 hours, and the data set from the G4 category contained 7,200 samples, representing 120 hours or five days of magnetic field fluctuations. Importantly, time-dependent variations of the X , Y , and Z components of \mathbf{B} were treated as separate "features" of the magnetic field data set.

Prior to being used as input for training the TimeGAN synthesizer, the data sets were pre-processed. Data pre-processing involved the following steps:

- 1) Each element of the time-series feature vector was scaled to be within the range of $[0, 1]$.
- 2) Sliding window technique was implemented across all time steps to create subsets, or input sequences, of the original time-series data set. All subsets combined formed the final data set used to train a TimeGAN model.
- 3) The final data set were then then "randomized" so that subsets were placed in no particular order. This step allows the data set to mimic independent and identically distributed (IID) sets. After randomization, temporal dynamics would still be preserved within each subset. Thus, these subsets may also be called "input sequences" in this paper.

C. Evaluation Metrics for Synthetic Data

Two metrics commonly used to assess the quality of output from neural network models were considered here, diversity and predictability.

Diversity: The desired outcome is for the synthesized data and real data to share similar feature distribution. After synthetic data generation, the diversity and feature distributions of the synthetic data were compared against real data using two dimensionality reduction techniques called, Principle Component Analysis (PCA) [14] and t-distributed stochastic neighbor embedding (t-SNE) [15]. PCA and t-SNE analyses were conducted to provide a 2-dimensional visual representation of how close synthetic distributions resemble actual distributions observed by the magnetometer. Though PCA and t-SNE are both dimensional reduction techniques commonly used for data visualization and analysis, they are used in conjunction here for correlation purposes. Unlike the more ubiquitous PCA, t-SNE uses a non-linear approach to reducing data complexity, and thus is capable of capturing more complex relationships between features. If the synthetic data completely inherits the feature distribution of the real data set, then the expected observation would be a complete overlap between synthetic and real data dimensionality reduction results.

Predictability: Ideally, the synthesized data should inherit the predictive characteristics of the real data. The quality and

usefulness of synthetic data can be determined by how well-matched its predictive characteristics are to those of the real input data. To evaluate this, a 2-layer GRU-based recurrent neural network (RNN) was trained on the first 50 percent of the synthetic data set and tested against the last 50 percent of the real data (“train on synthetic, test on real” framework [12]). The trained RNN model was then used to predict future values (“synthetic predictions”). Finally, predictions for synthetic data were then evaluated by use of the coefficient of determination R^2 and the mean absolute error (MAE) scores. Note that to ensure the RNN model is of good quality (so as to not impact the results and analysis of the trained TimeGAN model), the RNN model was also trained on the corresponding real data to form predictions (“real predictions”) and then evaluated accordingly.

R^2 is a measure of goodness-of-fit between predicted and true values, and thus can be used to indicate how well the trained RNN model can form predictions based on the input data. If the RNN model carries excellent prediction capability and the synthetic data perfectly inherits predictive characteristics of the real data set, then both R^2 scores for synthetic predictions and real predictions should be very close to 1.0.

MAE is a model evaluation metric used to compute the magnitude of error between predicted data and real data. Again, if the quality of the RNN model is good and both synthetic and real data sets share the same predictive characteristics, then the MAE values are expected to be very close to *zero*. Mathematically, estimations for R^2 and the MAE over n samples are represented by equations (6) and (7), respectively [16].

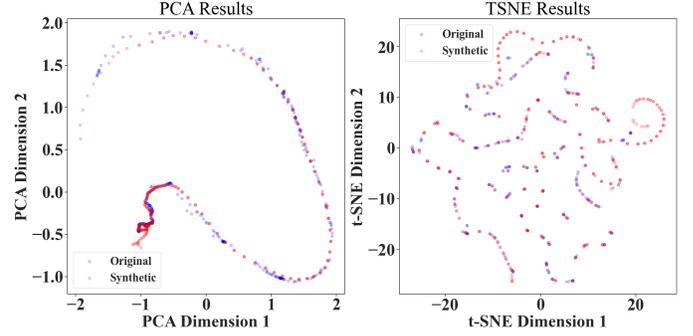
$$R^2(y, \hat{y}) = 1 - \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2} \quad (6)$$

