

Using LSTM To Perform Load Predictions For Grid-Interactive Buildings

Kyppy N. Simani, Yuval O. Genga, and Yu-Chieh J. Yen

Abstract—Energy consumption from the residential sector forms a large portion of the electricity grid demand. The growing accessibility of residential load profile data presents an opportunity for improved residential load forecasting and the implementation of demand-side management (DSM) strategies. Machine learning is a tool well-suited for predicting stochastic processes, such as residential power usage due to human behavior. Long short-term memory (LSTM) recurrent neural networks are especially suited for predicting time-series data such as electrical load profiles. This paper investigates the impact of LSTM hyperparameters to the predictive performance of models, which include the tradeoffs associated with training data size, horizon ratios, model fidelity, prediction horizon and computational intensity. This paper provides a framework to evaluate the choice of LSTM hyperparameters for understanding trade-offs in a practical application of load profile predictions for the context of Grid-interactive Efficient Buildings (GEBs).

Index Terms—Grid-Interactive Buildings, LSTM, Machine Learning, Load Forecasting, Demand-Side Management.

I. INTRODUCTION

LOAD forecasting is an essential component to access demand flexibility with grid-interactive efficient buildings (GEB). Fully automated GEBs are envisioned to balance the load requirements for the building and demand response signals from the grid [1]. GEBs with on-site generation and storage could manage interactions with the grid for buying or selling credits through energy markets. Effective load demand forecasting of the building would enable automated decision-making for a GEB controller to manage energy, optimize costs and ensure user comfort [2]. Fig. 1 demonstrates the high level interactions between GEBs and the grid.

The development of a high performance machine learning model traditionally involves an iterative, trial-and-error process that involves the selection and tuning of numerous hyperparameter values to train and test several prediction models. This can be a time-, data- and compute-intensive procedure. In a practical setting such as a GEB home, the electrical energy required to complete the model development could incur additional costs to the end-user. Furthermore, the

All authors with the School of Electrical & Information Engineering at University of the Witwatersrand, Johannesburg, South Africa. Corresponding author is Kyppy N. Simani (email: kyppy.simani@students.wits.ac.za), Yuval O. Genga (email: yuval.genga@wits.ac.za), Yu-Chieh J. Yen (email: yu-chieh.yen@wits.ac.za).

Based on “Using LSTM To Perform Load Modelling For Residential Demand Side Management”, by K.N Simani, Y.O. Genga, and Y-C.J Yen which appeared in the Proceedings of Southern African Universities Power Engineering Conference (SAUPEC) 2023, Johannesburg, 24 to 26 January. © 2023 SAIEE

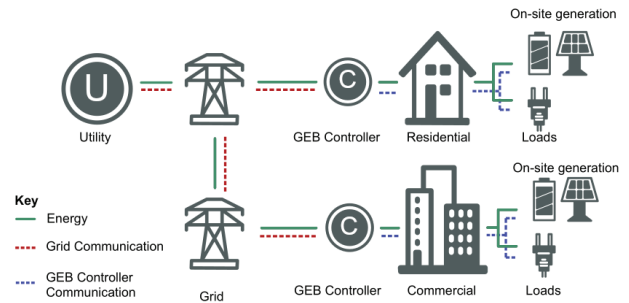


Fig. 1: Grid-interactive efficient building applications.

tuning of some hyperparameters, such as the load profile data resolution, would require the purchase of additional hardware.

The aim of this paper is to investigate the relationship and extent of LSTM hyperparameters for the envisioned application of load profile forecasting for GEBs. The evaluation of hyperparameters provided in this paper is envisioned to provide a framework towards informed design choices for LSTM application. The residential load profile uses real, measured household data [3]. A long short-term memory (LSTM) [4] machine learning algorithm is used to produce a series of prediction models in two experiments each varying specific hyperparameters values including: training data size, window size and prediction horizon length.

There exists literature exploring LSTM and various other neural network architectures to optimize load profile prediction accuracy and model performance [5]. However, there is a lack of research exploring the application of the relationships between various hyperparameters and model fidelity. For each experiment, LSTM prediction models are trained using various hyperparameters and fidelity performance is evaluated using Mean Absolute Error (MAE) and energy predictions.

