JOURNAL OF
EMERGING INVESTIGATORS

Article

Evaluating the effectiveness of synthetic training data for day-ahead wind speed prediction in the Great Lakes

Alexander Wycoff[1], Jonathan Stelman[2]
[1] Eden Prairie High School, Eden Prairie, Minnesota
[2] Ford Foundation, New York, New York

SUMMARY

With an estimated offshore potential wind energy capacity of 575 gigawatts, the Great Lakes region is a promising area for future wind energy development. Electric utilities that use wind energy rely on accurate day-ahead wind energy forecasts, mainly informed by predicted wind speed, to account for the variability of wind energy production. Since wind turbines have not yet been developed in the Great Lakes, observational wind data is collected at very few sites. Hence, utilities may seek to use synthetic wind data to pre-train accurate day-ahead wind speed prediction models. Yet, prior studies have only utilized synthetic data in predicting sub-hourly winds speeds. We hypothesized that training long short-term memory neural networks to predict day-ahead wind speeds on synthetic wind data instead of observational wind data would increase accuracy since synthetic wind data is available over longer time spans than observational wind data. We used observational data from the Lake Michigan Wind Assessment and synthetic data from the National Offshore Wind Great Lakes dataset and found that networks trained on synthetic data had a lower mean absolute percentage error score in day-ahead wind speed prediction than networks trained on observational data. We also optimized additional parameters of the networks for both synthetic and observational network types, further improving accuracy. The availability of a wind speed prediction model trained on synthetic data will reduce reliance on historical observational data at future sites of wind energy infrastructure, allowing utilities to swiftly adapt accurate prediction methods to new sites.

INTRODUCTION

The Great Lakes have an estimated offshore wind energy potential of 575 gigawatts (1). With most sites throughout the Great Lakes reaching annual average wind speeds of nine meters per second or greater, the Great Lakes region has significant opportunities for offshore wind energy development (2). However, challenges remain, such as the potential for ice cover causing ice jamming of turbines and the limited width of current ports that restrict the maximum size of wind turbines on certain lakes (3). Yet, offshore wind energy development in the Great Lakes has also been deemed technically and economically feasible in some areas (3). Wind energy development in the Great Lakes would also assist states in meeting their clean energy goals and provide economic benefits to nearby population centers (1). For example, government mandates such as New York's Climate Act have spurred demand for new renewable energy projects, in a bid to reduce greenhouse gas emissions (3). These new sources of renewable energy will be used to heat buildings, power electric transportation methods, and reduce industrial and agricultural emissions (3). The strong wind energy resources available in the Great Lakes have sparked interest in the development of offshore wind energy to meet emission reduction goals (3). When wind energy development begins at larger scales in the Great Lakes, electric utilities will need to be able to integrate this new source of energy effectively.

Researchers have proposed pathways to bring commercial wind energy to the Great Lakes within the next decade (2). However, the variability of wind energy production can make its integration burdensome, as electric utilities must adapt to changes between forecasted and realized wind energy (4). Grid operators rely on accurate day-ahead wind speed forecasts to make decisions about energy pricing, and to balance energy production from other sources (5). Since wind energy production is variable by nature, accurate wind energy forecasts can assist utilities in integrating wind energy reliably (6). Thus, more accurate wind speed predictions would allow utilities to improve wind energy production estimates, enhancing the reliability of wind energy. Overprediction of wind speeds forces utilities to compensate for shortfalls in production by sourcing energy from more expensive sources, whereas underprediction can result in energy waste as more electricity is produced than is necessary (6). Previous projects to improve the forecasting of wind speed production have demonstrated that improvements in predictive wind speed models can reduce wind energy overprediction and underprediction, directly corresponding to a decrease in the excess costs incurred by electric utilities (6).