$$MAE(y, \hat{y}) = \frac{1}{n} \sum_{i=0}^{n-1} |y_i - \hat{y}_i| \quad (7)$$

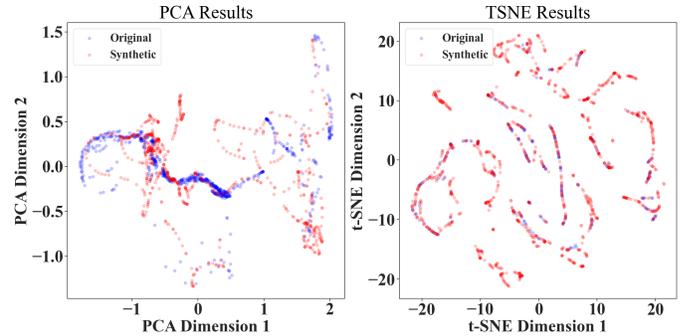
where y_i represents the true value of the i th element in the sample set, \hat{y}_i is the corresponding predicted value, and \bar{y} is the average value of y_i over over n samples.

III. EXPERIMENTAL RESULTS AND ANALYSIS

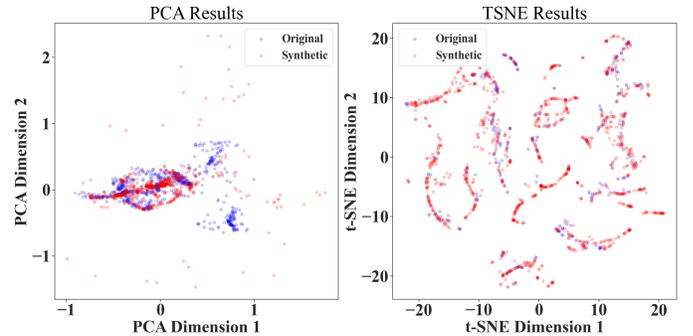
The objective of this section is to present and explain experimental observations. Unless otherwise noted, results presented reflect the following model hyperparameters and input properties for training TimeGAN: TimeGAN model was customized to have 128 network layers and 24 hidden dimensions; a learning rate of $5e-4$ was used; the batch size of 128 was used for training G3 and G4 data, and 120 for G1-G2; input noise dimension for the generator network was 32; input vector length for each subset is 30; finally, a maximum number of 50,000 training steps was used to train the model. These parameters were selected to optimize model performance while reducing computation time.



(a) PCA and t-SNE visualizations on G1 to G2 magnetic field data.



(b) PCA and t-SNE visualizations on G3 magnetic field data.



(c) PCA and t-SNE visualizations on G4 magnetic field data.

Fig. 2: PCA and t-SNE visualizations on synthetic data generated by TimeGAN (in red) and real magnetic field data (in blue). Each row provides visualizations for different severity GMD events.

A. Statistical Visual Analysis

As displayed in Fig. 2, PCA and t-SNE visualizations show good overlap between synthetic and real data across all three severity types, thus demonstrating that feature distributions between synthetic and real data are similar. For G1 to G2 data sets, it is observed that synthetic data shows better overlap with real data when compared with G3 and G4 partitions using PCA analysis. This observation could be explained by the fact that G1 to G2 magnetic field fluctuation magnitudes are less than those inherent to more severe storms such as G3 and G4, and thus making its distributions and statistics the easiest for TimeGAN to learn.

TABLE II
PREDICTION RESULTS WITH A 2-LAYER RNN MODEL

Prediction Scores using G4 Input Data			
Training Samples (Real or Synthetic)	R ²	MAE	# Epochs Required
Real	0.948	0.009	253
Synthetic	0.962	0.008	125
Prediction Scores using G3 Input Data			
Training Samples (Real or Synthetic)	R ²	MAE	# Epochs Required
Real	0.879	0.037	704
Synthetic	0.843	0.034	387

B. Quantitative Predictive Scoring

Sixty samples of G1-G2, 465 samples of synthetic G3, and 3,585 samples of synthetic G4 data were used to separately train a 2-layer RNN model to be used for prediction. Predictions for real data were used as the control to ensure the quality of the 2-layer RNN sequence-prediction model.