II. BACKGROUND

The power grid landscape is in a state of transition worldwide [6]. An increase in power demand coupled with an expanding energy mix adds variability and stability issues on the grid, presenting an opportunity for optimization through energy management. There are several demand-side management (DSM) techniques used to manage grid demand by shifting load or peak curtailment. These include the use of responsive loads, ripple control, price incentives or Time-of-Use (TOU) tariffs [7]. In addition, end users have greater control over supply and local usage, through on-site generation such as rooftop photovoltaics (PV), electric vehicles (EVs) and usage of smart technologies to control loads.

A grid-interactive building is envisioned to respond to utility and user requirements through an autonomous control unit, as shown in Fig. 1. Its function could include the evaluation of decisions for the following:

- response to DSM signals,
- trading on energy markets or for TOU tariffs,
- monitoring battery state-of-charge for planned outages,
- evaluating vehicle-to-grid energy exchange,
- selling excess on-site generation such as PV energy, or
- activating flexible loads, such as water heaters, and air conditioning to draw from the grid in times of excess supply or in anticipation for an occupant's requirements.

Its purpose is to ensure that the building function can continue without impacting user comfort, possibly providing opportunity to enhance it. Therefore, a predictive model is required to forecast load usage in a building. Accurate daily demand forecasts for a building would enable the building controller to interact intelligently with both the utility and the building occupants.

A. Grid Interactive Buildings

Anticipating load demand and usage profile patterns is an important aspect of effective demand side management [8]. These forecasts can offer insights that affect planning for future electrical infrastructure, investment decisions as well as reduce the risk of supply shortages [3]. Electrical utilities have traditionally performed short to long-term demand forecasts on the distribution/generation level to inform infrastructure planning and DSM policy [8]. The shift towards grid-interactive buildings which enable the automation of power measurements and energy management, consumers can access their own consumption information. This enables consumers to effectively manage their energy usage and implement DSM strategies specific to their consumption requirements. Additionally, it has been shown that equipping consumers with more information about their power consumption improves the likelihood and duration of their engagement with their consumption habits [9].

To achieve load shifting and peak reductions utilities can encourage consumers to modify their behavior by implementing TOU tariffs. Power monitoring of grid-interactive buildings enables consumers to take advantage of TOU tariffs. Furthermore, consumers with renewable energy resources such as PV-battery storage can take advantage of energy trading in addition to TOU tariffs.

The ability to predict load demand and energy consumption can augment the advantages of grid-interactive buildings by enabling more informed decision making with regards to TOU, and for prosumers, forecast energy supply and demand for more effective energy trading [10].

B. Machine Learning for Load Prediction

Machine learning algorithms are an effective tool for analyzing residential load data as it is able to detect latent patterns and estimate future outcomes based on large, historical data sets [4], [11]. LSTM neural networks have been shown to




Fig. 2: Training and validation data split for 100 days of data

be particularly suited for modeling time-series data as it can identify and retain pattern features for long data sets [4]. They have also been shown to outperform the prediction accuracy of traditional load forecasting techniques such as auto-regressive integrated moving average (ARIMA) models, as well as other neural network architectures for predictions on time-series data [11]. Additionally, LSTMs have been shown to demonstrate high forecast accuracy for time-series predictions using both single variate and multivariate input features [12].

Research has been done to explore means of optimizing LSTM parameter selection to improve model performance [13]. However, there is a lack of research exploring the relationships between the depth of a forecast and its impact on the the training feature requirements to maintain a high model fidelity.

III. METHODOLOGY

For this investigation an LSTM neural network is used to develop several load profile prediction models in two experimental scenarios. For each experiment the relationship between the model training duration, prediction accuracy and a specific hyperparameter is examined by training a series of models with varying values of each hyperparameter. Experiment 1 examines the impact of increasing training data. Experiment 2 examines the ratio of the look-back window size to the horizon length, which is the number of time steps ahead that the model will predict. The raw load profile dataset is sourced from active power values measured from a real household over a period of 4 years [15]. From this data four features are identified to form the LSTM training dataset: active power, time of day, day of the week and time of year. This training data is split into five training subsets, {100, 200, 300, 400, 500} where each member of the set is a number of days of load profile measurements. Each LSTM model is trained using a rolling forecast scheme of 12-hour long input windows to predict outputs of a given horizon length [14].