A comparison of previous studies found that deep learning approaches have surpassed traditional machine learning methods, obtaining a lower mean absolute percentage error (MAPE) in wind speed predictions on prediction timescales ranging from one minute to six hours (7). One deep learning method, the long short-term memory (LSTM) network, is a recurrent neural network that is particularly well-suited to modeling short- and long-term dependencies in time series data due to its low error in multi-step ahead predictions (8). Two previous studies have utilized multiple variations of LSTM networks to predict wind speeds (9-10). However, because the locations, data, and forecasting periods involved in each study varied dramatically, their comparability to each other is limited (9-10). Additionally, both studies relied only on observational data (9-10). Because wind turbines have not yet been developed in the Great Lakes, the observational wind

https://doi.org/10.59720/24-307

data typically needed to train wind speed prediction models may not yet be available at potential wind turbine locations. However, synthetic wind data, which is developed through mathematical simulations and has a high spatial resolution, could be used to pre-train wind speed prediction models when observational wind data is not available (5). Synthetic data is available at a spatial resolution of two kilometers throughout the entire Great Lakes region, making it much more widely available than observational data, and has been confirmed to realistically represent observational wind speed data over larger timescales (5).

While previous studies have used synthetic wind data to improve sub-hourly wind speed predictions, it is unknown whether synthetic wind data is beneficial for day-ahead wind speed prediction (5). In our study, we aimed to create more accurate day-ahead wind speed prediction models using synthetic wind data. We hypothesized that LSTM neural networks trained on synthetic wind data would have a lower MAPE in predicting day-ahead wind speeds than LSTM neural networks trained on observational wind data. We leveraged synthetic data from the National Offshore Wind (NOW-23) Great Lakes dataset generated through version 4.2.1 of the Weather Research & Forecasting (WRF) program and observational data captured during the Lake Michigan Wind Assessment off the coast of Muskegon, Michigan (11-12). We used multiple experiments to determine parameters for a comparison between networks trained on observational and synthetic wind data. We determined that LSTM networks trained with 50 epochs and a batch size of 8 resulted in the lowest variance for networks trained on observational wind data. We found that networks trained on synthetic data had a lower mean MAPE score in day-ahead wind speed prediction than networks trained on observational data. When little or no observational data is available, LSTM networks trained on synthetic data could be more accurate in day-ahead wind speed prediction. The usage of synthetic data to pre-train day-ahead wind speed prediction models could assist utilities in integrating new sources of wind energy developed in the Great Lakes.

RESULTS

We aimed to evaluate the effectiveness of synthetic data for day-ahead wind speed prediction in the Great Lakes, hypothesizing that LSTM neural networks trained on synthetic wind data would have a lower MAPE than networks trained on observational wind data. To test our hypothesis, we developed and tuned an LSTM network for time series prediction using data from either the Lake Michigan Wind Assessment observational data (April 2013 to November 2013) or the synthetic wind data from the NOW-23 Great Lakes dataset (January 2000 to November 2013) (11-12). For each network, MAPE scores were found by testing on the final month of data from the Lake Michigan Wind Assessment (December 2013). Data from the 2013 Lake Michigan Wind Assessment was taken from a buoy approximately ten kilometers from the eastern shoreline of Lake Michigan near Muskegon, Michigan, and data from the NOW-23 Great Lakes dataset was taken from the closest site to this buoy, approximately 520 meters away (**Figure 1**). The networks used the previous 24 hours of wind speed and direction data at a temporal resolution of 1 hour to predict the wind speed 24 hours later at a given site.

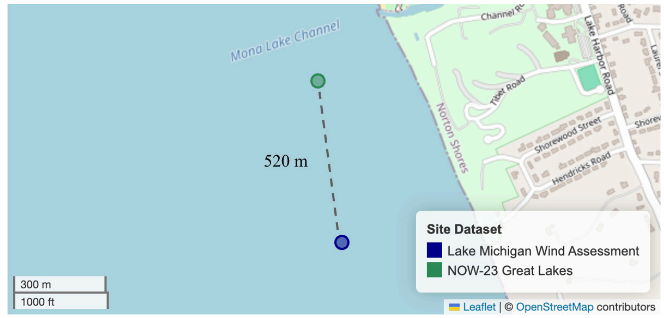

Figure 1: Map of Muskegon, Michigan with site locations. The physical locations of the sites selected from the NOW-23 Great Lakes dataset (synthetic data) and Lake Michigan Wind Assessment (observational data) are displayed over a map of Muskegon, Michigan. Map data is available from openstreetmap.org under the Open Database License.