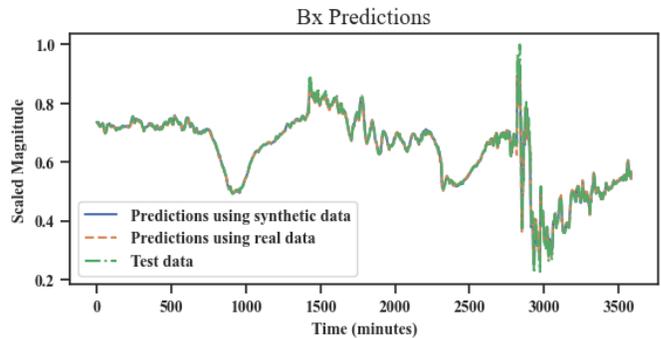
As shown in Table II, the prediction scores that assess how reliable synthetic data is in capturing temporal dynamics of the real data are remarkably impressive. Prediction scores for the G3 synthetic data are slightly poorer than that of G4, but this could be attributed to the smaller number of G3 training data available for the RNN model from which to learn. For both G3 and G4, the real and synthetic predictive scores are very similar in value, indicating excellent inheritance of temporal characteristics.

The number of epochs required for predictions made by the RNN model to converge to ground truth data is also shown for the record. Note that due to the small number of training samples available, the RNN-model was unable to produce reliable predictions for the G1-G2 scenario, and thus prediction results are not shown here.

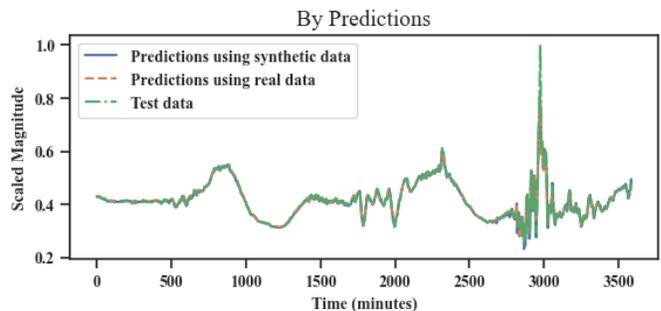
A complimentary graphical representation of prediction results using G4 magnetic field data is shown in Fig. 3. Although Fig. 3c shows an area where a large spike was underestimated in predictions, predicted curves generally demonstrate remarkable agreement with the ground truth test data, indicating that the 2-layer RNN is capable of forming reliable predictions. Most importantly, predictions on real and on synthetic data are almost completely aligned, indicating excellent inheritance of predictive characteristics from real data to synthetic data produced by TimeGAN.

C. The General Time-series Synthetic Data

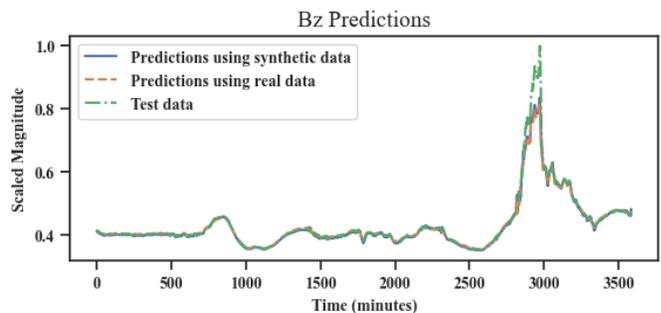
A comparison of sequence vectors for the time-series real and synthetic magnetic field data is shown in Fig. 4. The synthetic G3 data shown in Fig. 4a was generated using a training set with an input sequence length of 320, whereas the synthetic G1-G2 data shown in Fig. 4b used an input sequence length of 60. It was observed that the synthetic data for the G3 partition captures the general shape of the real magnetic field, but misses some higher-frequency temporal dynamics. On the other hand, synthetic data for the G1-G2 partition appears to inherit higher frequency fluctuations and also captures the



(a) Predictions on scaled B_X data.