A. Training Data Preparation

A public data set of measurements from a real household is sourced and used to train the LSTM models. The data set consists of timestamped measurements for aggregate active power and sub-metered energy measured between December 2006 and November 2010 [15]. The data is measured at a rate of one sample per minute for a total of approximately 4 years of samples. Missing values are substituted by copying values from the same time in the previous day.

The raw residential load profile data is used to identify several features for the final training data set. The active power values are scaled using a min-max normalization to

TABLE I: Values assigned to each training subset and ratio.

Subset	Training		Ratios		
	Training (days)	Validation (days)	Ratio	Window (hr)	Horizon (hr)
S_1	100	25	0.05	12	0.6
S_2	200	50	0.10	12	1.2
S_3	300	75	0.30	12	4.0
S_4	400	100	0.50	12	6.0
S_5	500	125	0.75	12	9.0
			1.00	12	12.0

create the first feature [16]. The date and time values for the timestamps of each load profile measurement are expanded into two feature sets. The first set expresses each timestamp as a time-of-day signal by mapping the date-time values to sine and cosine functions with a period of 24-hours. The second set expresses each timestamp as a time-of-year signal by mapping the date-time values to sine and cosine functions with a period of 365.24 days. The resultant training data is comprised of a multivariate, time-series dataset composed of four high-level features: active power, the time of day, day of the week and the time of year. Each time-based feature is mapped to a unique pair of values using a sine-cosine transform function.

The selected date ranges for the training and validation subsets of each experiment are chosen relative to a single reference data point in the dataset. For each subset the training and validation data are split using a ratio of 80% training and 20% validation as demonstrated in Fig. 2 for a training-validation subset consisting of 100 days. For each experiment the test data is selected as the 24-hour period immediately following the chosen reference point of each experiment. This is to ensure that each of the models trained using different lengths of days would have a common test set so that their performances could be fairly compared.

B. LSTM Model Design

A stacked LSTM architecture composed of 3 sequential neuron layers was chosen for this work. The first layer consists of 100 neurons, while the second and third layers contain 50 neurons each. The outputs of the final LSTM layer are sent to a linear-output feed forward layer which maps the intermediate LSTM layer outputs to a single output value i.e., the power forecast of the target time interval. The hyperbolic tangent (tanh) is chosen as the activation function [17].

For the training parameters a batch size of 32 is used. The number of epochs is kept constant at three. It is found that additional epochs accrued significant increases to model training time for relatively little gain in performance. The ADAM optimizer with a learning rate of 0.0001 is used [18]. Model development and training is performed using a desktop PC with a 3.70 GHz AMD Ryzen 5 5600X 6-Core processor with 16GB RAM, and a GeForce RTX 3070Ti graphics card with 8GB GDDR6X VRAM, using the Keras library with a Tensorflow back end.

C. Experimental Setup

Two experiments are performed to investigate the impact of selected hyperparameter variations on LSTM model prediction performance. In Experiment 1 models are trained using

increasing amounts of training data. In Experiment 2 the ratio of window size to prediction horizon is varied.

1) *Experiment 1:* The purpose of this experiment is to investigate the relationship between training data size, computational resource and model fidelity. It is well-established that more training data improves the performance of machine-learning models. This experiment aims to observe the extent to which increasing training data size improves model prediction against computational intensity.

The training data increments start from 1 day and are increased by 90-day increments (approximately 3 months) up to a maximum of 1080 days (3 years), so that complete seasons (i.e., summer, fall winter, spring) are incremented. This is to reduce variability due to changes in seasons. After three epochs the model training is complete. Each model is evaluated by performing predictions on the 24-hour set of test data measurements.

