



Review

Deep-learning-based multi-timescale load forecasting in buildings: opportunities and challenges from research to deployment

Sakshi Mishra*, Stephen M. Frank, Anya Petersen, Robert Buechler, Michelle Slovensky

Intelligent Campus Program, National Renewable Energy Laboratory, Golden, Colorado, United States

ARTICLE INFO

Article history:

Received 12 July 2024

Received in revised form

18 August 2024

Accepted 26 August 2024

Keywords:

Deep learning, Recurrent neural networks, Long short-term memory, Building load forecasting, Grid-interactive efficient buildings, Smart grid

*Corresponding author

Email address:

Sakshi.M@outlook.com

DOI: 10.55670/fpll.fuen.3.4.3

ABSTRACT

Electricity load forecasting for buildings and campuses is becoming increasingly important as the penetration of distributed energy resources (DERs) grows. Efficient operation and dispatch of DERs require reasonably accurate predictions of future energy consumption in order to conduct near-real-time optimized dispatch of on-site generation and storage assets. Electric utilities have traditionally performed load forecasting for load pockets spanning large geographic areas, and therefore, forecasting has not been a common practice by buildings and campus operators. Given the growing trends of research and prototyping in the grid-interactive efficient buildings domain, characteristics beyond simple algorithm forecast accuracy are important in determining the algorithm's true utility for smart buildings. Other characteristics include the overall design of the deployed architecture and the operational efficiency of the forecasting system. In this work, we present a deep-learning-based load forecasting system that predicts the building load at 1-hour intervals for 18 hours in the future. We also discuss challenges associated with the real-time deployment of such systems as well as the research opportunities presented by a fully functional forecasting system that has been developed within the National Renewable Energy Laboratory's Intelligent Campus program.

1. Introduction

The commercial buildings sector in the United States consumed 1,671.61 trillion Btu of energy in January 2020 alone (immediately prior to the onset of the COVID-19 pandemic) and a total of 18,177.95 trillion Btu in 2019 [1]. Overall, buildings account for nearly 40% of the total energy consumption in the United States [2], and building energy consumption is projected to expand by an annual 1.5% globally [3]. Even with many commercial facilities operating at reduced occupancy during the COVID-19 pandemic, preliminary industry reports indicate that commercial building energy is again increasing as operators increase ventilation rates [4]. Because buildings are among the largest consumers of energy globally, research to increase the design and operational energy efficiencies of the commercial sector (i.e., commercial buildings) will play a significant role in meeting energy and greenhouse gas emissions reduction targets. The penetration of behind-the-meter distributed energy resources (DERs) is also increasing, owing to the decreasing cost of renewable energy technologies [5]. Smart building operations can simultaneously reduce energy

consumption (and, by extension, greenhouse gas emissions) and optimize behind-the-meter DER dispatch to save money for building owners/operators via additional value streams such as demand management (peak shaving), energy arbitrage, user-initiated demand-response, and optimal electric vehicle charging. Properly deployed, these capabilities also benefit the utility grid by enhancing grid reliability and resilience, deferring or reducing capital expenditures required to upgrade the distribution grid, and helping balance the supply of renewable energy as its penetration increases. The U.S. Department of Energy's Grid-interactive Efficient Buildings initiative [6] describes the multi-faceted benefits of these smart building capabilities. Smart building controls are also an integral component of autonomous energy grids [7], where they provide opportunities for granular device-level controls. Proportional-integral-derivative (PID) controllers are still widely used in building controls for heating, ventilating, and air-conditioning (HVAC) systems, largely because of their simplicity and fast solution times, as they employ numerical methods to determine the controlling parameters [8].

However, the time-varying system dynamics characteristics of HVAC control systems result in the inconsistent performance of traditional controllers, often leading to less-economic performance. Moreover, the objectives of building control systems have shifted from being solely occupant-comfort and economic operations to being grid-interactive prosumers. As prosumers, building energy management systems have goals such as (i) maximizing the use of on-site distributed energy resources, (ii) managing consumption profiles based on time-of-use energy rates, and (iii) using on-site energy storage to reduce the peak demand for cost savings. Such a complex set of objectives requires the use of advanced control techniques like model-predictive control [9]. Implementation of intelligent dispatch and model-predictive control algorithms that help achieve the above-stated goals to reduce overall energy consumption and operational costs of the building nearly always requires a forecast of the energy consumption of the building. These algorithms use these building load predictions to schedule the dispatch of flexible loads, on-site clean energy resources, and/or energy storage systems to optimize a desired metric, such as utility cost. Forecast accuracy is a key driver of model-predictive control algorithm effectiveness [8-10]. Prediction of building energy consumption patterns is also important for detecting faults or operational anomalies in energy systems. Therefore, building load forecasting systems are an essential component of smart buildings.

Building load forecasting is not a new research question; much research work in the literature has addressed this problem space [11-13]. Commercial load forecasting tools for building energy predictions are also available [14,15]. The journey of successful research outcome to industry deployment, however, inevitably passes through the research-prototyping phase, and advanced building load forecasting systems are no exception. This is especially relevant for data-driven modeling projects where the type, quality, and pre-processing of data largely dictate the performance of the forecasting system. The work presented in this article sheds light on the end-to-end pipeline of building load forecasting systems- in a campus-wide prototyping setup- using advanced data-driven methods (recurrent neural networks). The following paragraphs delineate the types of modeling algorithms used in the literature for the building energy forecasting problem and specifically discuss works that have utilized data-driven modeling methods.

