

Article

Interval Load Forecasting for Individual Households in the Presence of Electric Vehicle Charging

Raiden Skala ¹, Mohamed Ahmed T. A. Elgalhud ¹, Katarina Grolinger ^{1,*} and Syed Mir ²¹ Department of Electrical and Computer Engineering, Western University, London, ON N6A 5B9, Canada² London Hydro, London, ON N6A 4H6, Canada

* Correspondence: kgroling@uwo.ca; Tel.: +1-519-661-2111 (ext. 81407)

Abstract: The transition to Electric Vehicles (EV) in place of traditional internal combustion engines is increasing societal demand for electricity. The ability to integrate the additional demand from EV charging into forecasting electricity demand is critical for maintaining the reliability of electricity generation and distribution. Load forecasting studies typically exclude households with home EV charging, focusing on offices, schools, and public charging stations. Moreover, they provide point forecasts which do not offer information about prediction uncertainty. Consequently, this paper proposes the Long Short-Term Memory Bayesian Neural Networks (LSTM-BNNs) for household load forecasting in presence of EV charging. The approach takes advantage of the LSTM model to capture the time dependencies and uses the dropout layer with Bayesian inference to generate prediction intervals. Results show that the proposed LSTM-BNNs achieve accuracy similar to point forecasts with the advantage of prediction intervals. Moreover, the impact of lockdowns related to the COVID-19 pandemic on the load forecasting model is examined, and the analysis shows that there is no major change in the model performance as, for the considered households, the randomness of the EV charging outweighs the change due to pandemic.

Keywords: residential load forecasting; Long Short-Term Memory; Bayesian Neural Network; Bayesian optimization; interval forecasting; EV charging; analysis of variance



Citation: Skala, R.; Elgalhud, M.A.T.A.; Grolinger, K.; Mir, S. Interval Load Forecasting for Individual Households in the Presence of Electric Vehicle Charging. *Energies* **2023**, *16*, 4093. <https://doi.org/10.3390/en16104093>

Academic Editors: Antonio Gabaldón, María Carmen Ruiz-Abellón and Luis Alfredo Fernández-Jiménez

Received: 20 March 2023

Revised: 8 May 2023

Accepted: 10 May 2023

Published: 15 May 2023



Copyright: © 2023 by the authors. Licensee MDPI, Basel, Switzerland. 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/>).

1. Introduction

Affordable and reliable sources of electricity enable the sustainable growth of strong economies and can improve the average person's quality of life [1] by providing reliable access to appliances, medical equipment, communication, entertainment, and other devices. The dependence on power grids to provide electricity is increasing due to the continuous integration of novel electronic devices into every aspect of modern life [2] as these devices rely on a reliable source of electricity. The mainstream adoption of electric vehicles (EVs) away from traditional internal combustion engine (ICE) vehicles for consumer use is set further to entrench reliance on access to electricity due to the increased electricity demand for charging. While this transition can be a positive step in reducing carbon emissions [3], embracing EVs will shift transportation energy requirements from petroleum-based products to electric grids. Of specific interest to this paper, is the demand created by charging EVs in residential households, which is frequently utilized due to its relative affordability and convenience.

As countries, such as Canada, plan to ban the sale of new ICE vehicles by 2035 [4], preparations are required to ensure the success of this shift. This includes actions such as installing charging stations, increasing electricity generation capacity, investing into battery technologies, and improving infrastructure throughout the grid to handle the higher loads required by EV charging. The capability of electricity distribution companies to accurately forecast the hourly electricity consumption of residential households that own EVs is instrumental in the transition to EVs as it assists the utility companies to

anticipate and manage increased energy demand, plan sufficient capacity to meet expected demand, and ensure grid stability. Failure in predictive ability poses a risk to the balance between electricity supply and demand, which can impose serious threats to grid stability, human life, and overall economic interest [5]. A loss in the balance between supply and demand can lead to power outages, brownouts, and other disruptions. In turn, these electricity disruptions can severely disrupt critical infrastructures including transportation, communications, and financial systems, as well as essential services such as emergency response and healthcare.

There have been extensive efforts to create predictive load forecasting models using machine learning (ML) with historical energy consumption data collected by smart meters or similar technologies, often combined with meteorological information [6,7]. In recent years, deep learning techniques, especially those based on Recurrent Neural Networks (RNNs) have been outperforming other techniques [8,9]. While these studies had great successes in terms of accuracy for a variety of use cases, they do not specifically address short-term forecasting for individual households in presence of EV charging [10] which introduces challenges due to variations of charging patterns.

Moreover, most predictive models for load forecasting generate point predictions instead of an interval for their expected electricity demand [10], which limits the usefulness of the forecast in decision-making. By providing only a single value for expected electricity demand, these models fail to convey the range of potential outcomes and the degree of uncertainty associated with each prediction. Furthermore, providing only a point forecast, without the range values, may not offer sufficient information for effective risk management. In contrast, interval forecasting approaches provide decision-makers with a more nuanced understanding of the possible outcomes, enabling them to make more informed and effective decisions. For example, considering the full range of possible outcomes, instead of a single value, allows the stakeholders to plan for different scenarios and better mitigate risks.

To address these drawbacks, this paper proposes a probabilistic interval forecasting approach for predicting the hourly electricity demand in households with EV charging. By using probabilistic methods, our approach generates a range of likely outcomes rather than a single-point estimate which provides a more comprehensive understanding of the potential effects of EV charging on household electricity demand, gives information about the uncertainty associated with the predicted value due to dynamic charging behaviors, and offers decision-makers a more complete picture of the forecasted demand. The interval predictions are generated with Long Short-Term Memory Bayesian Neural Networks (LSTM-BNNs). LSTM was chosen as it is well-suited for capturing temporal dependencies in data while BNN was added to estimate the probability distribution of expected values for interval predictions. LSTM-BNN was trained using historical household electricity consumption data and local temperature data. To assess the effectiveness of the proposed LSTM-BNN model, its performance, measured using four metrics, is compared to the performance of the standard point prediction LSTM model. Additionally, due to the impact of the COVID-19 pandemic on electricity consumption patterns, the point and interval models have been examined on two datasets: one with the lockdown period and one without. The results show that the accuracy greatly varies among households, but for each household, the proposed LSTM-BNN achieves similar accuracy to point forecasts while providing the advantage of prediction intervals.

The remainder of the paper is organized as follows: Section 2 provides background information on LSTM and BNN techniques and introduces the four common performance measurements used for gauging the effectiveness of regression models while Section 3 reviews related work on load forecasting and interval forecasting. The proposed LSTM-BNN interval forecasting approach is described in Section 4 followed by the evaluation presented in Section 5. Finally, Section 6 concludes the paper.

2. Background

This section begins by introducing Long Short-Term Memory (LSTM) networks and Bayesian Neural Networks (BNNs), followed by a discussion of the four performance measures commonly used for assessing regression models.