Parameter optimization

Through multiple experiments, we discovered that using two or fewer years of synthetic training data resulted in higher MAPE and thus were not comprehensive enough to train an LSTM network effectively, while using 10 or more years of training data yielded diminishing returns in accuracy (**Figure 2**). Networks with ten or more years of training data plateaued around 29.3% MAPE, while networks with two or fewer years had over 30% MAPE (**Figure 2**). Furthermore, we observed that larger increases in the number of epochs used to train networks on synthetic data also had diminishing returns in accuracy, with the best accuracy being 28% MAPE with 100 epochs, and that additional epochs worsened accuracy in models trained on observational wind data (**Figures 3, 4**).

We also experimented with the batch size used to train networks, discovering that batch size had little effect on the accuracy of networks trained on synthetic data but was influential in determining the accuracy of networks trained on observational wind data (**Figures 4, 5**). For the observational network type, the MAPE scores for batch sizes of 8 and 128, respectively, were 30.5% and 33.6%, a 3.1% difference, while

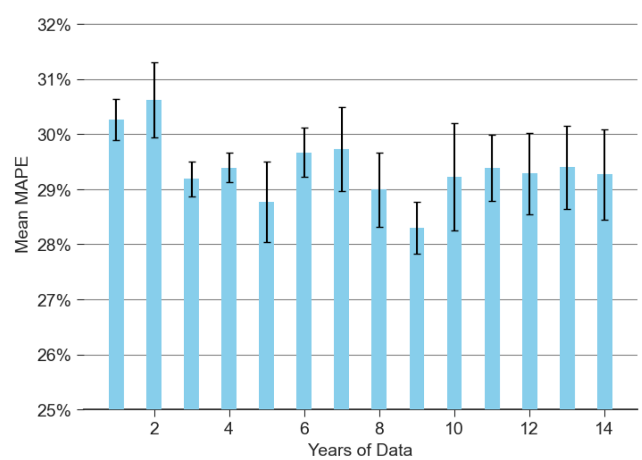

Figure 2: Effect of years of data on network mean absolute percentage error (MAPE) scores. The mean MAPE scores compared to years of synthetic data from the NOW-23 dataset used to train long short-term memory (LSTM) networks (n = 10). LSTM networks were trained for 50 epochs with a batch size of 128. Error bars represent the standard deviation of the MAPE scores.

JOURNAL OF
EMERGING INVESTIGATORS
https://doi.org/10.59720/24-307

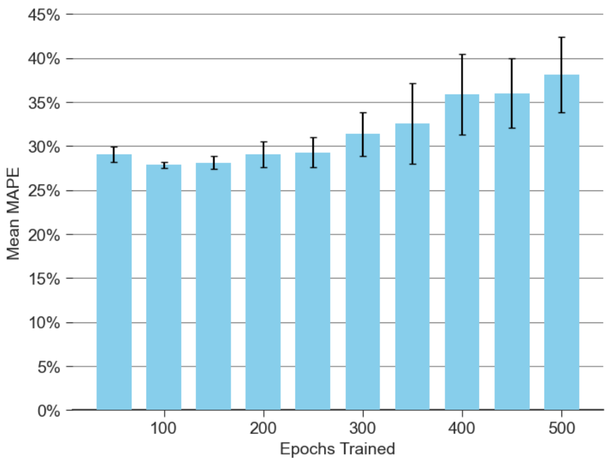

Figure 3: Effect of the number of epochs trained on network mean absolute percentage error (MAPE) scores. Mean MAPE scores for the number of trained epochs using synthetic data from the NOW-23 dataset. Long short-term memory networks (n = 10) were trained on synthetic data with a batch size of 128. Error bars represent the standard deviation of the MAPE scores.

there was a difference of only 0.5% for the synthetic network type (**Figure 4**). Thus, networks trained on observational wind data often overfit to training data at higher numbers of epochs because of the limited amount of data and require smaller batch sizes to better learn their datasets. Due to this, we used 50 epochs and a batch size of 8 to train networks on observational and synthetic wind data for comparison since these parameters resulted in the lowest standard deviation in MAPE of 1.09% for the observational network type (**Figure 4**).