As building systems become more advanced through next-generation sensors, controls, connectivity, and communications, they produce a large volume of empirical data available to building operators for decision-making. This data, along with meteorological parameters, can be harnessed to predict building and campus energy consumption. There are two main approaches available for this prediction task: (i) physics-based (or “white box”) modeling and (ii) data-driven (or “black box”) modeling. (Physics-based models are termed “white box” because the inner workings of the model are typically open to the modeler, whereas data-driven models are termed “black box” because their prediction logic is often opaque.) A third category, “grey box” models [16], represents a combination of the physics-based and data-driven approaches. They combine a partial theoretical structure with

data, offering relatively simpler model architectures. Grey box models are generally trained or fit like data-driven models. Physics-based building models such as EnergyPlus® [17,18] model the physical relationships between the building characteristics (construction details, operation schedules, shading information) and environmental parameters (sky conditions) to calculate building energy consumption [19]. Data-driven models, on the other hand, make predictions by learning the pattern empirically from historical data. Data-driven models have two major subcategories: (i) statistical models and (ii) machine learning models. Deep learning models are a subset of machine learning models that are capable of learning complex nonlinear relationships between the inputs and the predicted variable(s). Recurrent neural networks (RNNs) are an advanced variant of deep neural networks that are capable of incorporating temporal dependencies between the input and output variables. Long short-term memory networks (LSTM) are a variant of RNNs that are effective at capturing longer-term temporal dependencies in the data sets.

There are numerous advanced applications of predictive analytics in the renewable energy field [20]- spanning the energy system from generation (solar and wind forecasting) to consumption (smart buildings energy forecasting to fault predictions). Because the focus of this article is smart buildings’ predictive analytics application, herein, we discuss literary works focused on data-driven building load forecasting. Recent research has established the utility of neural networks for forecasting the energy consumption of both individual buildings and groups of buildings. For example, a machine-learning-based forecasting model is presented in [21], whereas [22] covers a review of the intelligent system for power load forecasting. A Polish power system study with a deep learning approach is presented in [23], and [24] takes a deeper look at a hybrid system based on LSTM for short-term power load forecasting; similarly, reference [25] presents an RNN-based robust short-term load forecasting framework. An online adaptive RNN-based load forecasting algorithm with smart meter data is presented in [26]. Short-term load forecasting for urban loads using an artificial neural network is tackled in reference [27]. An LSTM-based method is improved for short-term load forecasting [28]. On the research end of the building load forecasting spectrum, the existing literature has extensively focused on assessing the effectiveness of machine learning (specifically deep learning) models for the building load prediction problem. The commercial end of building load forecasting spectrum, on the other hand, focuses on using proven and tested traditional methods in prepackaged software products. Both these extreme ends of the building load forecasting problem space do not address the “research prototyping” step for incrementally adopting cutting-edge prediction methodologies. That is to say, none of these works have covered a holistic picture of the deep-learning-based prediction system deployment process—in a research setup—which starts with the data collection and curation process and continues to critically assess the challenges and opportunities with its end-to-end deployment and continued use. The work presented in this paper contributes a perspective on how to move data-driven algorithms from research to practical deployment by presenting a case study

of designing, building, and deploying a research prototype of a deep-learning-based building load forecasting system. The current section lays out the introduction and background behind the presented work, and the rest of the paper is organized as follows: Section II provides a brief introduction of the real-time testbed (Intelligent Campus) utilized for deploying the deep learning models. Section III discusses the algorithm employed in modeling, methodology, and benchmarking. Section IV presents the challenges faced, lessons learned, and opportunities that a fully functional and deployed load forecasting system presents for modern buildings. Section V concludes the paper with a discussion of outlook. This paper is an extended version of our preprint [29].

2. Intelligent campus

The Intelligent Campus platform at the National Renewable Energy Laboratory (NREL) was established to collect historical and real-time building performance data to support analytics that enhance operational awareness and decision-making with respect to energy use. The platform was developed using open standards and protocols (Figure 1) with the intent to provide an architecture that is readily transferable to other campuses [30]. Since its inception, the Intelligent Campus platform has evolved into a program providing an ecosystem for the use of the NREL campus as a living laboratory. The Intelligent Campus program offers nascent energy analytics technologies and algorithms a pathway to mature toward commercialization. The Intelligent Campus program includes an interdisciplinary portfolio of projects aimed at prototyping the future of smart, sustainable, resilient, and self-healing buildings. By providing an environment for technology demonstration in real facilities, the Intelligent Campus program is bridging the gap between theoretical research into building efficiency technologies (which are often published but rarely deployed) and practical application. Intelligent Campus pilot projects include detailed monitoring and analysis in order to identify and improve technology shortcomings.