2) *Experiment 2:* The purpose of this experiment is to investigate the relationship between training data size, horizon window and model fidelity. Depending on the application and frequency of exchange signals, a load forecast of 24 hours is considered ideal. The aim is to determine the extent to which training data size can improve the model accuracy for a range of prediction horizons. The prediction horizon length is determined as a fraction of the window size using:

$$\frac{h}{w} = r \quad (1)$$

where h is the prediction horizon length, w is the window length and r is the horizon-window ratio. The set of horizon-window ratios used for this experiment are {0.05, 0.1, 0.3, 0.5, 0.75, 1.0}.

For this experiment a set of five training data sizes is chosen {100, 200, 300, 400, 500}. For each set of training data, a model is produced using each combination of window size and prediction horizon length as presented in Table I. Finally, each model is evaluated by performing predictions on the 24-hour set of test data measurements.

D. Evaluation Metrics

The Mean Absolute Error (MAE) is used to evaluate the forecast accuracy of each LSTM prediction model [19]. MAE is simple to interpret as the result is in the same unit as the predicted data. It is also robust against outliers in the dataset. The MAE is a score that determines the average absolute error between the observed and model predicted values and is formulated as:

$$\frac{\sum_{i=1}^n |y_i - x_i|}{n} \quad (2)$$

where, y_i is a predicted value, x_i is the observed value and n is the total number of data points. An MAE value of 0 represents a perfect prediction, therefore values closer to 0 indicate higher model fidelity and better model performance.

To quantify the energy prediction performance a residual energy (RE) value is determined. This is calculated by taking