Network comparison

We evaluated the MAPE scores for each network on the final month of data from the Lake Michigan Wind Assessment (December 2013), which was withheld from training. While the mean MAPE of LSTM networks trained on synthetic wind data was lower than that of networks trained using observational data, it was not by a large margin (**Figure 6**). To determine statistical significance, we recreated our experiment by retraining each network type 30 times to compile a distribution of MAPE scores for both (**Figure 6**). The sample mean MAPE of networks trained on observational and synthetic wind data were 30.66% and 29.06%, respectively. The standard deviation of MAPE scores for networks trained on observational and synthetic wind data were 1.09% and 0.79%, respectively. A *z*-test for skewness showed the MAPE score distribution of networks trained on observational data was significantly different from a normal distribution (z = 4.133, $p < 0.001$). We observed that LSTM networks trained on synthetic data had a lower mean MAPE score in day-ahead wind speed prediction than networks trained on observational data (Welch's *t*-test, $t = 6.395$, $p < 0.001$).

As a benchmark, we also compared the LSTM networks trained on synthetic data to a persistence model, which uses the last known wind speed measurement to predict the next. The LSTM networks trained on synthetic data were more effective, with their MAPE score of 29.06% being 35% lower than the mean score of 44.95% achieved by a persistence model.

DISCUSSION

Our study aimed to utilize synthetic data from the NOW-23 Great Lakes dataset and LSTM networks to create a practical wind speed prediction model. We hypothesized that LSTM neural networks trained on synthetic wind data would have a lower MAPE in predicting day-ahead wind speeds than LSTM neural networks trained on observational wind data. We tested LSTM networks using data from the Lake Michigan Wind Assessment and found convincing evidence for our hypothesis that LSTM networks trained using synthetic data had a lower MAPE in day-ahead wind speed prediction than those trained using observational data ($p < 0.001$). We found three to ten years of synthetic training data to be optimal and a fewer number of epochs helpful in reducing overfitting.
Although the use of dropout layers and regularization help to reduce overfitting, we still saw network overfitting based on some training parameters (13). For networks trained on

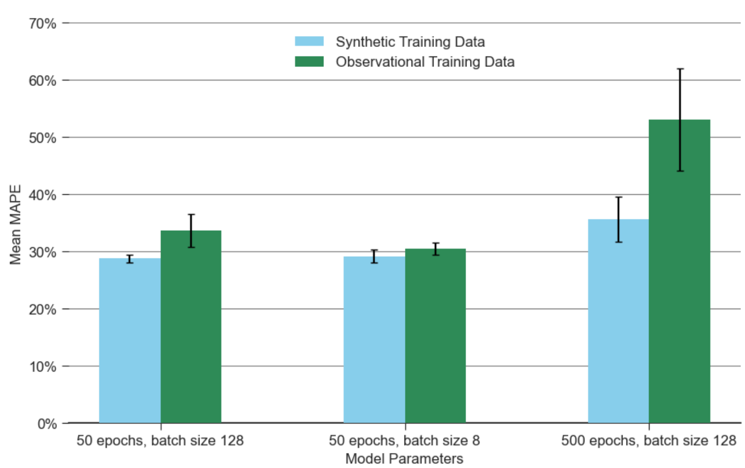

Figure 4: Comparison of mean absolute percentage error (MAPE) scores for models with varying parameters. Mean MAPE scores for networks trained with synthetic data from the NOW-23 dataset or observational data from the Lake Michigan Wind Assessment over different epochs and batch sizes. Long short-term memory networks (n = 30) were trained for each training data type and parameter combination. Error bars represent the standard deviation of the MAPE scores.

https://doi.org/10.59720/24-307

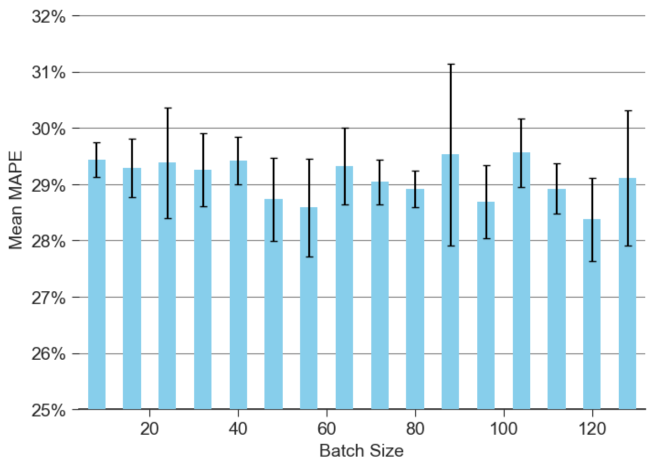

Figure 5: Effect of batch size on network mean absolute percentage error (MAPE) scores. Mean MAPE scores for networks trained with data from the NOW-23 Great Lakes dataset. Long short-term memory networks (n = 10) were trained for 50 epochs. Error bars represent the standard deviation of the MAPE scores.