2.1 Data: measurement, collection, and curation

Advanced analytics and model-predictive control algorithms for smart buildings require reliable, internally consistent building performance data. NREL collects a variety of performance data at its South Table Mountain campus. The campus has a single common utility electric meter for all facilities, including the campus’s central heating and chilled water plant. Eighteen campus buildings have whole-building electrical meters; several other small buildings do not have dedicated electrical meters. Several newer buildings have submetering by end-use, per the requirements of ASHRAE 90.1 [31]. Facilities that consume heating and chilled water from the central plant have thermal meters for heating and chilled water consumption. NREL captures the electric and thermal meter interval data using a central energy management information system (EMIS). The EMIS also collects interval data from the campus building automation system and syncs data from several cloud data sources via web application programming interfaces, including utility data from the campus main electricity meter and NREL’s on-site research weather station [32]. To maximize accessibility for operations and research, Intelligent Campus organizes collected interval data per the Project Haystack standard [30,33]. These interval data are available to serve as inputs for predictive load models.

2.2 Predictive analytics

Intelligent Campus’s predictive analytics project applies recent machine learning and deep learning advances to develop and continuously improve load forecasting capability for NREL’s South Table Mountain campus. Intelligent Campus executes research, prototyping, and enhancement in a cycle so that practical deployment lessons can be folded back into early-stage research, increasing the ultimate effectiveness and eventual impact of the work. Intelligent Campus’s building load forecasting work builds from prior NREL research that predicts expected building performance based on exogenous inputs [34].

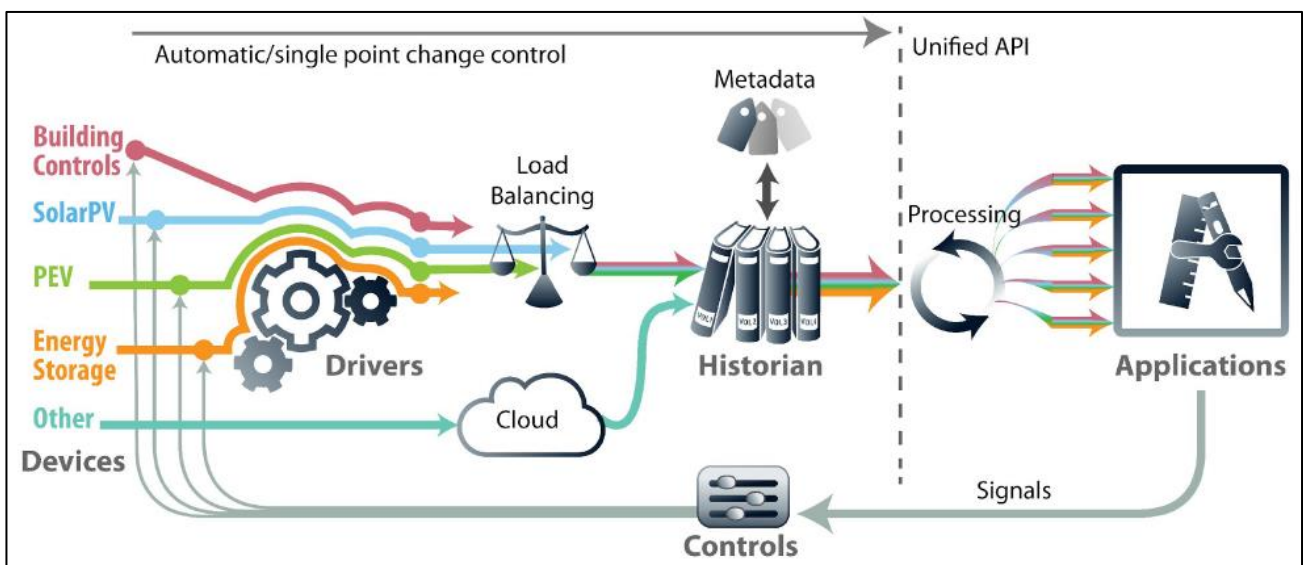


Figure 1. End-to-end data collection architecture (from [29])

As an academic discipline, load forecasting approaches can be studied using static postprocessed data sets. Transferal of these results to a real-world setting requires, at a minimum, some underlying sensor infrastructure, a mapping of sensor data streams to model inputs, and quality control on the measured data being fed to the model. Therefore, the focus of the current effort is twofold: (i) to automate, replicate, and scale the predictive algorithms and (ii) to improve accuracy by leveraging state-of-the-art neural networks.

3. Load forecasting using deep learning

3.1 Deep learning: long short-term memory network

Artificial neural networks are universal function approximators. They are capable of representing complex nonlinear relationships in high-dimensional data sets such as the one being employed for building load forecasting in this work. Deep learning algorithms, which employ multiple hidden layers, have powerful generalizing capabilities. That is, they are able to make reasonably accurate predictions for previously unseen scenarios because of their ability to learn the intricate structures in large data sets. Feedforward neural networks, sometimes referred to as “vanilla” deep neural networks, have the fundamental drawback of independence among the time-series samples/data points, which makes them an ineffective choice of algorithm for time-series-based prediction problems. This is because the entire state of the feedforward neural network is cleared after processing each time-series sample, which means that the network starts mapping the inputs and output for the next time step from scratch, thereby failing to account for the impact of the previous time step’s input variables. Recurrent neural networks (RNN) are advanced types of deep neural networks that overcome feed-forward neural networks’ limitations. They are different from feedforward neural networks because of the presence of additional directed edges that introduce temporal memory components, enabling them to capture complex nonlinear relationships between the temporally related inputs and outputs across multiple time steps. This makes them effective at modeling time-series forecasting problems. A long short-term memory (LSTM) network [35], shown in Figure 2, is an RNN variant that is able to learn long-term dependencies between the input features and the predicted variable [36].