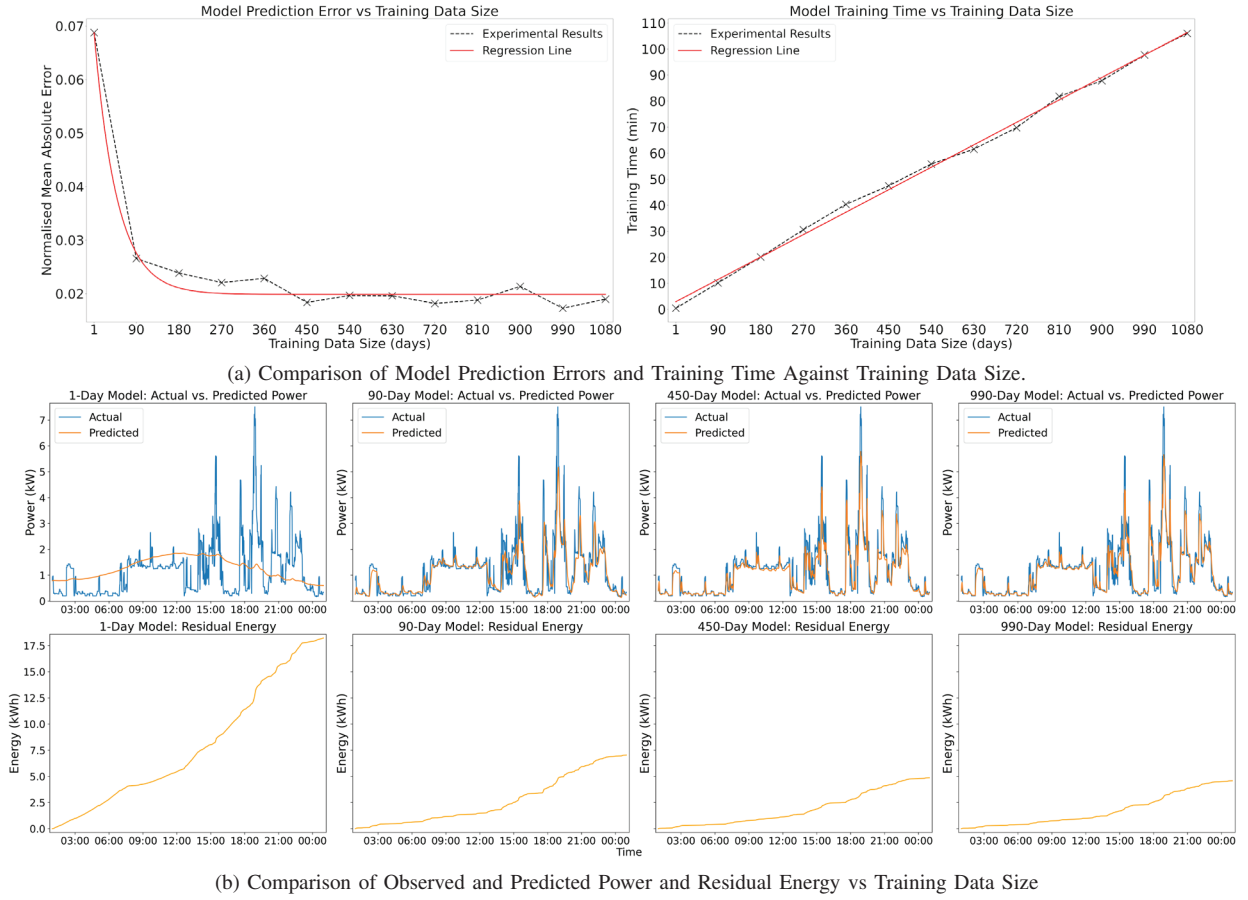


Fig. 3: Results for training data against MAE and residual energy (Experiment 1).

the numerical integral of the absolute difference between the power-time graphs of the observed and predicted values:

$$RE = \sum_{n=0}^{N-1} \frac{1}{2} (|c[n] - d[n]| + |c[n+1] - d[n+1]|) (\Delta x)_n \quad (3)$$

where N is the total number of samples, $c[n]$ is the observed power and $d[n]$ is the predicted power at sample $[n]$, and Δx is the time step size at sample n .

A RE value of 0 indicates an exact energy prediction. RE values greater than 0 indicate larger amounts of erroneously predicted energy and poorer prediction model performance.

IV. RESULTS

The results for each experiment are discussed in the sections to follow.

A. Experiment 1

A comparison of the MAE against training data size for Experiment 1 is shown in Fig 3 (a) and (b). The results show an exponentially decaying relationship between the MAE and training data size. A regression analysis derives the following equation:

$$MAE = 0.050e^{-0.021t} + 0.020 \quad (4)$$

where t is the training data size expressed in days.

A comparison of the model training time against training data size is shown in Fig 3 (a). It can be observed that for smaller training data sizes the training duration is also small. As the training data size increases the model training duration also increases at an approximately linear rate.

Regression analysis of this plot confirms a linear relationship between the training duration and training data size. The equation for the linear model is:

$$T = 0.096t + 2.804 \quad (5)$$

where T is the training duration size expressed in minutes and t is the training data size expressed in days.

A choice of training data size needs to be a balanced trade-off between model fidelity (MAE) and computational intensity (training duration).

A comparison of the predicted load profiles and cumulative predicted energy error for the 1, 90, 450 and 990-day prediction models are shown in Fig 3 (b). As the training data size increases there is a reduction in the amount of energy that is incorrectly predicted.

B. Experiment 2

A sample of the predicted load profiles for the range of training days for a chosen horizon ratio is presented in Fig. 4. A comparison of the MAE values for the Experiment 2