2.1. Long Short-Term Memory Neural Network

Neural networks are a type of machine learning model inspired by the human brain: they use interconnected artificial neurons to learn and process information by mimicking the way biological neurons signal to one another [11]. A recurrent neural network (RNN) is a type of neural network designed to process sequential data by using internal memory and recurrent connections, allowing it to capture temporal dependencies and patterns in the data. A Long Short-Term Memory (LSTM) neural network model is similar to RNN models in that it can capture temporal relationships by using an internal memory mechanism to keep track of past inputs and selectively remember or forget certain information. The main difference between an LSTM and RNN model is that LSTM models have additional structures, such as gating mechanisms, that provide better control over the flow of gradients and help prevent the vanishing and exploding gradient problems that can occur in standard RNNs, making them more effective for modeling longer sequences of data. LSTM computation at time t is given as follows:

$$f_t = \sigma(W_{fx}x_t + W_{fh}h_{t-1} + b_f) \quad (1)$$

$$i_t = \sigma(W_{ix}x_t + W_{ih}h_{t-1} + b_i) \quad (2)$$

$$o_t = \sigma(W_{ox}x_t + W_{oh}h_{t-1} + b_o) \quad (3)$$

$$\bar{C}_t = \varphi(W_{cx}x_t + W_{ch}h_{t-1} + b_c) \quad (4)$$

$$C_t = f_t \odot C_{t-1} + i_t \odot \bar{C}_t \quad (5)$$

$$h_t = o_t \odot \varphi(C_t) \quad (6)$$

Equations (1)–(3) depict the computation at the forget f_t , input i_t , and output o_t gates respectively, while Equations (4)–(6) determine the cell state C_t and hidden state h_t . The sigmoid (σ) and tanh (φ) functions contribute to controlling exploding gradients by keeping values between zero to one and negative one to one respectively. The current cell input x_t and previous cell hidden state $h_{(t-1)}$ are the inputs received by the LSTM cell. The gate biases b_f , b_i , b_o , and b_c , the current cell weight matrices W_{fx} , W_{ix} , W_{ox} , and W_{cx} and the hidden state cell weight matrices W_{fh} , W_{ih} , W_{oh} , W_{ch} of each LSTM cell are adjusted throughout the training process using backpropagation through time with the goal of minimizing the loss between the predicted and true values. The use of \odot indicates computing the elementwise Hadamard product of two matrices.

Due to its ability to capture temporal dependencies over long periods of time, the LSTM model has been very successful in many domains including load forecasting [12,13]. For the same reason, we use the LSTM cells in the proposed LSTM-BNN interval forecasting approach.

2.2. Bayesian Neural Network

The Bayesian Neural Network (BNN) model [14] relies on Bayesian inference to determine the posterior predictive distribution with the ultimate goal of quantifying the uncertainty introduced by the models so as to explain the trustworthiness of the prediction. This is achieved by incorporating previous inputs X and outputs Y as well as model parameters ω in Bayes' theorem as follows:

$$P(\omega|X, Y) = \frac{P(Y|X, \omega) \cdot P(\omega)}{P(Y|X)} \quad (7)$$

Here, $P()$ indicates the probabilities and $P(\cdot|\cdot)$ are conditional probabilities.

By computing the integral of the full posterior distribution, given in (7), multiple times using different samples from the model parameters, a distribution can be generated for a predicted value y_{new} using new inputs x_{new} :

$$P(y_{new}|x_{new}, X, Y) = \int P(y_{new}|x_{new}, \omega) \cdot P(\omega|X, Y) d\omega \quad (8)$$

However, due to the full posterior probability being computationally demanding for deep neural networks, alternative approaches are required to make the use of Bayesian inference feasible in practice. Zhang and Mahadevan [15] demonstrated Monte Carlo dropout remaining active while a network generates predictions to be sufficient for approximating the posterior predictive distribution as it minimizes the relative entropy between the approximate and true posterior distributions while remaining computationally feasible. Consequently, our approach takes advantage of BNN and the dropout technique to generate interval load forecasts for households with EVs.

2.3. Performance Metrics

The four prominent performance metrics that are used for evaluating the margin of error between a prediction made by a machine learning model and the true value are: Mean Absolute Percent Error (MAPE), Mean Square Error (MSE), Root Mean Square Error (RMSE), and Mean Absolute Error (MAE) [12,15]. While there are other metrics specifically dealing with evaluating probabilistic forecasts [16], we primarily use the mentioned metrics as they allow us to compare point and interval forecasts. These metrics are calculated according to the following equations:

$$\text{MAPE} = \frac{100\%}{N} \sum_{i=1}^N \frac{|y_i - \hat{y}_i|}{y_i} \quad (9)$$

$$\text{MSE} = \frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2 \quad (10)$$

$$\text{RMSE} = \sqrt{\left(\frac{1}{N}\right) \sum_{i=1}^N (y_i - \hat{y}_i)^2} \quad (11)$$

$$\text{MAE} = \frac{1}{N} \sum_{i=1}^N |y_i - \hat{y}_i| \quad (12)$$

where y_i is the true value of the i -th sample, \hat{y}_i is the predicted value for the i -th sample and N is the total number of samples.

MAPE has an advantage over the other three metrics as it is a scale-independent metric representing the error as a percentage of the actual value and therefore suitable for comparing models on datasets of different value scales. MSE and RMSE metrics are both based on the Euclidean distance to determine the level of error between predicted and true values. The difference between the MSE and RMSE metrics is that MSE provides more severe punishment for predictions that are very different from the true value. MAE is used to measure the mean absolute difference between predictions and true values and is less severe at penalizing large differences between predicted and true values than MSE or RMSE. To obtain a different view of forecasting accuracy, our study employs all four metrics.

3. Related Work

This section first reviews recent load forecasting studies focusing on those based on machine learning and then discusses techniques for interval predictions in different domains.

3.1. Electricity Load Forecasting

This subsection first reviews recent load forecasting studies for a diversity of consumers including residential households and buildings. This provides insights into state-

of-the-art models and represents directions for forecasting in the presence of EV changing. Next, related work in predicting EV charging in various settings is examined.

An LSTM-based model for short-term load forecasting on the individual household level was proposed by Kong et al. [17]. They [17] found that a significant hurdle to creating forecasts at the household level is the large degree of diversity and volatility in energy consumption between households when compared to making forecasts at the substation level. This difficulty in residential forecasting due to load variability and concept drift also aligns with the findings of Fekri et al. [12,18].

Residential load forecasting was also investigated by Zhang et al. [19]: while Kong et al. [17] used LSTM-based approach Zhang et al. [19] employed Support Vector Regression (SVR). In their study, Zhang et al. [19] investigated predicting daily and hourly electricity consumption for 15 households with data obtained from smart meters. The accuracy of the load predictions varied significantly across households, depending on the variability of energy-related behaviors among occupants. Daily load estimates were generally more accurate, as they mitigated the randomness in hourly changes.

L'Heureux et al. [20] presented transformer-based architecture for electrical load forecasting. They adapted the transformer model from the Natural Language Processing (NLP) domain for load forecasting by modifying the NLP transformer workflow, adding N-space transformation, and designing a novel technique for handling contextual features. They examined the proposed transformer-based architecture on 19 different data streams and with four different forecasting horizons. For most data streams and forecasting horizons, the transformer accuracy was better than the Seq2Seq network; however, for 12-h forecast Seq2Seq was slightly better.

A multi-node load forecasting was investigated by Tan et al. [21]: they proposed multi-task learning combined with a multi-modal feature module based on an inception-gated temporal convolutional network for node load prediction. The feature extraction module captures the coupling information from the historical data of the node, while the multi-task learning utilizes a soft sharing mechanism to leverage the shared information across nodes to improve the forecast accuracy. Experimental results demonstrate the effectiveness of the proposed method in accurately forecasting load demand across multiple nodes.