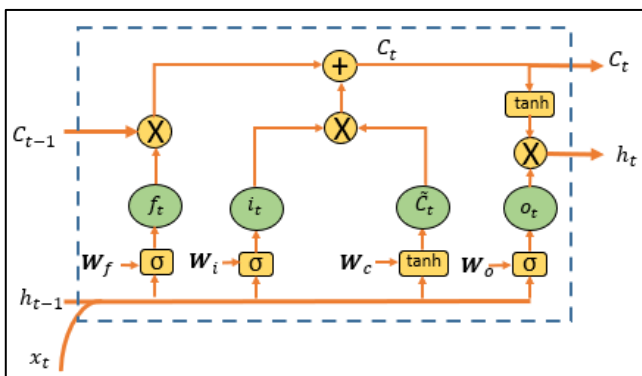


Figure 2. LSTM memory cell diagram (from [29,36])

Figure 2 displays a typical memory cell for an LSTM network. In the figure, C denotes the cell states at different points in time t ; f_t , i_t , and o_t denote the forward, input, and output gates, respectively; $h^{(t)}$ is the hidden state at time t ; and W_f , W_i , W_c , W_o are the weights. Further explanation can be found in [36].

3.2 Performance measure

The evaluation metric used to measure the performance of the forecasting algorithm is the mean squared error (MSE). The difference between the ground truth (i.e., the actual future time-step value of the measurement being predicted) and forecasting values (i.e., output of the deep learning model) is calculated using MSE:

$$MSE = \frac{1}{n} \sum_{j=1}^n (AV_j^{(t)} - \widehat{OV}_j^{(t)})^2 \tag{1}$$

where $AV_j^{(t)}$ is a vector of actual (ground truth) values and $\widehat{OV}_j^{(t)}$ is a vector of forecasted (output) values.

3.3 Case study: building load forecasting on a research campus

In this case study, the energy consumption at the main electricity meter for the Café building on the NREL South Table Mountain campus is predicted at an hourly resolution. There are six input features used for training the model: relative humidity, barometric pressure, dry bulb temperature, global horizontal irradiance, total cloud cover, and wind speed. These data are from NREL’s Solar Radiation Research Laboratory data set [32]; however, they also represent measurements commonly available from high-quality weather stations worldwide. The model is run for 200 epochs with the data spanning a year’s interval: 10- and 2-months train/test split, respectively. We use a single-layered network with 35 neurons. The following MSE plot (Figure 3) shows the train and test losses over several thousand iterations. Despite the difference between final train and test losses, no overfitting was observed in the experiment.

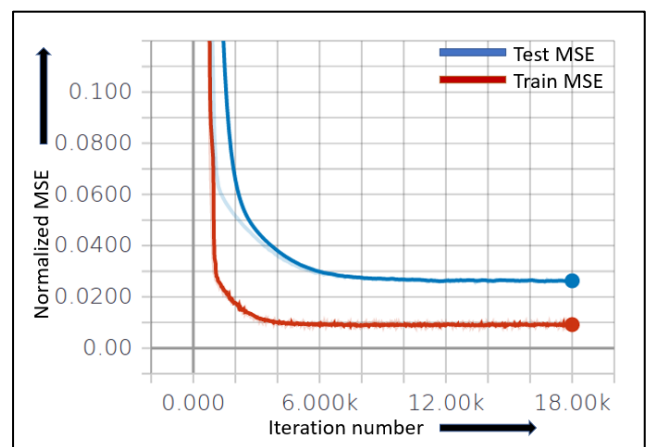


Figure 3. Mean squared error (MSE) error plot for train and test data (from [29])

The model can produce load forecasts at multiple future time steps, as shown in Figure 4 for the Café building's main power. The models use the most recent 12 hours of weather data to capture the transient effect of weather on energy consumption and use the time of day and day of the week to capture time-based energy consumption patterns, such as those related to building occupancy. Because neural networks can learn non-linear relationships, the trained models are able to infer complex behaviors, such as the pre-occupancy energy spike in the Café due to early morning food preparation. Each model outputs a prediction of building load from the current time step to 18 hours in the future in 1-hour increments. This input-output structure applies to all buildings that were studied on the NREL campus. Model performance was not seen to change significantly as the length of the forecast window increased/decreased, but the training time was significantly impacted by the length of the training window (number of hours of historical data used as input to the model).

4. Operational load forecasting system deployment

A primary research objective of the case study was to establish deployment and continuous improvement pathways for the developed algorithms. That is, developed algorithms should be able to be rapidly deployed for testing within the Intelligent Campus platform and iteratively improved with minimal retooling. The Intelligent Campus team intends to expand the algorithm to forecast other quantities (such as building electricity meters, thermal meters, and photovoltaic system generation) and to support additional exogenous variables. Therefore, from a software development point of view, the system's input-output architecture must be flexible, replicable, and scalable. Due to the nature of the deep learning algorithms used, ongoing access to data is an integral part of the continued operation of the system. Periodic retraining of the models as more data gets collected has the potential to increase model accuracy.

The following two subsections describe the challenges and opportunities that open up for smart buildings' research and prototyping once those challenges are addressed. Understanding the deployment aspect of a deep-learning-based forecasting system, therefore, completes the re-search-prototyping loop.