synthetic data, we found that using two or fewer years of training data lowered accuracy, indicating that using this little synthetic training data may lead to network overfitting and might not be a representative enough sample of wind patterns. For both network types, we also found that using greatly beyond 50 epochs, or iterations of training on an entire dataset, led to increased error, particularly when networks were trained on observational data. This indicates that the networks became overfit to the training data at a higher number of epochs.

The difference in performance between the synthetic and observational network types might be explained by the larger number of samples in the synthetic dataset than the observational dataset and the inclusion of the relevant predictive month in the synthetic dataset. The synthetic data contained a total of about 184,000 observations, while the observational data contained only about 5,700. Furthermore, the relevant predictive month of December was contained in only the synthetic wind data, as the observational data spanned only April 2013 to November 2013. Since wind patterns are typically subject to seasonal variations, the representation of patterns during the relevant predictive month may have contributed to the increased performance seen with synthetic data.

Previous studies have approached wind speed prediction using machine learning, deep learning, and artificial intelligence (7). Yet, these approaches have typically been confined to areas where wind speed data is historically available, limiting the extent to which they can be applied. A previous study focused on day-ahead wind speed prediction with LSTM networks reported a maximum improvement in mean absolute error over a persistence model of approximately 17% across all models tested (10). Persistence models use the last known wind speed measurement to predict the next and are a typical benchmark for the performance of wind speed prediction models. When evaluated on observational data withheld from training, our LSTM networks trained on synthetic data had a mean MAPE score approximately 35% lower than that obtained using a persistence model. However, a direct comparison of our results to this study's is not possible since it was set at an unrealistic elevation for wind turbines of 20 meters (10). As of 2018, the average hub height for turbines in the United States was about 88 meters (14).

Considering the data available, our study is limited in scale and generalizability. While we considered day-ahead predictions with an hourly sampling rate in this study, future studies could train additional models at different prediction timescales, compare models trained on data sampled at different rates, and train similar networks on different prediction timeframes to suit various applications. We also considered only one region from the NOW-23 dataset, so further studies on this topic may also seek to extrapolate results to the additional regions of the NOW-23 dataset. Additionally, future research could utilize the optimizations to network parameters made in this study to improve the accuracy of the LSTM model presented. Nevertheless, by its nature, synthetic data is limited in its realism to observational data. While it realistically represents observational data at the timescale used in this study, it is not a perfect indicator of actual wind features (5). Synthetic data may not always be

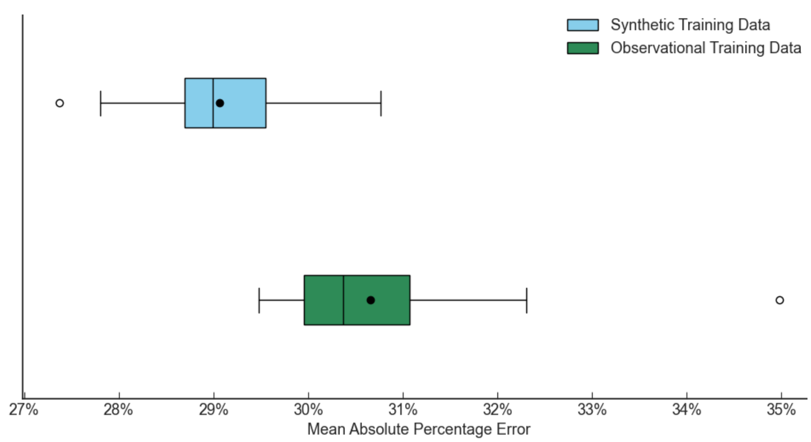

Figure 6: Box plots of distributions of mean absolute percentage error (MAPE) scores. The distribution of mean MAPE scores for long short-term memory (LSTM) networks (n = 30) trained with synthetic data from the NOW-23 dataset or observational data from the Lake Michigan Wind Assessment. The center line denotes the median. The black circles denote means and white circles denote outliers in each distribution. Network types were trained using 50 epochs and a batch size of 8. Welch's t-test shows a significant difference in MAPE scores between network types ($t = 6.395$, $p < 0.001$).