(b) Predictions on scaled B_Y data.



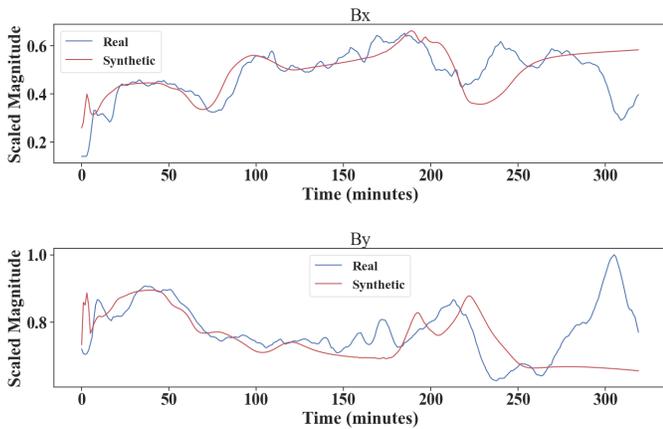
(c) Predictions on scaled B_Z data.

Fig. 3: Predictions using scaled G4 magnetic field data. Extrapolations based upon synthetic and real training data are both compared against corresponding data from the actual test data set.

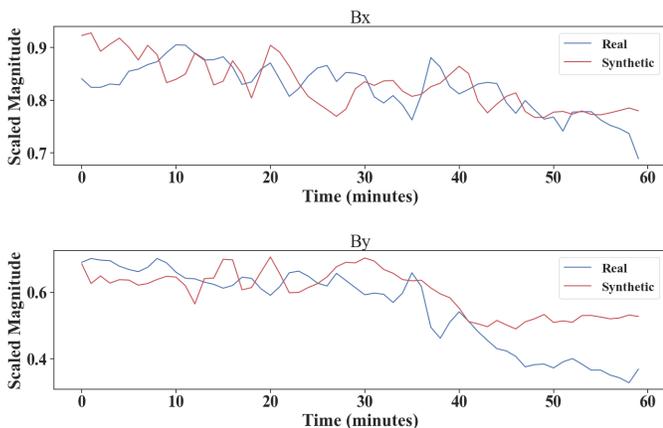
general temporal pattern of the real data better than that of 4a. This could be attributed to using a much shorter input sequence vector length (60 vs. 320) to train the TimeGAN model. A longer vector length may contain more temporal dynamics per input sequence which may have necessitated more training steps in order for TimeGAN to generate similar quality synthetic data than the authors allowed.

IV. DISCUSSION AND FUTURE WORKS

The larger goal is to synthesize data that mimics severe GMD storms, but the scarcity of geomagnetic field data from severe GMD events makes this challenging. Although neural networks show promise in modeling input environments, they may not be capable of capturing all of the important physical



(a) B_X and B_Y for a G3 Storm



(b) B_X and B_Y for a G1-G2 Storm

Fig. 4: Comparison of sequence vectors for the scaled time-series real and synthetic magnetic field data.

boundaries just by training on available data. Moreover, generating severe GMDs might not be possible without incorporating physics that dictates the true distribution of geomagnetic field fluctuations from severe GMDs. A potential avenue for future investigation could be incorporating physics that determines geomagnetic field boundaries into neural networks to create realistic models for severe GMDs.

Within the space weather and geophysics community, there has been a number of studies modeling historically severe GMD events. One such study published by Winter et al. constructs high-resolution magnetic fields representing the 1859 Carrington event (the most extreme GMD storm ever recorded) by utilizing modern high-resolution magnetometer observations of less severe GMD events and low-resolution magnetometer observations of the Carrington event. Thus, another potential application of TimeGAN would be to synthesize more variety of data representing severe GMD events by training on magnetic fields constructed by [17].