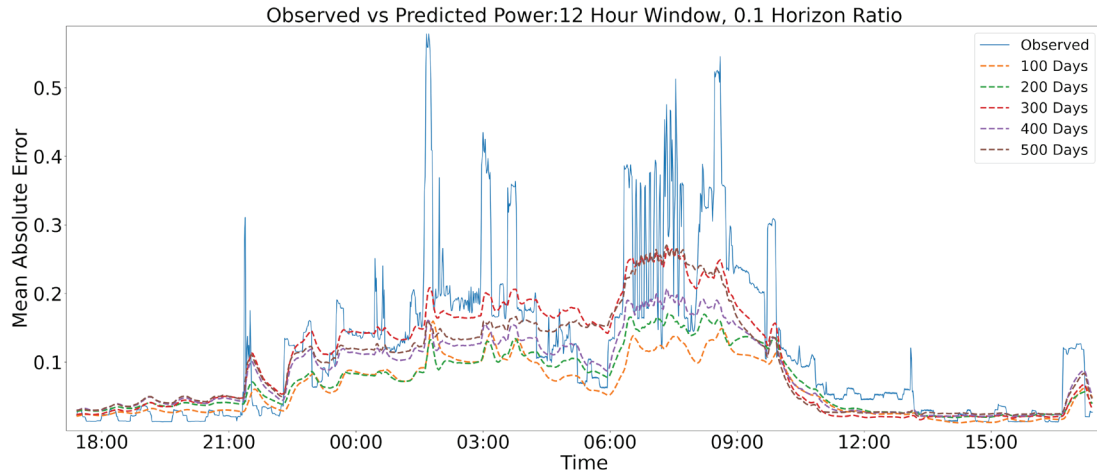


Fig. 4: Comparison of observed vs predicted power profiles for a 0.1 ratio prediction model (Experiment 2).

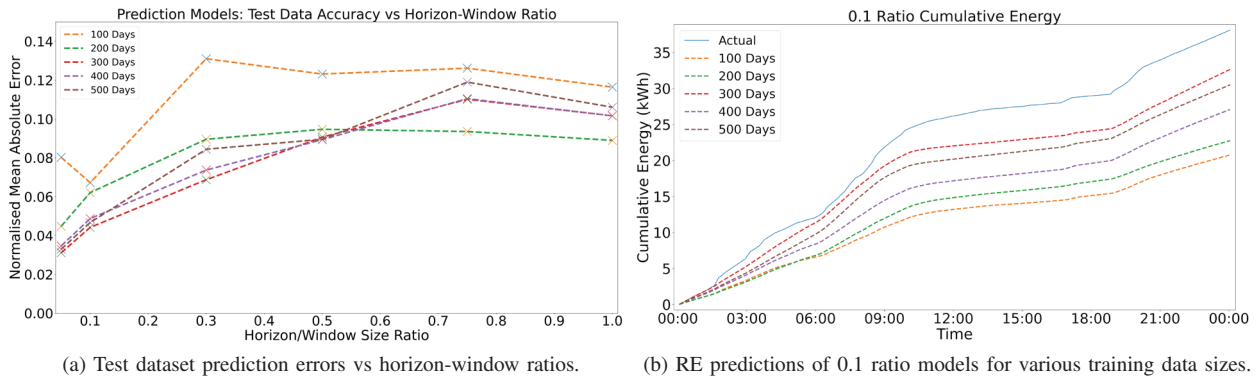


Fig. 5: Prediction errors and residual energy (Experiment 2).

model predictions on the test data set is shown in Fig. 5 (a). As the prediction horizon increases the predictions tend to become more conservative and consequently underestimate their values. All models struggled to predict the larger peaks. This trend worsens as the prediction horizon increases.

For the ratio of 0.1 (relating to a 1.2 hr horizon from Table I), the models demonstrate relatively high prediction accuracy. The highest performing model is the 300 day training data. This suggests that there is an optimal training data size for horizon window. Training days for more than 300 days, approximating a year, suggest diminishing returns.

In Fig. 5 (b), a cumulative energy plot using a 0.1 ratio model is shown. Models of each training data size show fair agreement for the first six hours, and produce an under-prediction for energy as time is extended.

V. DISCUSSION AND FUTURE WORK

From Experiment 1, an exponential decay relationship is observed for model fidelity (MAE) against training data size. The relationship between computational intensity (training time) and training data is shown to be linear. After 90 days (approximating a season), it is shown that there is a marginal increase in model prediction fidelity, while computational intensity increases. This results in marginal gains for the cost of higher computational resources as training data is increased.

For a practical application in a GEB, an energy management system set up with a 90-day training model is a good starting point with an MAE of 0.37kW and an energy prediction error of 28%. More training days provide marginal gains at higher computational cost. If retraining every day, larger training day sizes become infeasible.

From Experiment 2, a horizon window ratio of 0.1, relating to a 1.2 hr horizon, shows good agreement. In addition, the 300 day model demonstrates the best performance, suggesting that more training days may result in diminishing returns.