https://doi.org/10.59720/24-307

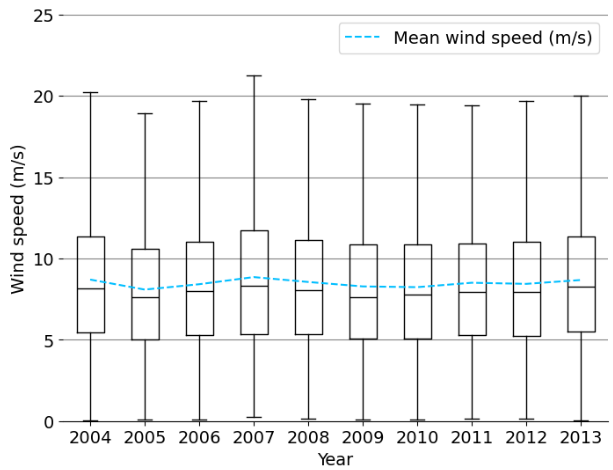

Figure 7: Wind speed distributions from 2004 to 2013 for synthetic data. Distributions of wind speeds in meters per second (m/s) in the NOW-23 synthetic dataset from 2004 to 2013 were plotted by year with outliers not shown (n = 8760). Mean wind speeds by year are shown by the dotted blue line.

available over a recent time period for a given site, so it may need to be regenerated using new data that reflect recent wind patterns.

Creating more accurate wind speed prediction models can help mitigate the inefficiencies due to predictable variances in wind energy production. Our work contributes to the trend of utilizing deep learning for wind speed prediction, demonstrating that LSTM networks can achieve higher accuracy in day-ahead prediction when using synthetic data. Furthermore, sites viable for wind energy production that lack historical observational data can utilize wind speed prediction models trained on synthetic data, which can be further finetuned as observational data becomes available. Future wind energy infrastructure in the Great Lakes region will benefit from the greater availability of accurate wind prediction models, encouraging further development.

MATERIALS AND METHODS

Data

We used synthetic data from the National Renewable Energy Laboratory's NOW-23 Great Lakes dataset simulated at an elevation of 80 meters, which was generated using the Weather Research & Forecasting (WRF) program and validated with light detection and ranging (LiDAR) data from Lake Michigan (15). We used observational data from a buoy approximately 10 kilometers from the eastern shoreline of Lake Michigan in the 2013 Lake Michigan Wind Assessment near Muskegon, Michigan, at an elevation of 75 meters (12). We utilized both datasets at a temporal resolution of one hour. We used a set of about 750 observational samples from December 2013 as the testing dataset (12). The date column of the observational data was normalized to match that of the synthetic data, and the wind speed and direction columns were cleaned by replacing missing data values with the last known value from the column (16). Missing data made up approximately 5% of the data in the testing set and outliers made up about 1%.

The use of weather variables other than wind features in wind speed forecasting is generally not associated with improvements in prediction accuracy (7). Hence, to train our LSTM networks, we selected only wind speed in meters per second and wind direction in degrees as features. These observations were taken at either 75 or 80 meters, reflecting the height of most turbines in the United States, which was about 88 meters in 2018 (14).

We retrieved synthetic wind data from 2000 to 2013 for the NOW-23 Great Lakes dataset from the National Renewable Energy Laboratory developer network API and concatenated for the Muskegon site using a script available on GitHub (16). To create the NOW-23 Great Lakes dataset, researchers ran 16 WRF setups over a one-year period and selected the best-performing setup combination in predicting wind speeds when validated using LiDAR data (15). The WRF setup provides the boundary conditions and mathematical models used in generating the dataset (15). The dataset was created by concatenating multiple one-month segments of WRF simulations, which ran with a startup period beginning two days before the start of each month (15).