Ribeiro et al. [22] investigated short- and very short-term load forecasting for warehouses and compared several machine learning and deep learning models including linear regression, decision trees, artificial neural networks, and LSTM models. In their experiments RNN, LSTM, and GRU cells achieved comparable results. Jian et al. [23] also worked on very short-term load forecasting: they proposed a framework based on an autoformer which combines decomposition transformers with auto-correlation mechanism. Multi-layer perceptron layers are added to the autoformer for an improved deep information extraction. In their experiments, the proposed deep-autoformer framework outperformed several deep-learning techniques on the task of very short-term residential load forecasting.

An encoder-decoder RNN architecture with a dual attention mechanism was proposed by Ozcan et al. [6] to improve the performance of the RNN model. The attention mechanism in the encoder helps identify important features whereas the attention in the decoder assists the context vector and provides longer memory. In their experiments, the encoder-decoder RNN architecture achieved improved accuracy in comparison to LSTM; however, the computation complexity was increased.

Short-term load forecasting has been investigated by Sun et al. [8]: they proposed a framework based on LSTM and an enhanced sine cosine algorithm (SCA). The authors enhanced the performance of the SCA, a meta-heuristic method for optimization problems, by incorporating a chaos operator and multilevel modulation factors. In experiments, they compared the modified SCA with several other population intelligence algorithms including particle swarm optimization and the whale optimization algorithm and showed that SCA improves performance.

There are very few studies concerned with load forecasting for EV charging demand and they mostly consider scenarios such as parking lots, fleets, and regional demand. For

example, Amini et al. [24] investigated forecasting of EV charging demand for parking lots. Their approach used an Autoregressive Integrated Moving Average (ARIMA) model with driving patterns and distances driven as inputs to determine the day-ahead demand of the conventional electrical load and charging demand of EV parking lots. Two simulated test systems, 6-bus and IEEE 24-bus systems, were used to examine the effectiveness of the proposed approach.

Yi et al. [25] highlighted the importance of accurate demand forecasting for planning and management of electric vehicle charging infrastructure. They presented a deep learning-based method for forecasting the charging demand of commercial EV charging stations by utilizing LSTM as a base for the Seq2Seq model and combining it with a clustering technique. The evaluation on over 1200 charging sites from the State of Utah and the City of Los Angeles showed that the proposed method outperforms other forecasting models such as ARIMA, Prophet, and XGBoost. For forecasting EV charging demand at charging stations in Colorado, Koohfar et al. [26] proposed a transformer-based deep learning approach. The proposed approach was compared to time-series and machine learning models including ARIMA, SARIMA, LSTM, and RNN. While for longer time horizons the transformer outperformed other techniques, for short-term forecasting (7 days ahead), LSTM and transformer achieved comparable results.

A multi-feature data fusion technique combined with LSTM was proposed by Ad-uama et al. [27] to improve the EV charging station load forecasting. They generate three sets of inputs for LSTM consisting of load and weather data pertaining to different historical periods. These three sets of data are then passed to the LSTM models which generate three predictions, and, finally, the LSTM outputs are combined using a data fusion technique. In their experiments, the proposed fusion-based approach achieved better accuracy than traditional LSTM in predicting EV charging station demand.

Zheng et al. [28] were interested in predicting the overall load from EVs in the city of Shenzhen, China. They recognize the diversity of charging patterns and therefore break down the fleet into four groups: private EVs, taxis, busses, and official EVs. Their approach provides a mid-and-long term EV load charging model based on the current utilization of EVs in Shenzhen using probabilistic models for EV charging profiles and forecasting EV market growth in the city using the Bass model. As they are concerned with the regional EV demand, some of the randomnesses of the individual EV charging is remedied through aggregation. Similarly, Arias and Bae [29] considered forecasting load for groups of EVs. Specifically, they take advantage of historical traffic data and weather data to formulate the forecasting model. First, traffic patterns are classified, then factors influencing traffic patterns are identified, and finally, a decision tree formulates the forecasting model.

Strategies for handling growing EV charging demand were investigated by Al-Ogaili et al. [10]. They classify the EV control strategies into scheduling, clustering, and forecasting strategies recognizing that precise estimates of charging are critical for fault prevention and network stability. They note that the stochastic nature of EV charging demand requires advanced forecasting techniques, commonly combined with the need for extensive data including historical charging data, weather, and travel patterns, which may not be readily available. Forecasting studies Al-Ogaili et al. [10] examined include predictions for groups of EVs or geographical regions, charging stations, and specific types of EVs (e.g., busses).

The reviewed studies [6,8,12,17–23] on generic load forecasting represent the state-of-the-art in energy forecasting but their behavior in presence of EV charging has not been examined. Nevertheless, they represent a great foundation for forecasting EV charging load. On the other hand, EV-related studies [10,24–29] do consider EV charging but they do so for a group of EVs, parking lots, charging-station, or regions, and do not confider forecasting load for individual households in presence of EVs. In contrast, we focus on predicting power consumption for individual households in presence of EV charging. Moreover, in contrast to point predictions provided in the aforementioned studies, our study offers interval predictions.

3.2. Interval Predictions

This subsection reviews approaches that have been taken by authors across different domains to create regression models that provide an interval for predictions. In contrast to point predictions, interval predictions quantify uncertainties and provide additional information for decision-making.

Interval predictions were generated for electricity spot pricing by Maciejowska et al. [30] for the British power market using factor quantile regression averaging. First, point predictions are obtained with a collection of models including autoregressive models, threshold autoregressive models, semiparametric autoregressive models, neural networks, and others. Next, point predictions generated by the mentioned models are combined using quantile regression averaging to provide final interval forecasts. The proposed approach performed better than the benchmark autoregressive model.

Shi et al. [31] considered interval predictions for forecasting wind power generation to quantify uncertainties in renewable energy generation. They train an RNN model with two outputs, one for the upper and one for the lower bound of a regression interval of predictions using the Lower and Upper Bound Estimation (LUBE) method. A new cost function incorporating prediction interval was designed and the dragonfly algorithm was introduced to tune the parameters of the RNN prediction model. One of the major challenges associated with training neural networks using the LUBE method is the difficulty in achieving convergence and occasionally the model may not converge [32]. Consequently, Kabir et al. [32] developed a customizable cost function to improve the convergence of LUBE models and assist in constructing prediction intervals with neural networks.

Zhang and Mahadevan [15] proposed interval forecasting for flight trajectory prediction and safety assessment by combining deep learning with uncertainty characterized by a Bayesian approach. Two types of Bayesian networks (BNN), feedforward neural networks and LSTM networks, are trained from different perspectives and then blended to create final predictions. In both BNNs, the dropout strategy quantifies model prediction uncertainty. The BNN approach was also successful in the work of Niu and Liang [33] where they improve nuclear mass and single-neutron separation energy prediction accuracy for determining nuclear effective reactions. In their experiments, Niu and Liang [33] demonstrate that a Bayesian approach can be combined with various forecasting techniques to improve nuclear mass predictions.

The reviewed studies [15,30–33] created interval predictions with various machine learning and statistical methods in various domains; however, none of them considered forecasting household electricity load in the presence of EV charging. Like our study, the works of Zhang and Mahadevan [15] and Niu and Liang [33] also employed BNN techniques to create interval prediction but they used it for very different use cases than load forecasting (flight trajectory [15] and nuclear mass predictions [33]).

4. Interval Load Forecasting in Presence of EV Charging

This section presents the problem formulation and methodology of the proposed interval forecasting for household load prediction in the presence of EV charging. The approach uses only historical energy consumption data obtained from smart meters and weather data which makes it practical and scalable for real-world applications as there is no need to collect data regarding EV charging habits or EV specifications.