4.1 Challenges

4.1.1 Data availability and quality control

"Garbage in, garbage out"—a common phrase used in the machine learning community—captures the importance of this step of data collection and curation. Assembling a machine learning training data set begins with proper commissioning and maintenance of meters, sensors, and other data inputs. Data should have proper units and scaling. (Although in some circumstances, a machine learning model could be successfully trained from improperly scaled or otherwise corrupted input data, if those data stream errors were corrected in the future, the model would need to be retrained).

Automated accumulation of new training data (without extensive human review) presents data quality challenges. NREL's Intelligent Campus team has implemented quality check mechanisms for the measurements coming from the sensors in real time before they are fed to the model. One example is removing the outliers identified by the values lying outside of three standard deviations. Another is operating basic fault detection rules that check for stuck sensors (constant value for long periods of time), logically out-of-range data (such as significant solar radiation recorded at night), and similar anomalies. These checks not only improve machine learning model quality but also provide an opportunity to detect and correct errors in the measurement systems and equipment.

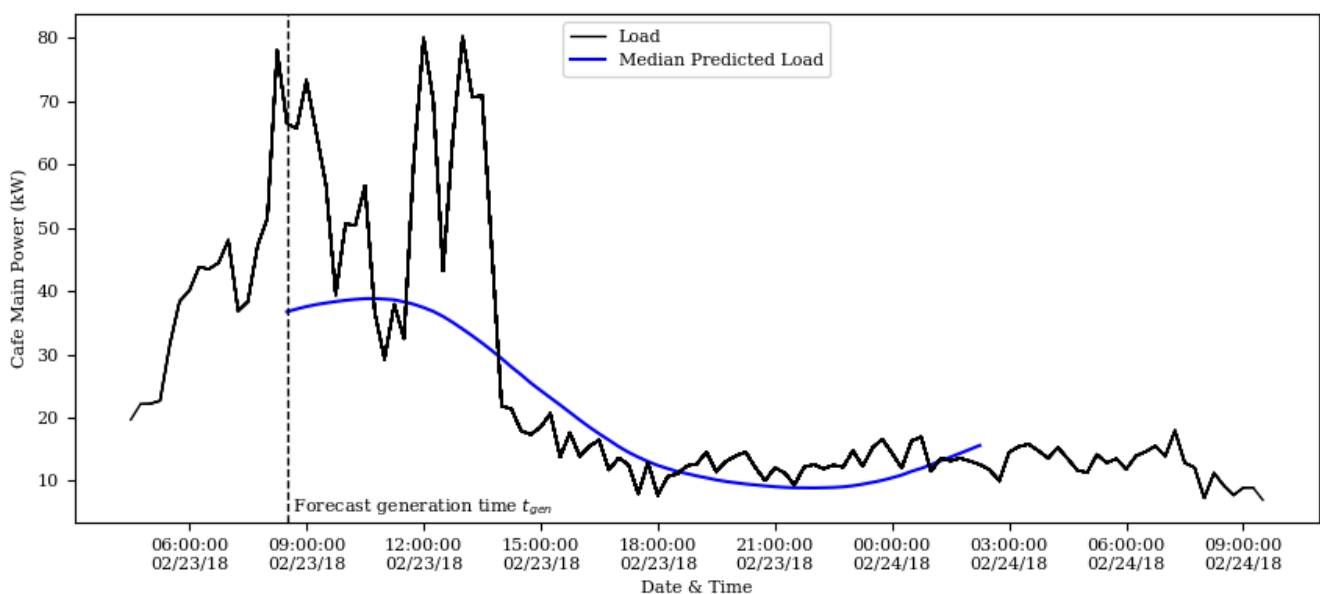


Figure 4. Multi-time horizon forecast (from [29])

4.1.2 Integration with EMIS and scalability

The practical implementation of a forecasting system requires reliable interfaces for transferring data to the model and for transferring model outputs to an analytics system that consumes them. The forecasting system developed for this case study communicates with the NREL EMIS via an application programming interface defined by the Project Haystack standard [33]. The EMIS provides forecast inputs, and the forecast system, in turn, writes forecasts back to the EMIS. Adherence to a metadata standard and an open application programming interface allows the forecasting system to be rapidly retrained and redeployed for new sites. Any database that adheres to the Project Haystack standard and correctly tags the required input points can be connected to a new instance of the deep learning algorithm, and a new model can be trained with minimal effort. When the forecasting system runs in integration with EMIS, it is also important to be able to redeploy the model (with either architectural or training updates) without service interruption. Because forecasts cover a span of time but are frequently updated, managing caching is a key challenge. In addition, the forecast read/write system must be flexible enough to handle arbitrary forecast inputs and outputs. Additionally, given the computationally intensive nature of deep learning models, scalability in terms of available computing power for retraining the models is another interesting challenge.