For a practical application in a GEB, an energy management system set up with a 1.2 hr window can produce reasonable predictions. This is for data sampled per minute. The training data for this performance only requires 300 days. In this case, more training data does not result in better performance.


Improvements on horizon length could be investigated with reductions in sampling resolution. Per minute samples may limit the horizon due to the number of prediction steps required to reach an hour (or more). In addition, improvements to the LSTM design could produce better performance. This study is not focused on the LSTM design, but rather the hyperparameter impact on performance.

VI. CONCLUSION

Effective demand-side management using autonomous on-site control requires load demand prediction at a suitable accuracy and horizon. Grid-interactive efficient buildings have the potential to enable consumers to optimize around DSM strategies such as TOU tariffs as well as engage in energy trading. Load forecasting is an integral part of any automated decision-making. The results of this investigation show that hyperparameter choices need to be made to balance the trade-offs of a practical implementation of load forecasting. For this LSTM design and this dataset, the residence could start with 90 training days with one-step horizon and increase performance to 300 training days for 1.2 hrs horizon. This work presents a framework for analyzing LSTM hyperparameters for prediction performance in preparation for a practical implementation towards autonomous control.

REFERENCES

- [1] Xue, X., Wang, S., Sun, Y., & Xiao, F. (2014). An interactive building power demand management strategy for facilitating smart grid optimization. *Applied Energy*, 116, 297–310. <https://doi.org/10.1016/j.apenergy.2013.11.064>
- [2] Pinto, G., Kathirgamanathan, A., Mangina, E., Finn, D. P., & Capozzoli, A. (2022). Enhancing energy management in grid-interactive buildings: A comparison among cooperative and coordinated architectures. *Applied Energy*, 310, 118497. <https://doi.org/10.1016/j.apenergy.2021.118497>
- [3] Sabapathi, D., Pawar, Y. S., Patnaik, S., Sivanantham, E., Prabhu, D. K., & Prakash, N. B. (2023). Management in Industrial Sectors using Neuro-Fuzzy Controller and Deep Learning. *Proceedings of the 3rd International Conference on Artificial Intelligence and Smart Energy, ICAIS 2023*, 1432–1437. <https://doi.org/10.1109/ICAIS56108.2023.10073714>
- [4] Lipton, Z. C., Berkowitz, J., & Elkan, C. (2015). A Critical Review of Recurrent Neural Networks for Sequence Learning. <https://arxiv.org/abs/1506.00019v4>
- [5] Luo, J., Zhang, Z., Fu, Y., & Rao, F. (2021). Time series prediction of COVID-19 transmission in America using LSTM and XGBoost algorithms. *Results in Physics*, 27. <https://doi.org/10.1016/j.rinp.2021.104462>
- [6] Fang, X., Misra, S., Xue, G., & Yang, D. (2012). Smart grid - The new and improved power grid: A survey. *IEEE Communications Surveys and Tutorials*, 14(4), 944–980. <https://doi.org/10.1109/SURV.2011.101911.00087>
- [7] Palensky, P., & Dietrich, D. (2011). Demand side management: Demand response, intelligent energy systems, and smart loads. *IEEE Transactions on Industrial Informatics*, 7(3), 381–388. <https://doi.org/10.1109/TII.2011.2158841>
- [8] Zhao, H., Tang, Z. (2016). The review of demand side management and load forecasting in smart grid. *Proceedings of the World Congress on Intelligent Control and Automation (WCICA)*, 2016-September, 625–629. <https://doi.org/10.1109/WCICA.2016.7578513>
- [9] Hargreaves, T., Nye, M., & Burgess, J. (2013). Keeping energy visible? Exploring how householders interact with feedback from smart energy monitors in the longer term. *Energy Policy*, 52, 126–134. <https://doi.org/10.1016/j.enpol.2012.03.027>
- [10] Radovanovic, A., Nesti, T., & Chen, B. (2019). A Holistic Approach to Forecasting Wholesale Energy Market Prices. *IEEE Transactions on Power Systems*, 34(6), 4317–4328. <https://doi.org/10.1109/TPWRS.2019.2921611>
- [11] Siami-Namini, S., Tavakoli, N., & Siami Namin, A. (2019). A Comparison of ARIMA and LSTM in Forecasting Time Series. *Proceedings - 17th IEEE International Conference on Machine Learning and Applications, ICMLA 2018*, 1394–1401.
- [12] Karim, F., Majumdar, S., Darabi, H., & Harford, S. (2019). Multivariate LSTM-FCNs for time series classification. *Neural Networks*, 116, 237–245. <https://doi.org/10.1016/j.neunet.2019.04.014>
- [13] Bouktif, S., Fiaz, A., Ouni, A., & Serhani, M. A. (2018). Optimal deep learning LSTM model for electric load forecasting using feature selection and genetic algorithm: Comparison with machine learning approaches. *Energies*, 11(7).
- [14] Yuan, S., Wang, C., Mu, B., Zhou, F., & Duan, W. (2021). Typhoon Intensity Forecasting Based on LSTM Using the Rolling Forecast Method. *Algorithms* 2021, Vol. 14, Page 83, 14(3), 83. <https://doi.org/10.3390/A14030083>
- [15] Individual household electric power consumption - UCI Machine Learning Repository (n.d.). Retrieved July 30, 2023, from <https://archive.ics.uci.edu/dataset/235/individual+household+electric+power+consumption>
- [16] Grierson, S., Thomson, C., Papadopoulos, P., & Buchanan, B. (2021). Min-max Training: Adversarially Robust Learning Models for Network Intrusion Detection Systems. *Proceedings - 2021 14th International Conference on Security of Information and Networks, SIN 2021*. <https://doi.org/10.1109/SIN54109.2021.9699157>
- [17] Farzad, A., Mashayekhi, H., & Hassanpour, H. (2019). A comparative performance analysis of different activation functions in LSTM networks for classification. *Neural Computing and Applications*, 31(7), 2507–2521. <https://doi.org/10.1007/S00521-017-3210-6/TABLES/11>
- [18] Kingma, D. P., & Ba, J. L. (2014). Adam: A Method for Stochastic Optimization. *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*. <https://arxiv.org/abs/1412.6980v9>
- [19] Willmott, C. J., & Matsuura, K. (2005). Advantages of the mean absolute error (MAE) over the root mean square error (RMSE) in assessing average model performance. *Climate Research*, 30(1), 79–82. <https://doi.org/10.3354/CR030079>