Problem Statement: Consider a time series of historical data for a household with EV charging represented as a sequence of input-output pairs (x_t, y_t) where x_t is a vector of features describing the state of the electricity consumption at time t including contextual factors such as temperature, time of day, day of the week, and day of the year, and y_t is a vector of real-valued electricity consumption values for this household at time t . The goal is to learn a probabilistic model $p(\hat{y}_{t+1}|x_{t+1}, D)$, where D represents historical observations and \hat{y}_{t+1} represents the predicted value, that can predict the output for a new input with uncertainty quantification represented as an interval I .

This interval is created by generating multiple predictions through different network configurations to obtain the Bayesian approximation of the predicted value (interval center) as shown in Equation (13). The minimum and the maximum of the interval are computed as shown in Equations (14) and (15) respectively:

$$E[\hat{y}_{t+1}] = \frac{1}{N} \sum_{i=1}^N \hat{y}_{t+1}^i \quad (13)$$

$$I_{min} = E[\hat{y}_{t+1}] - \sigma_{\hat{y}_{t+1}} \quad (14)$$

$$I_{max} = E[\hat{y}_{t+1}] + \sigma_{\hat{y}_{t+1}} \quad (15)$$

where N is the number of predictions generated for the time step $(t + 1)$ and $\sigma_{\hat{y}_{t+1}}$ is the standard deviation of the predictive distribution computed as follows:

$$\sigma_{\hat{y}_{t+1}} = \sqrt{\frac{\sum_{i=1}^N (\hat{y}_{t+1}^i - E[\hat{y}_{t+1}])^2}{N}} \quad (16)$$

The overall interval forecasting process is shown in Figure 1, while details of each component are described in the following subsections.

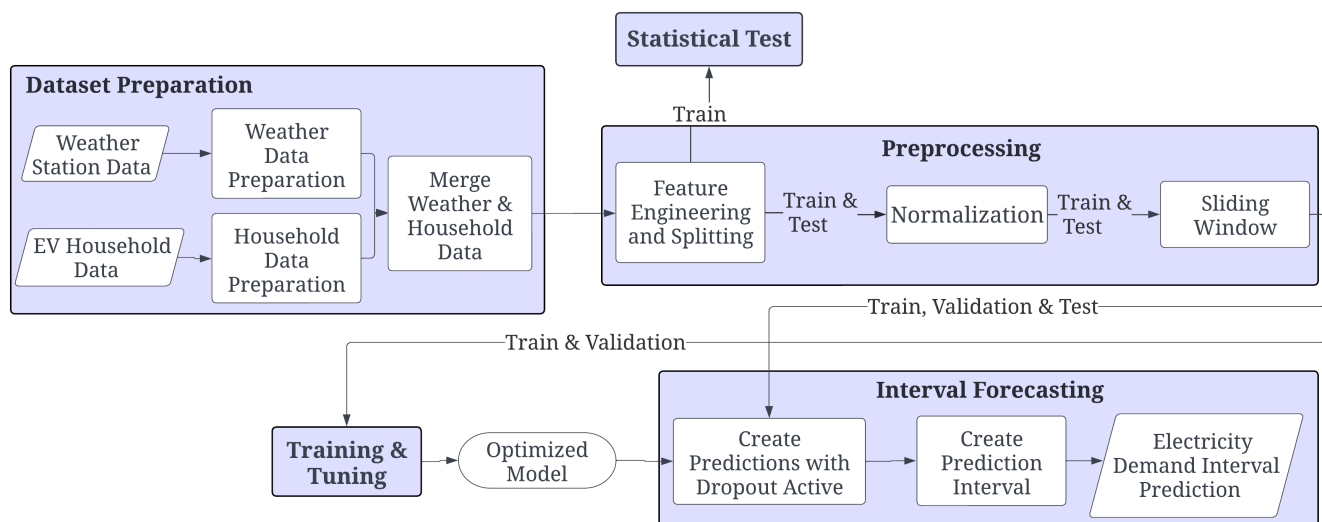


Figure 1. Overview of the proposed LSTM-BNNs interval forecasting approach for household load prediction in presence of EV charging.

4.1. Dataset Preparation

The two types of datasets being used are weather station data and historical household electricity consumption data. Each dataset undergoes preparation individually before they are merged.

The weather station data consists of multiple datasets from multiple weather stations in the approximate geographical areas surrounding the EV household. The features used from the weather station data are the hourly timestamps and the temperature recordings as the temperature is often considered the most influential weather factor in load forecasting [34]. The Weather Data Preparation conducted on the individual weather station datasets and shown in Figure 1 includes filling in missing temperature readings and combining all the weather station data. Missing temperature readings are filled using weighted averaging of the nearest complete temperatures. Since beyond the city, there are no additional geographical details given for any of the EV households, the temperature data from all stations are combined by averaging the temperatures from several weather stations to create a single average temperature dataset. In order to match the timestamps in the EV

Household datasets, the timestamps in the average temperature dataset are adjusted to adhere to daylight savings time (DST).

The household data here refers to hourly data recorded by the smart meter or a similar device. The two initial features in this data are the consumption period and the electricity consumed within that period. The consumption period for the household data initially contains both the start date and time, and the end date and time of the current electricity consumption period. These data, as indicated in Figure 1, undergo Household Data Preparation which involves isolating and only keeping the start date-time of the consumption period so that it can be merged with the weather station data. The electricity consumption feature from the initial dataset remains unchanged.

As part of the processing for both the weather station and household data, an additional time feature is generated. This feature is necessary because at the end of DST each year, the time is set backward one hour, resulting in two instances of the same date-time. This creates a conflict in merging the weather station and household data using only the date-time feature, as there are duplicate non-unique date-times that have no distinguishing differences. The additional time feature is added to both the household and weather station data to indicate if the specific date-time falls in DST or not. This removes the merging ambiguity for the duplicate date-times, as only the first occurrence of the date-time will occur during DST.

After the initial preparations are completed, each of the individual household datasets is merged with the average temperature dataset using the date-time and the additional time feature. In preparation for machine learning, the merged dataset proceeds to the preprocessing step.

4.2. Preprocessing

After the weather and household datasets are merged, the dataset undergoes the following preprocessing steps: feature engineering, splitting the data into train and test sets, and normalizing the train and test sets. Feature engineering step takes advantage of the recorded interval start date-time from the original household data to generate nine features as shown in Table 1. The purpose of creating additional features is to provide context information to the machine learning model, enabling it to generate better predictions.

Table 1. Features extracted from the reading date and time.

Date-Time Feature	Feature Description
Day of week	Range from 0 to 6 starting on Monday
Quarter of year	Range from 1 to 4 starting with January to March
Month of year	Range from 1 to 12 starting in January
Day of the Year	Range from 1 to 365 (366 for leap year)
Day of the Month	Range from 1 to 31
Week of the Year	Range from 0 to 53 for weeks starting on Monday
Hour of the Day	Range from 0 to 23 starting at 12:00AM
Year	Range from 2018 to 2020
DST	Value of 1 (not DST) or 0 (DST)

Following feature engineering, the dataset was split into training and testing sets: the last 10% of readings are assigned to the test set and the remaining to the training set. Then, a portion of the training set was separated to use as the validation set for model selection. As a result of the validation set creation, the distribution of the data becomes 80% for training, 10% for validation, and 10% for testing.