4.1.3 Online vs. offline training

Training is an important consideration for deep-learning-based forecasting systems because deep-learning models' accuracies have been shown to keep rising with more training. Offline training is a relatively simpler way of implementation in which the machine learning models are retrained by a manual process where an engineer copies the old model to a local machine, trains it further with the new data, and transfers it back to the servers where the systems are deployed. Offline training, therefore, requires more human hours over the course of the system's operation because human intervention is needed to train the models further at regular intervals. The online training approach, on the other hand, is an automated process. At its simplest, an online training system trains models on a schedule (for example, every 1 to 2 months) without any manual intervention [37]. Though setting up the automated training pipeline initially requires additional planning and coding, this online training eliminates the maintenance dependency on human intervention to move the models offline, retrain them, and upload them back into the forecasting system pipeline regularly. Thus, online training can increase operational efficiency and reduce opportunities for human error to impact the forecasting system. Moreover, because a robust metering platform has already been established as the first step in the process of our prototype of Intelligent Campus, high-quality data are constantly accumulating and can be used to increase the accuracy rates of our models. For our final implementation, we have chosen the online training option.

4.1.4 Continuous improvement

As described above, deep learning models are capable of increasing accuracy as more data is made available for

training. Also, in the initial stages, when the accuracy of the model is not satisfactory, given the shortage of data points to train it, large changes in predictive behavior can be off-putting to end users. Therefore, along with iterative training, we conduct model enhancement with hyperparameter tuning as the operational circumstances change over time. This is an infrequent yet desirable step.

4.1.5 Generalizability

Another typical challenge in deploying deep learning models for multiple prediction points (i.e., building-wise energy consumption forecasts) is generalizability. Because each prediction point has its own model, the deployment architecture must address the need for automated model generation. To address the challenge of generalizability, the forecasting system's data preprocessing and postprocessing modules are designed with the capability to take in a generic variable named "point-id" that is used as a reference to fetch the respective machine learning model and input data streams.

4.2 Opportunities

An automated end-to-end pipeline for load forecasting deployed on-site, with an online training mechanism in place, enables many opportunities to harness the predicted data for furthering the smart buildings' research. In the following subsections, we describe three potential applications.

4.2.1 Demand management

Motivations for electricity demand management (peak shaving), including utility demand charge reductions and reduced carbon footprint. (When the peak of multiple buildings and campuses coincides with the utility peak, lower-efficiency "peaking" power plants fired by gas or diesel are used to support the grid, which translates to proportionally higher greenhouse gas emissions.) Reducing peak demand requires the ability to predict the timing and magnitude of peak demand, then preemptively shift load or dispatch energy storage assets to avoid the peak.

4.2.2 Energy arbitrage

Energy arbitrage is another cost-saving opportunity available for grid-interactive buildings in cases where on-site energy storage is available. Energy arbitrage is performed by participating in the energy markets by charging energy storage (for example, batteries) when electricity prices are low and discharging storage to sell power back to the grid when prices spike. Even in utility markets where bidirectional electricity exchange is not possible, a form of energy arbitrage is possible by buying energy when prices are low, storing the energy locally, and dispatching the energy locally to meet load when energy prices are high. Energy arbitrage can be accomplished by employing near-real-time or real-time optimization algorithms. However, the real-time optimized storage dispatch requires reasonable foresight into future on-site energy consumption, thereby making building load forecasts an important component of the system.

4.2.3 Load flexibility

A third opportunity to save energy costs for building operations is shifting the load based on the time-of-use energy rates. Building load can be classified as critical, which consists of loads that cannot be deferred, and flexible, which

consists of loads whose consumption can be scheduled within a given time window. Examples of critical loads include lights, fans, microwaves, computing systems, etc.; these loads must be available for use when needed. Examples of flexible loads include dishwashers, laundry, and even heating and cooling (buildings can be precooled/preheated). With adequate foresight, flexible loads can be shifted to minimize energy costs. This also requires predicting the building's future energy consumption.

5. Conclusion

The paradigm of grid-interactive efficient buildings, i.e., dynamically operating buildings that constantly exchange information and energy with the utility grid, is increasingly gaining traction in both the research community and in industry. Such smart buildings will operate in harmony with the grid to make electricity more affordable and integrate a larger share of DERs while meeting the comfort and productivity needs of the buildings' occupants. NREL's Intelligent Campus is a living laboratory that is dedicated to accelerating the research and deployment of various pieces of the smart buildings puzzle, serving as a vital testbed for the innovative solutions that push the frontier of smart building operations. In this article, we describe and critically discuss a deep learning building load forecasting system that is an important module of smart buildings' overall controls. Through the critical discussion, we highlight that the characteristics beyond simple algorithm forecast accuracy are important in determining the algorithm's true utility for smart buildings. We demonstrate that metering, collection, and curation of data is a crucial part of the prototyping process of such an advanced forecasting system. Using LSTMs, the deep learning model can predict the building and campus load for an 18-hour' time horizon. The architecture is flexible with the time resolution; sub-hourly predictions can be enabled with minimal modifications to the code. Using Intelligent Campus as the platform for testing the research findings in real-time settings, we shed light on the deployment aspect of such advanced forecasting systems. We conclude that a well-architected design of the forecasting system is key to its effectiveness. To take a forecasting algorithm from the research phase to deployment for practical use, building engineers will need to focus on the end-to-end pipeline, which requires considerations for multiple submodules such as access to data, data handling, seamless integration with the existing building data platform, model training, and update mechanism, periodic performance checks, and model enhancement with hyperparameter tuning as operational circumstances change. The end-to-end forecasting system presented in this work forms a robust platform for furthering research in the building energy predictive analytics problem space. We plan to expand this work in the following directions:

5.1 Probabilistic forecasting

Rather than generating a point forecast for every time step, a range can be outputted with the estimated minimum and maximum values serving as the upper and lower bounds for a prediction band within which the actual consumption in the future time step will fall. Having a range instead of point forecasting is especially helpful in the fault detection application where consistent outliers (outside of the band

points) indicate a potential issue that can be further investigated by the building's operations team. Moreover, probabilistic forecasting using deep neural networks is an active inquiry front on the research end of the forecasting technology's spectrum.