We retrieved observational wind data from April 2013 to December 2013 for the Lake Michigan Wind Assessment from the Atmosphere to Electrons website (12). Synthetic and observational wind data were then split into training and testing groups and normalized using min-max normalization. We trained LSTM neural networks on either synthetic or observational data to predict day-ahead wind speeds at the same location near the coast of Muskegon, Michigan (**Figure 1**).

Layer	Output Shape	Number of Parameters
LSTM	(None, 16)	1,280
Dropout	(None, 16)	0
Dense	(None, 8)	136
Dense	(None, 1)	9

Table 1: Keras long-short term memory (LSTM) network architecture. The architecture of the LSTM network used in this study as given by Keras. The parameters are the totals of the weights and biases associated with each layer. An output shape of (None, 16) indicates that one or more lists of length 16 are passed as output from a layer.

Python setup

We used TensorFlow, a machine learning library, and Keras, a deep learning library, run on Jupyter Notebook in Python 3.11 to train the networks used in this study (17-18). Each network used for comparison was trained for 50 epochs, cycling through the entire training dataset 50 times, with a batch size of 8, which is how many samples from the dataset were taken before updating the network's weights in each cycle. Each network was also composed of the same architecture, which utilized an LSTM layer, two densely connected layers, and a dropout layer to reduce overfitting to training data by randomly dropping model weights (**Table 1**). Multi-layer neural networks with dropout layers show better convergence when using the Adam optimizer, so it was chosen as the optimizer for training (19). Each network layer containing weights was regularized using L2 regularization to reduce overfitting further by penalizing any excessively large weights in the neural networks. In all, each network contained about 1,400 total parameters (**Table 1**). Each network's MAPE accuracy was calculated using functions from the Scikit-learn library (20).

Parameter optimization

To minimize the MAPE of networks trained on synthetic data, we experimented with the number of epochs, years of data, and batch size used in training the networks on the Muskegon site. In three separate experiments, we tested years of data spanning from 1 year to 14 years in 1-year intervals, epochs spanning from 50 to 500 in 50 epoch intervals, and batch sizes spanning from 8 to 128 in intervals of 8. We altered one parameter while holding the other two constant. For the parameters held constant, 14 years of data were used, 50 epochs, and a batch size of 128. To evaluate the effect of altering these variables, we considered a network's MAPE score when tested on the month of observational data (December 2013) that had been withheld from training.

We used the results of these experiments to inform the parameters of the networks trained on synthetic data we tested, which utilized ten years of data spanning from 2004 to 2013. The distributions of wind speeds and mean wind speeds in the synthetic dataset were consistent throughout this timeframe (**Figure 7**). We also compared networks trained on synthetic and observational data with various training parameters, which led us to train the synthetic and observational networks we compared with 50 epochs and a batch size of 8. We trained these LSTM networks on synthetic wind data 30 times and observational wind data 30 times to compile distributions of the MAPE of each type of network.

Statistical tests

We used Welch's t-test to compare the distributions of MAPE scores for networks trained on synthetic and observational data. A z-test for skewness showed the score distribution of observational networks was significantly different from a normal distribution. Because of this, we used a bootstrap t-distribution to interpret the result of the t-test.

ACKNOWLEDGMENTS

We greatly appreciate Renaissance Learning, which provided valuable mentorship and insight from multiple professionals throughout the research process.