Next z-score normalization was used to reduce the dominance of large features and improve convergence. This technique was chosen over other normalization techniques as it is good at handling outliers present during peak electricity consumption events. The

z-score normalization transforms the feature to the mean of 0 and the standard deviation of 1 as follows:

$$z_{ij} = \frac{x_{ij} - \mu_j}{\sigma_j} \quad (17)$$

where x_{ij} and z_{ij} are the initial unscaled and scaled values of the i -th sample of the j -th feature respectively, and μ and σ are the mean and standard deviation of all the samples of the j -th feature respectively. Note that the mean and standard deviations are calculated only on the training set to avoid data leakage.

Next, the sliding window technique is employed to prepare data for the machine learning model and to provide the model with a fixed number of previous electricity consumption, time-date, and temperature features as the inputs for predicting the next time step. This is accomplished by creating an input window that contains all features including the time-step, temperature, and consumption data within the window size [18]. For instance, for a window of size w , electricity consumption together with all other features for the past w time steps are used as the input for predicting the next energy consumption values. The window slides for s steps to create the next sample. The advantage of the windowing technique for electricity forecasting is in allowing the model to consider the demand for recent time steps when making predictions. The exact window size w is determined within the optimization process.

The sliding window technique is applied to each training, validation, and test set. After this step, the samples have the dimension of $w \times f$ where f is the number of features.

4.3. Training and Tuning

As shown in Figure 1, the Training and Tuning stage follows the Preprocessing step. The deep learning technique LSTM was selected as the machine learning model because in recent years it has demonstrated great successes in load forecasting and outperformed other forecasting techniques [12,18]. The hyperparameter search was carried out with Bayesian optimization as unlike grid or random search, this method performs a more directed exploration of a defined tuning space by selecting hyperparameters that lead to a local, ideally global, minimum loss [35]. The Bayesian optimization achieves a minimum loss by using the posterior distribution of the Mean Squared Error (MSE) loss function determined by previous models to guide the selection of new hyperparameter combinations. This directed selection process minimizes the time and computational resources needed for the exploration of the defined hyperparameter space [35].

In this work, the search space explored included window size, batch size, the number of LSTM layers, the number of neurons in each LSTM layer, the learning rate for the Adam optimizer, and the dropout probability. The window size determines the number of previous time steps to be used as the input to the network while the batch size specifies the number of training windows a single batch contains. The number of LSTM layers and the number of LSTM neurons are adjusted to find a balance between increasing model complexity and variance to fit the training set while maximizing the model's ability to generalize and make accurate predictions when given novel data points. The learning rate for the Adam optimizer is a critical parameter for training each model as it determines the rate at which updates are made to the weight and bias parameters of the model. Using a learning rate that is suitable for finding minima in the loss function enables the model to converge efficiently.

Finally, the dropout probability hyperparameter is used to prevent overfitting: it determines the probability that neurons in a layer will randomly be given zero values. Using a dropout probability that is too high can have a detrimental effect on overall performance as it could result in too many inactive neurons and prevent the model from learning. As the dropout technique is typically used only to reduce overfitting to training data, it is typically disabled when the model is making predictions (inference time). However, within the proposed LSTM-BNN model, dropout is also active while making predictions as it is

the key component to creating the probabilistic interval predictions as described in the following subsection.

The LSTM model is trained and tuned using Bayesian optimization for each individual household independently. Once training and tuning are completed, the model is ready to proceed to the Interval Forecasting step.

4.4. Interval Forecasting

This subsection described how the trained and tuned LSTM model is used to create the prediction intervals. The approach is inspired by the works of Zhang and Mahadevan [15] and Niu and Liang [33], and as those works, it employs the BNN technique to generate intervals. However, they used the BNN technique for different applications and with different networks.

With the active dropout, the trained model makes a sufficiently large number of predictions for each sample of the dataset. Due to dropout being active, the model has a high likelihood to produce a different point prediction each time it makes a prediction even though all inputs are the same. This variation in predictions is because while the dropout is active there is a probability that any component, excluding input and output neurons, can be removed from the prediction calculation. The varying point predictions due to the use of different components in prediction calculations allow for the construction of an interval prediction as a variational approximation of Bayesian inference for the model uncertainty [15].

After multiple predictions were made for the same input samples, the mean and standard deviation for each sample is determined using the collection of point predictions the model created. Finally, the interval prediction is given as one standard deviation above and below the mean value of the point predictions for each sample.

A summary of the four steps taken in creating the interval prediction is given in the following steps:

- Step 1:** Make multiple predictions for a given input.
- Step 2:** Compute the mean and standard deviation of the predictions for each input sample.
- Step 3:** Center the interval at the mean value.
- Step 4:** Define the upper bound of the interval as one standard deviation above the mean, and the lower bound as one standard deviation below the mean.

All the models generated by Bayesian optimization are evaluated on the training and validation sets while only the best-performing model for each household selected on the validation set is evaluated on the test set. In other words, the model selection is carried out on the validation set.

4.5. Statistical Tests

Household energy consumption is dependent on the behaviors of its occupant, and as such changes when those behaviors change. We are interested to examine the effect the COVID-19 related pandemic lockdowns had on households with EV charging. As the Shapiro–Wilk test showed that the datasets do not follow a normal distribution, Mann–Whitney U test was used to examine electricity consumption with and without lockdown in order to determine if lockdowns created a change in household electricity consumption habits to the extent that it is statistically different and could impact the predictive capacity of the model.

To carry out this analysis, two datasets are considered: the first dataset, referred here as the *lockdown dataset*, contains the entire original EV household dataset including data from the lockdown as well as before lockdowns. The second dataset, *non-lockdown dataset* is a subset of the EV household dataset containing only data collected outside of the lockdown period. Both the lockdown and non-lockdown datasets go through the same preparation, preprocessing, and prediction processes for the creation of the prediction model.

As seen in Figure 1, evaluation using Mann-Whitney is carried out before the normalization is applied. The test is performed after the dataset is split into components training, validation, and testing. A comparison of the differences in the Mann-Whitney results helps us determine if there is a greater similarity between training and test sets for the lockdown dataset or for the non-lockdown dataset which is a subset of the lockdown dataset.

There are three possible outcomes for the comparison of the Mann-Whitney results for the lockdown and non-lockdown datasets. The first is that there is no significant difference between training and test sets for either the lockdown or non-lockdown datasets. The second possible outcome is that there is a greater difference between the lockdown dataset training and test sets than for the non-lockdown dataset. And the third possibility is that there is a greater difference between the non-lockdown dataset training and test sets than for the lockdown dataset.

The Mann-Whitney results are compared with the final performance of the models that are trained on the lockdown and non-lockdown datasets to observe whether there is a correlation between differences in the datasets and model predictive performance. The analysis of the Mann-Whitney results and the predictive performance of the model will improve our understanding of the conditions under which the generated models are reliable. Understanding when a model is reliable is critical for mitigating the risks of a blackout because it ensures that decisions are made based on reliable forecasting information.

5. Evaluation

This research was carried out in collaboration with London Hydro, a local electrical distribution utility for the city of London, Ontario, Canada. The real-world dataset provided by London Hydro was shared through Green Button Connect My Data (CDM), a platform for secured sharing of energy data with the consumer's consent. Through work like this, London Hydro is preparing for the increased proliferation of EVs and the corresponding increase in electricity demand. London Hydro needs home EV charging data to identify nonwire solutions such as scheduling charging during off-peak hours when there is solar generation.

In this evaluation, we consider four households with EVs and refer to them as EV1, EV2, EV3, and EV4. The time period ranges for all four households' recordings are very similar, as given in Table 2. Note that the time period ranges non-lockdown dataset are the same for all four households. The Weather Station Data was obtained from Environment and Climate Change Canada and consists of two datasets from two weather observation stations in the London area that were merged by averaging their temperature readings, as discussed in Section 4.2.