5.2 Predictive maintenance

Current practice in the building industry is to perform "schedule-based" maintenance for fault prevention, which is not effective in terms of flagging the equipment vulnerable to a fault beforehand. The next generation of the fault-prevention mechanism is "condition-based" or "preventive" maintenance. This involves using statistical analysis to assess the health of the patients and estimate the probability of failure in order to inform the operators. The presented forecasting system can be further developed for predictive-maintenance applications where the occurrence of a specific fault on the system, which may occur due to unusual wear-and-tear on equipment, can be predicted well in advance, thereby preventing equipment failure from triggering faults.

Author contributions

Conceptualization, Sakshi Mishra, Stephen M. Frank, and Anya Petersen; method-ology, Sakshi Mishra and Stephen M. Frank; software, Anya Petersen and Sakshi Mishra; validation, Stephen M. Frank and Robert Buechler; formal analysis, Sakshi Mishra; resources, Stephen M. Frank and Michelle Slovensky; data curation, Stephen M. Frank; writing—original draft preparation, Sakshi Mishra, Stephen M. Frank, and Anya Petersen; writing—review and editing, Robert Buechler and Michelle Slovensky; visualization, Sakshi Mishra, Robert Buechler, and Marj Schott; supervision, Michelle Slovensky; project administration, Anya Petersen; funding acquisition, Michelle Slovensky. All authors have read and agreed to the published version of the manuscript.

Ethical issue

The authors are aware of and comply with best practices in publication ethics, specifically concerning authorship (avoidance of guest authorship), dual submission, manipulation of figures, competing interests, and compliance with policies on research ethics. The authors adhere to publication requirements that the submitted work is original and has not been published elsewhere in any language.

Funding

This research was funded by the U.S. Department of Energy Office of Energy Efficiency and Renewable Energy Building Technologies Office.

Acknowledgments

This work was authored by NREL, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Funding provided by the U.S. Department of Energy Office of Energy Efficiency and Renewable Energy Building Technologies Office. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. The authors wish to acknowledge the National Renewable Energy Laboratory's continued support in enhancing the research and prototyping capabilities of the Intelligent Campus program. The authors would like to thank Marj Schott (NREL) for contributing Figure 1 to the article.

Data availability statement

Data sharing does not apply to this article as no datasets were generated or analyzed during the current study.

Conflict of interest

The authors declare no potential conflict of interest.