Received: July 29, 2024
Accepted: June 18, 2025
Published: December 21, 2025

REFERENCES

1. Laurie, Carol. "Exploring Offshore Wind Energy Opportunities in the Great Lakes." *National Renewable Energy Laboratory*, 26 Jun. 2023, https://www.nrel.gov/news/program/2023/exploring-offshore-wind-energy-opportunities-in-the-great-lakes.html. Accessed 2 Oct. 2024.
2. Musial, Walter, et al. *Great Lakes Wind Energy Challenges and Opportunities Assessment*. National Renewable Energy Laboratory, 31 Mar. 2023, https://doi.org/10.2172/1968585.
3. "Great Lakes Wind Energy Feasibility Study." *New York State Energy Research and Development Authority*, https://www.nyserda.ny.gov/All-Programs/Clean-Energy-Standard/Clean-Energy-Standard-Resources/Great-Lakes-Wind-Feasibility-Study. Accessed 2 Oct. 2024.
4. Konda, Srikanth Reddy, et al. "Dynamic Energy Balancing Cost Model for Day Ahead Markets With Uncertain Wind Energy and Generation Contingency Under Demand Response." *IEEE Transactions on Industry Applications*, vol. 54, no. 5, pp. 4908-16, 6 June 2018, https://doi.org/10.1109/TIA.2018.2844363.
5. Perr-Sauer, Jordan, et al. "Short-term wind forecasting using statistical models with a fully observable wind flow." *Journal of Physics: Conference Series*, vol. 1452, 2020, https://doi.org/10.1088/1742-6596/1452/1/012083.
6. Turner, David, et al. "Evaluating the Economic Impacts of Improvements to the High-Resolution Rapid Refresh (HRRR) Numerical Weather Prediction Model." *Bulletin of the American Meteorological Society*, vol. 103, no. 2, 12 Feb. 2022, pp. E198-E211, https://doi.org/10.1175/BAMS-D-20-0099.1.
7. Alves, Décio, et al. "The Potential of Machine Learning for Wind Speed and Direction Short-Term Forecasting: A Systematic Review." *Computers*, vol. 12, no. 10, 13 Oct. 2023, https://doi.org/10.3390/computers12100206.
8. Lindemann, Benjamin, et al. "A survey on long short-term memory networks for time series prediction." *Procedia CIRP*, vol. 99, 2021, pp. 650-55, https://doi.org/10.1016/j.procir.2021.03.088.
9. Neshat, Mehdi, et al. "A deep learning-based evolutionary model for short-term wind speed forecasting: A case study of the Lillgrund offshore wind farm." *Energy Conversion and Management*, vol. 236, 2021, https://doi.org/10.1016/j.enconman.2021.114002.
10. Araya, Ignacio, et al. "A Multi-Scale Model based on the Long Short-Term Memory for day ahead hourly wind speed forecasting." *Pattern Recognition Letters*, vol. 136, 2019, pp. 333–40, https://doi.org/10.1016/j.patrec.2019.10.011.
11. Bodini, Nicola, et al. *2023 National Offshore Wind data set (NOW-23)*. National Renewable Energy Laboratory, Jan. 1 2020, United States, https://doi.org/10.25984/1821404.
12. Standridge, Charles. *Lake Michigan Buoy near Muskegon, MI / Raw Data*. 2023, United States, https://doi.org/10.21947/1877893.
13. Ookura, Soichiro, and Hiroyuki Mori. "An Efficient Method for Wind Power Generation Forecasting by LSTM in Consideration of Overfitting Prevention" *International Federation of Automatic Control*, vol. 53, no. 2, 14

https://doi.org/10.59720/24-307

Apr. 2021, pp. 12169-74, https://doi.org/10.1016/j.ifacol.2020.12.1008.
14. Lantz, Eric, et al. "Increasing Wind Turbine Tower Heights: Opportunities and Challenges." National Renewable Energy Laboratory, May 2019, United States, https://doi.org/10.2172/1515397.
15. Bodini, Nicola, et al. "The 2023 National Offshore Wind Data Set (NOW-23)." *Earth System Science Data*, 26 Apr. 2024, https://doi.org/10.5194/essd-2023-490.
16. Wycoff, Alex, "Wind Speed Prediction with Synthetic Data." GitHub. https://www.github.com/alexwycoff/windspeed. Accessed 2 Oct. 2024.
17. TensorFlow Developers. *Tensorflow*. v2.20.0, *Zenodo*, 13 Aug. 2025, https://doi.org/10.5281/zenodo.16852354.
18. Chollet, François, et al. *Keras*. v3.11.3, 2015, https://keras.io.
19. Kingma, Diederik, and Jimmy Ba. "Adam: A Method for Stochastic Optimization." *arXiv*, 2014, arXiv:1412.6980v9, https://doi.org/10.48550/arXiv.1412.6980.
20. Pedregosa, Fabian, et al. "Scikit-learn: Machine learning in Python." *Journal of Machine Learning Research*, vol. 12, 2011, pp. 2825-2830.