A challenge to using TimeGAN is balancing model capability with computational time. To maximize the probability of success in generating statistically accurate data, it may

become necessary to increase model complexity and the size of the training data set. Both approaches necessitate longer computation times, which hinder efficient optimization of model hyperparameters and generally slow down experiments. For example, training the embedding, supervised and joint networks using 50,000 training steps on an input data size of 7,170 input sequence vectors of dimensions 30×3 required approximately 29 hours when utilizing the NVIDIA GeForce RTX 2080 Ti GPU.

V. CONCLUSIONS

This study explores the use of a novel machine learning approach called TimeGAN to generate synthetic data that adheres to statistics of actual geomagnetic fields produced by a GMD. The quality and reliability of synthetic data sets were validated using statistical visualization techniques as well as prediction scoring. Empirical results demonstrate that the synthetic data successfully captures feature distributions and temporal dynamics of the input geomagnetic field samples, showing promise for TimeGAN to be used for producing a variety of synthetic geomagnetic scenarios to be used for GMD simulations and analysis. This study serves as a part of the larger goal to model severe GMD events.

REFERENCES

- [1] "Geomagnetic storms." [Online]. Available: <https://www.swpc.noaa.gov/phenomena/geomagnetic-storms>
- [2] R. Pirjola, "Review on the calculation of surface electric and magnetic fields and of geomagnetically induced currents in ground-based technological systems," *Surveys in Geophysics*, vol. 23, 2002.
- [3] V. D. Albertson, J. M. Thorson, and et al., "Solar-induced-currents in power systems: Cause and effects," *IEEE Transactions on Power Apparatus and Systems*, vol. PAS-92, no. 2, 1973.
- [4] P. Dehghanian and T. J. Overbye, "Temperature-triggered failure hazard mitigation of transformers subject to geomagnetic disturbances," in *2021 IEEE Texas Power and Energy Conference (TPEC)*, 2021.
- [5] T. J. Overbye, T. R. Hutchins, and et al., "Integration of geomagnetic disturbance modeling into the power flow: A methodology for large-scale system studies," in *2012 North American Power Symposium (NAPS)*. IEEE, 2012.
- [6] G. Cauley and M. Lauby, "High-impact low-frequency event risk to the North American bulk power system," *North Amer. Electr. Rel. Corp.(NERC), Atlanta, GA, USA, Tech. Rep.*, 2010.
- [7] "NOAA space weather scales." [Online]. Available: <https://www.swpc.noaa.gov/noaa-scales-explanation>
- [8] D. Kerridge, "Intermagnet: Worldwide near-real-time geomagnetic observatory data," in *Proceedings of the Workshop on Space Weather, ESTEC*, vol. 34, 2001.
- [9] I. Goodfellow, J. Pouget-Abadie, and et al., "Generative adversarial networks," *Communications of the ACM*, vol. 63, no. 11, 2020.
- [10] C. Zhang, S. R. Kuppannagari, and et al., "Generative adversarial network for synthetic time series data generation in smart grids," in *IEEE Int. Conf. on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*, 2018.
- [11] X. Zheng, B. Wang, and L. Xie, "Synthetic dynamic PMU data generation: A generative adversarial network approach," in *Int. Conf. on Smart Grid Synchronized Measurements and Analytics (SGSMA)*, 2019.
- [12] J. Yoon, D. Jarrett, and M. Van der Schaar, "Time-series generative adversarial networks," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [13] YData AI, "ydata-synthetic," <https://github.com/ydataai/ydata-synthetic>. git, 2020.
- [14] F. B. Bryant and P. R. Yarnold, "Principal-components analysis and exploratory and confirmatory factor analysis." 1995.
- [15] L. Van der Maaten and G. Hinton, "Visualizing data using t-SNE." *Journal of machine learning research*, vol. 9, no. 11, 2008.

- [16] F. Pedregosa, G. Varoquaux, and et al., "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, 2011.
- [17] L. M. Winter, J. Gannon, and et al., "Spectral scaling technique to determine extreme Carrington-level geomagnetically induced currents effects," *Space Weather*, vol. 15, no. 5, 2017.