For comparison of lockdown to non-lockdown data, additional four subsets of the EV household electricity consumption datasets are created by removing the lockdown data following the start of lockdowns on 1 March 2020. For each of the four households, two individual LSTM-BNN predictor models were trained and tuned. The first model for each household is trained and evaluated on the entire dataset which contains lockdown and non-lockdown electricity consumption data, and the second model for each household is trained and evaluated using only the non-lockdown electricity consumption data.

Table 2. Data ranges for lockdown and non-lockdown datasets.

Households	Lockdown	Non-Lockdown
EV1	2018/07/22 00:00 to 2020/07/21 11:00	2018/07/22 00:00 to 2020/02/29 23:00
EV2	2018/07/22 00:00 to 2020/07/21 00:00	
EV3	2018/07/22 00:00 to 2020/07/21 10:00	
EV4	2018/07/22 00:00 to 2020/07/21 09:00	

All experiments were coded in Python with the use of the PyTorch machine learning framework and the Ray Tune library for model training and tuning. This remainder of this section consists of three subsections: first, the results of Mann-Whitney analysis

are discussed, next the hyperparameter search space is defined and training behavior is summarized for model optimization, and finally, the predictive performance is analyzed.

5.1. Statistical Test Results

Two trials are completed for each household, one for the full dataset with lockdown data and the other for the subset without lockdown data. After initial preparations are completed according to the described methodology, the dataset for non-lockdown was split into train and test sets, similar in proportions to those used for the complete dataset. The shift in behavior due to lockdowns was analyzed to determine if there was a statistically significant difference in the distribution of the electricity consumption between train and test sets for lockdown and non-lockdown conditions.

In order to interpret the results of the Mann-Whitney, the null hypothesis was established. The null hypothesis in this scenario is that there is no statistically significant difference between the training and test set for any of the datasets. For the significance level of 5%, the null hypothesis was rejected for cases where the P-value of the Mann-Whitney is less than 0.05 (5.00×10^{-2}).

The p -value results of the Mann-Whitney analysis comparing training and test datasets shown in Table 3 confirm that the null hypothesis can be rejected for all datasets, as they fall significantly below the threshold value of 5.00×10^{-2} . Therefore, the Mann-Whitney results indicate that there is a statistically significant difference between the training and test datasets regardless of lockdowns for all households.

Table 3. Mann-Whitney p -value results for the lockdown and non-lockdown datasets.

Households	p -Values Lockdown	p -Values Non-Lockdown
EV1	1.66×10^{-23}	9.48×10^{-5}
EV2	8.25×10^{-3}	9.33×10^{-19}
EV3	5.64×10^{-188}	1.07×10^{-31}
EV4	9.22×10^{-43}	3.40×10^{-8}

5.2. Model Training and Tuning

For each of the datasets outlined in Table 2, 80 models were considered using Bayesian optimization within a defined hyperparameter search space. The hyperparameters tuned for the model were batch size, window size, the number of hidden layers, the number of neurons in the hidden layers, the Adam optimizer learning rate, and the dropout probability. The defined search space for each of the hyperparameters is summarized in Table 4. Every model that was trained had its performance evaluated using the performance metrics outlined in Section 2.3.

Table 4. Hyperparameter search space.

Hyperparameter	Search Space
Batch Size	32, 64, or 128
Window Size	12, 24, 48, or 72
Hidden Layers	1, 2, or 3
Hidden Neurons	64 or 128
Learning Rate	1×10^{-4} , 1×10^{-3} , or 1×10^{-2}
Dropout Probability	5×10^{-2} , 1×10^{-1} , or 1.5×10^{-1}

The input and output layers were each set to a fixed size. The number of neurons in the input layer was set by the number of features in the input dataset, and the output layer has a single neuron for the regression prediction output. Different sliding window sizes w were used to provide the model with varying numbers of previous time steps to use as inputs. The window size is an important consideration because using a different number of previous time steps may help the models capture distinct patterns in each household's electricity consumption. Each model predicts the energy consumption one hour ahead.

The options for dropout probability tuning were based on the tuning range used by Zhang and Mahadevan [15] for creating BNN models. A dropout of zero was also included to act as a benchmark for how a non-BNN neural network would perform for each of the households. The hidden layer space and an epoch of 150 were determined by referring to the hyperparameters used for LSTM models for electricity load forecasting used by Kong et al. [17]. A sample of the training behavior of the optimal LSTM-BNN for EV3 and the non-lockdown dataset model is shown in Figure 2. From this figure, it can be seen that the early stopping could be beneficial in the reduction of computational resources, as a very minor improvement of the validation performance can be observed beyond approximately 80 epochs.

In experiments, the proposed LSTM-BNN interval forecasting is compared to the point forecasting. Both use LSTM as their base model and both undergo exactly the same dataset preparation, preprocessing, and training and tuning steps, as described in Sections 4.1, 4.2 and 4.3, respectively. The difference is that in point forecasting, at inference time, the dropout is not active, and, therefore, point forecasting results in a single precision for each time step. In contrast, the proposed LSTM-BNN generates multiple predictions and forms an interval with the BNN technique.

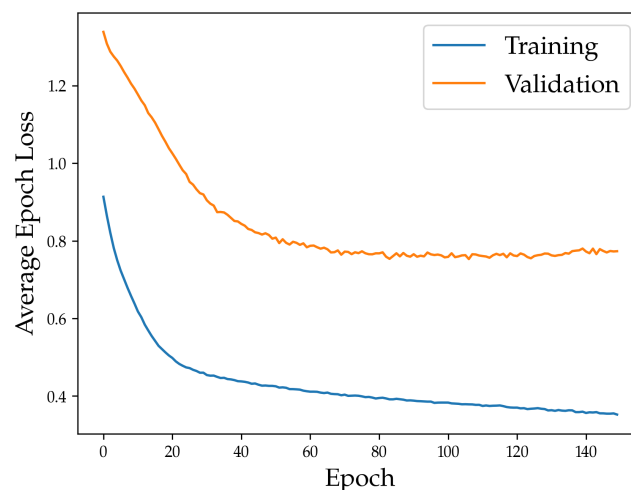


Figure 2. Training and validation loss of the optimal LSTM-BNN EV3 non-lockdown model.

5.3. Predictive Performance Analysis

The performance results are explored in three parts: first, the overall performance of point and interval prediction models are examined, followed by the analysis of the performance among households. Next, interval forecasts are compared to point forecasts, and lockdown is compared to non-lockdown. Finally, the correlation between Mann-Whitney results and model performance is examined.

5.3.1. Overall Performance

Table 5 shows the average MAPE values for the four households for each of the two datasets, lockdown and non-lockdown, and for each of the two approaches, point and interval prediction approaches. For interval forecasts, the four performance metrics were calculated using the mean of the generated interval from the set of the forecast generated with the Bayesian technique as described in Section 4.4. For each household, dataset, and point/interval approach, the model was tuned, and the results from the tuned models were averaged and reported in this table. All MAPE values are significantly higher than those reported in the literature [12,18], but that is to be expected as EV charging behavior adds remarkable variability and randomness to power consumption pattern compared with office buildings or households without EVs. In general, excluding electricity consumption lockdown data from March 2020 or later did not create an improvement in the predictive performance of the models. While an increase in the error between actual and predicted

values is expected between training and testing, there is a much greater increase for the non-lockdown dataset than for the lockdown dataset.