Kypny N. Simani was born in Nairobi, Kenya in 1993. He received a BEng. in electronic engineering from the University of Pretoria, South Africa in 2018. He is currently an MSc. candidate at the University of the Witwatersrand, South Africa. He holds a research fellowship at Brown University.

He has worked as a software developer specializing in digital communications services. Since 2022, he has been conducting research on the application of machine learning to optimize energy trading using grid-tied, residential PV systems. His research interests include energy trading, energy management, and machine learning.



Yuval O. Genga received a B.Sc degrees in electrical and information engineering from the University of Nairobi, Nairobi, Kenya in 2012. He holds an M.Sc and Ph.D. degree from the University of the Witwatersrand, South Africa. He is currently a lecturer at the University of the Witwatersrand. His research interests revolve around the practical applications of Machine Learning (ML) and Artificial Intelligence (AI) with a specific emphasis on optimizing and enhancing traditional approaches.

Within this domain, he is actively engaged in multiple projects that leverage Deep Learning techniques, particularly in computer vision, time-series data prediction, and telecommunication systems.



Yu-Chieh J. Yen was born in Taipei, Taiwan in 1986. She received her BSc, MSc and PhD degrees in electrical engineering from the University of Witwatersrand, South Africa. She is currently an academic in the School of Electrical & Information Engineering at the University of the Witwatersrand. She is a research fellow of Oxford University.

She specializes in modeling of energy systems for various contexts including demand-side management and grid-interactive buildings, with a particular focus on electric water heaters. Her interests lie in sustainable development to enable access to energy and STEM education. She has informed on national policy related to DSM and water heating regulation.