References

- [1] Total Energy - Monthly Energy Review, U.S. Energy Information Administration,, 02 05 2020. [Online]. Available: <https://www.eia.gov/totalenergy/data/monthly/index.php>. [Accessed 02 05 2020].
- [2] U. DOE, Buildings energy data book. Available: <https://data.openei.org/submissions/144>
- [3] P. G. Jordan, "Global Markets" in Solar Energy Markets, 2014, pp. 127-133.
- [4] Energy-Spectral, During COVID-19 Lockdown, Office Buildings are Consuming Even More HVAC Energy Than When Fully Occupied, Spectral Energy, 2020. [Online]. Available: <https://spectral.energy/news/covid19-office-building-energy-consumption-increasing/>. [Accessed 25 07 2020].
- [5] P. Donohoo-Vallett, P. Gilman, D. Feldman, J. Brodrick, D. Gohlke, R. Gravel, A. Jiron, C. Schutte, S. Satyapal, T. Nguyen, P. Scheihing, B. Marshall and S. Harman, "Revolution Now The Future Arrives for Five Clean Energy Technologies," U.S. DOE, 2016.
- [6] M. Neukomm, V. Nubbe and R. Fares, "Grid-interactive Efficient Buildings Technical Report Series," DoE Office of Energy Efficiency & Renewable Energy, 2019.
- [7] B. Kroposki, E. Dall'Anese, A. Bernstein, Y. Zhang and B.-M. Hodge, "Autonomous Energy Grids," 2017.
- [8] A. Afram and F. Janabi-Sharifi, "Theory and applications of HVAC control systems - A review of model predictive control (MPC)," Building and Environment, vol. 72, pp. 343-355, 2014.
- [9] E. Khanmirza, A. Esmaeilzadeh and A. H. D. Markazi, "Predictive control of a building hybrid heating system for energy cost reduction," Applied Soft Computing, vol. 46, pp. 407-23, 2016.
- [10] T.-T. Nguyen, H.-J. Yoo and H.-M. Kim, "Analyzing the Impacts of System Parameters on MPC-Based Frequency Control for a Stand-Alone Microgrid," Energies, vol. 10, no. 4, 2017.
- [11] Y. Iino, M. Murai, D. Murayama, I. Motoyama, S. Kuzusaka and K. Ueta, "Hybrid modeling with physical and JIT model for building thermal load prediction and optimal energy saving control," in ICCAS-SICE, Fukuoka, Japan, 2009.
- [12] F. Zhang and X. Zhou, "Gray-Regression Variable Weight Combination Model for Load Forecasting," in 2008 International Conference on Risk Management & Engineering Management, Beijing, China, 2008.
- [13] Y. Iwafune, Y. Yagita, T. Ikegami and K. Ogimoto, "Short-term forecasting of residential building load for distributed energy management," in 2014 IEEE International Energy Conference (ENERGYCON), Cavtat, Croatia, 2014.
- [14] Itron, "Energy Forecasting". <https://na.itron.com/>
- [15] Enverus, "Electric Gas and Load Forecast". <https://www.enverus.com/>
- [16] B. Sohlberg and Jacobsen, "Grey Box Modelling - Branches and Experiences," IFAC Proceedings Volumes, vol. 41, no. 2, pp. 11415-11420, 2008.
- [17] D. B. Crawley, L. K. Lawrie, F. C. Winkelmann, W. F. Buhl, Y. H. C. O. Pedersen, R. K. Strand, R. J. Liesen, D. E. Fisher, M. J. Witte and J. Glazer, "EnergyPlus: creating a new-generation building energy simulation program," Energy and Buildings, vol. 33, no. 4, pp. 319-331, 2001.
- [18] EERE, "EnergyPlus," energy.gov, 2020. [Online]. Available: <https://www.energy.gov/eere/buildings/downloads/energyplus-0>.
- [19] S. S. Kwok and E. W. M. Lee, "A study of the importance of occupancy to building cooling load in prediction by intelligent approach," Energy Conversion and Management, vol. 52, no. 7, pp. 2555-2564, 2011.
- [20] S. Mishra, A. Glaws and P. Palanisamy, "Predictive Analytics in Future Power Systems: A Panorama and State-Of-The-Art of Deep Learning Applications," in Optimization, Learning, and Control for Interdependent Complex Networks, SpringerLink, 2020, pp. 147-182.
- [21] J. Hwang, D. Suh and M.-O. Otto, "Forecasting Electricity Consumption in Commercial Buildings Using a Machine Learning Approach," Energies, vol. 13, no. 22, 2020.
- [22] I. S. Jahan, V. Snasel and S. Misak, "Intelligent Systems for Power Load Forecasting: A Study Review," Energies, vol. 13, no. 22, 2020.
- [23] T. Ciecchulski and S. Osowski, "Deep Learning Approach to Power Demand Forecasting in Polish Power System," Energies, vol. 13, no. 22.
- [24] Y. Jin, H. Guo, J. Wang and A. Song, "A Hybrid System Based on LSTM for Short-Term Power Load Forecasting," Energies, vol. 13, no. 23, 2020.
- [25] G. Chitalia, M. Pipattanasomporn, V. Garg and S. Rahman, "Robust short-term electrical load forecasting framework for commercial buildings using deep recurrent neural networks," Applied Energy, vol. 278, 2020.
- [26] M. N. Fekri, H. Patel, K. Grolinger and V. Sharma, "Deep learning for load forecasting with smart meter data: Online Adaptive Recurrent Neural Network," Applied Energy, vol. 282, 2021.
- [27] N. Gautam, A. S. Mayal, V. S. Ram and A. Priya, "Short Term Load Forecasting Of Urban Loads Based On Artificial Neural Network," in 2nd International Conference on Power and Embedded Drive Control (ICPEDC), 2019, 2019.
- [28] J. Cui, Q. Gao and D. Li, "Improved Long Short-Term Memory Network Based Short Term Load Forecasting," Hangzhou, China, 2020.
- [29] S. Mishra, S. M. Frank, A. Petersen, R. Buechler and M. Slovensky, "Deep-Learning-Based, Multi-Timescale Load Forecasting in Buildings: Opportunities and Challenges from Research to Deployment (preprint)," vol. 1, no. 1, pp. 1 - 13, 2020.
- [30] D. Cutler, S. Frank, M. Slovensky, M. Sheppy and A. Petersen, "Creating an Energy Intelligent Campus: Data

- Integration Challenges and Solutions at a Large Research Campus," in ACEEE Summer Study on Energy Efficiency in Buildings, Pacific Grove, CA, 2016.
- [31] ASHRAE, "Energy Standard for Buildings Except Low-Rise Residential Buildings, Standard 90.1-2019," 2019.
- [32] T. Stoffel and A. Andreas, "NREL Solar Radiation Research Laboratory (SRRL): Baseline Measurement System (BMS)," National Renewable Energy Laboratory, Golden, Colorado (Data), 1981.
- [33] "Project Haystack," 2020. [Online]. Available: <https://project-haystack.org/>.
- [34] G. P. Henze, S. Pless, A. Petersen, N. Long and A. T. Scambos, "Control limits for building energy end use based on frequency analysis and quantile regression," Energy Efficiency, vol. 8, p. 1077–1092, 2015.
- [35] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," Neural Computation, vol. 9, no. 8, pp. 1735-1780, 1997.
- [36] S. Mishra and P. Palanisamy, "An Integrated Multi-Time-Scale Modeling for Solar Irradiance Forecasting Using Deep Learning," arxiv.org, 2019.
- [37] Cloud-Architecture-Center, Google, "MLOps: Continuous delivery and automation pipelines in machine learning," 2021.



This article is an open-access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).