The average interval prediction performance shows that model performance was better overall on the full dataset that included lockdown data. Point predictions for the lockdown were also better than for the non-lockdown for the test dataset. This result is somewhat surprising considering that it would be expected that electricity consumption would be more difficult to predict in a lockdown environment than in a non-lockdown environment when the historical data used for creating models are based mainly on non-lockdown behavior. A possible reason for this is that with the full dataset, the model has more data to learn from.

Table 5. Average model performance comparison.

Dataset	Model Type	Train MAPE	Validation MAPE	Test MAPE
Lockdown	Point	44.8121	62.2463	63.4173
	Interval	37.0744	59.8753	68.7626
Non-Lockdown	Point	37.2093	61.1446	71.5252
	Interval	44.1562	68.2564	122.3360

Figure 3 shows an example of interval forecast results; specifically actual energy consumption and predicted interval forecasts for EV3 and the lockdown dataset. It can be observed that the prediction interval varies throughout time indicating uncertainties in the forecasted values.

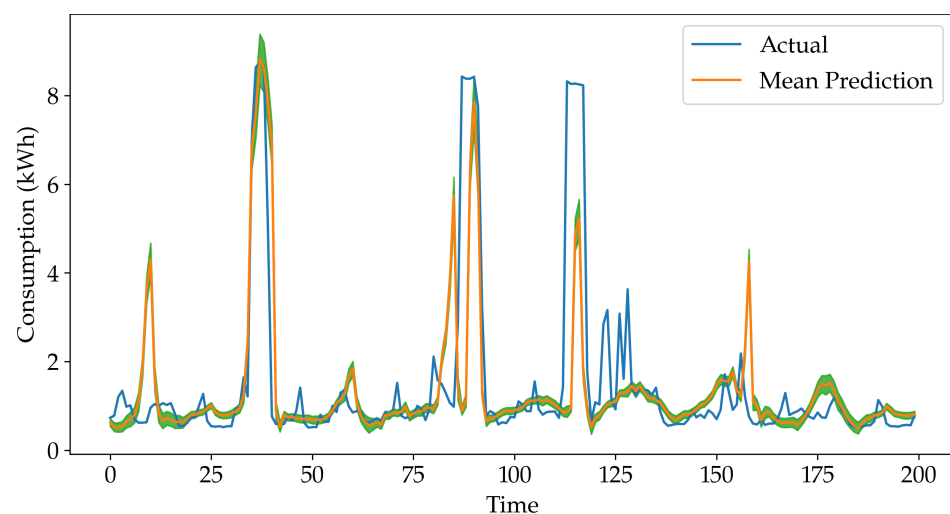


Figure 3. The first 200 predictions made on the validation set of the optimized EV3 model for the lockdown dataset. The green areas indicate one standard deviation above and one standard deviation below the mean prediction value.

While the focus of this work is the comparison of interval and point predictions, here we further examine interval forecast for the example EV3. Prediction Interval Coverage Probability (*PICP*) measures the fraction of actual values that lay within the prediction interval. In the case of the proposed LSTM-BNN, this will vary depending on how many standard deviations σ are forming the interval. For the interval of $\pm 1\sigma$, *PICP* is only 0.28, while increasing to $\pm 2\sigma$, $\pm 3\sigma$, and $\pm 5\sigma$, increases *PICP* to 0.51, 0.68, 0.82, respectively. Note that similar to MAPE values from Table 5, these *PICP* values indicate high errors caused by the randomness of EV charging.

5.3.2. Performance among Households

The results for the best-performing model for each individual household, for point and interval forecasts, are shown in Tables 6 and 7 for lockdown and non-lockdown datasets

respectively. Note that models are created for each household individually, and not for a group of homes. The results include all four metrics MAPE, MSE, RMSE, and MAE for all datasets train, validation, and test. The best-performing point and interval prediction models in terms of test set MAPE were for households EV3 and EV4 respectively on the lockdown dataset. Although there are some cases where the models generate better predictions, such as test MAPE of about 31% for EV3 with lockdown dataset and point predictions, most others exhibit higher error. Despite the success of LSTM models in forecasting electricity consumption for offices and households [12,36], the results show that in the presence of EV charging their accuracy is greatly reduced. As noted by other studies [12,17] the variability in household electricity consumption makes creating accurate predictions at this granularity challenging, which seems to only be exacerbated further with the additional consideration of EV charging. This variability results in much higher MAPE, irrelevant of the dataset or the household, than those typically reported in the literature. However, the literature commonly considers offices, schools, or a group of buildings that have much more predictable energy consumption patterns. Moreover, MAPE can produce very high values when the actual values are close to zero. Note that, while MSE, RMSE, and MAE are included in Tables 6 and 7, we are not comparing among households based on those metrics as they are dependent on the scale of the actual values.

Table 6. Lockdown results.

Metric	Model Type	Household							
		EV1		EV2		EV3		EV4	
		Point	Interval	Point	Interval	Point	Interval	Point	Interval
MAPE	Train	70.0221	35.7423	43.0054	56.4060	26.7295	25.0654	39.4916	31.0839
	Validation	78.0701	80.1722	92.6025	87.0372	36.3606	37.5774	41.9519	34.7142
	Test	47.0062	70.9079	122.2161	103.2104	31.5572	57.0495	52.8898	43.8828
MSE	Train	0.3033	0.0470	0.2293	0.3027	0.3726	0.3395	0.7405	0.3801
	Validation	0.2600	0.2528	0.8229	0.7407	1.3065	1.4637	1.1385	1.2946
	Test	0.4742	0.4492	2.0359	1.8019	1.5855	4.2337	1.8502	1.9896
RMSE	Train	0.5507	0.2168	0.4789	0.5502	0.6104	0.5827	0.8605	0.6165
	Validation	0.5099	0.5028	0.9071	0.8606	1.1430	1.2098	1.0670	1.1378
	Test	0.6886	0.6703	1.4269	1.3423	1.2592	2.0576	1.3602	1.4105
MAE	Train	0.2687	0.1190	0.2870	0.3348	0.3281	0.3152	0.4771	0.3519
	Validation	0.2089	0.2121	0.5672	0.5249	0.4907	0.5043	0.5545	0.4847
	Test	0.3617	0.3758	0.7978	0.7444	0.6656	1.1306	0.7871	0.6899

Table 7. Non-lockdown results.

Metric	Model Type	Household							
		EV1		EV2		EV3		EV4	
		Point	Interval	Point	Interval	Point	Interval	Point	Interval
MAPE	Train	71.9591	71.3261	26.4764	35.0559	32.2495	33.4092	18.1520	36.8335
	Validation	79.8842	116.1025	79.5442	81.3031	39.4762	35.7723	45.6739	39.8479
	Test	84.0878	266.6882	62.2114	72.7572	65.6180	94.3859	74.1836	55.5127
MSE	Train	0.9714	0.4387	0.0768	0.1281	0.4747	0.8219	0.0845	0.8166
	Validation	0.3199	0.4667	1.1516	1.2964	2.3181	1.8993	2.0753	1.7791
	Test	0.5117	0.5249	0.8616	0.8402	2.7516	3.1830	2.3104	1.2918
RMSE	Train	0.9856	0.6623	0.2771	0.3580	0.6890	0.9066	0.2908	0.9037
	Validation	0.5656	0.6832	1.0731	1.1386	1.5225	1.3781	1.4406	1.3338
	Test	0.7153	0.7245	0.9282	0.9166	1.6588	1.7841	1.5200	1.1366
MAE	Train	0.4498	0.3085	0.1641	0.2026	0.3912	0.4663	0.1792	0.4672
	Validation	0.2312	0.3241	0.6561	0.6889	0.6265	0.5881	0.7280	0.6617
	Test	0.2681	0.4952	0.5562	0.5458	0.7869	0.9827	0.8580	0.6530

To analyze household differences, Figure 4a shows the MAPE values for point and interval predictions for each household for lockdown and lockdown data while Figure 4b does the same for the non-lockdown scenario. It can be observed that there are some

households that were easier for the predictive models to capture. The EV3 and EV4 household datasets produced the best-performing models but the models trained on the EV1 and EV2 households were much less successful. The predictive performance of the model trained on the EV2 lockdown dataset which produces the highest p -value highlights that there is no direct relationship between Mann-Whitney results and model performance. The Mann-Whitney results from Table 3 indicate that in all scenarios, there is a significant difference between the training and test dataset which could be one of the reasons for the high MAPE results.

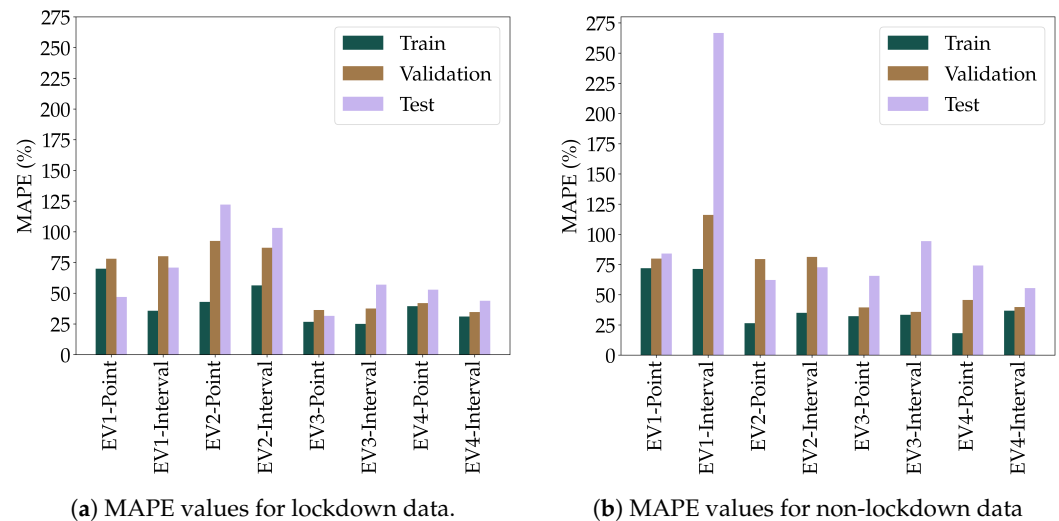


Figure 4. MAPE values for lockdown and non-lockdown data.

To analyze this data from a different perspective, Figure 5 shows the MAE values for the four EVs for point and interval prediction in lockdown and non-lockdown periods. Since MAE is scale dependent and consumption scales vary among households, different approaches should be compared for each household individually and not among households. For four scenarios, point forecasting achieves lower error than interval forecasting; however, interval forecasting has the advantage of providing uncertainty information.

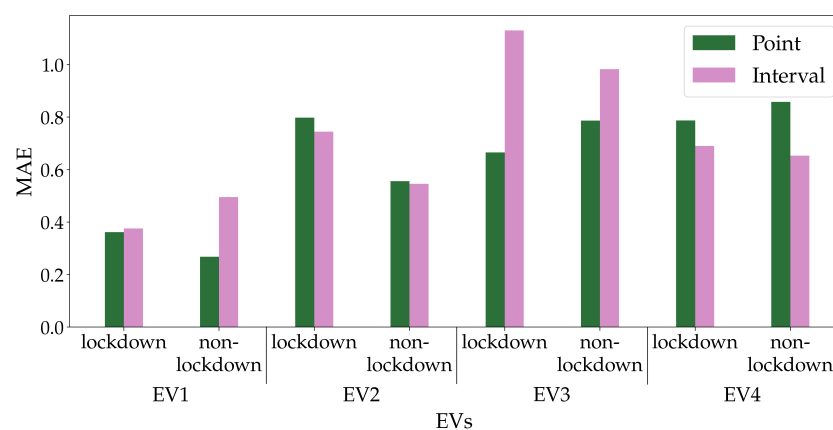


Figure 5. MAE values for lockdown and non-lockdown data.

5.3.3. Interval Forecasts Compared to Point Forecasts

Figure 6a compares interval forecasts to point forecasts in terms of MAPE for the lockdown test dataset. In addition to point and interval MAPE for each household, it includes the averages for the four households. It can be observed that the average MAPE for the point forecast is about 5% lower than the average MAPE for interval forecasts. Also, this figure highlights that EV2 was the hardest household to predict for both point and

interval approaches. At the same time, EV3 was the most straightforward prediction for point forecasting while EV4 achieved the lowest prediction error for interval forecasting. Figure 6b shows the same comparison but for non-lockdown test data. Again, MAPE for interval forecasting is higher than for point forecasting but the difference between average MAPE for interval and point forecasts is much larger for the non-lockdown dataset than for the lockdown dataset. Interval forecasting for EV1 performed especially poorly which raised the average MAPE for interval forecasts.

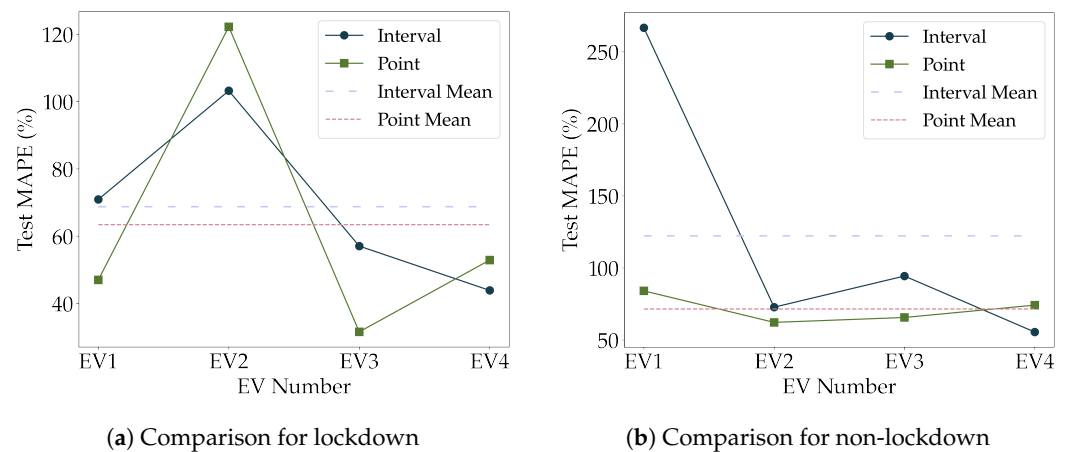


Figure 6. Comparison between interval and point forecasts for lockdown and non-lockdown data.

5.3.4. Lockdown Compared to Non-lockdown

Figure 7a,b compare forecasting with lockdown and non-lockdown for point and interval forecasting respectively. Figure 7a shows that average point MAPE for non-lockdown data is about 8% higher than for lockdown. All lockdown MAPE values except EV2 were much lower than non-lockdown predictions. Similarly, Figure 7b indicates that interval prediction for lockdown achieves better results than for non-lockdown except EV2. Overall, accuracy varies greatly among households.

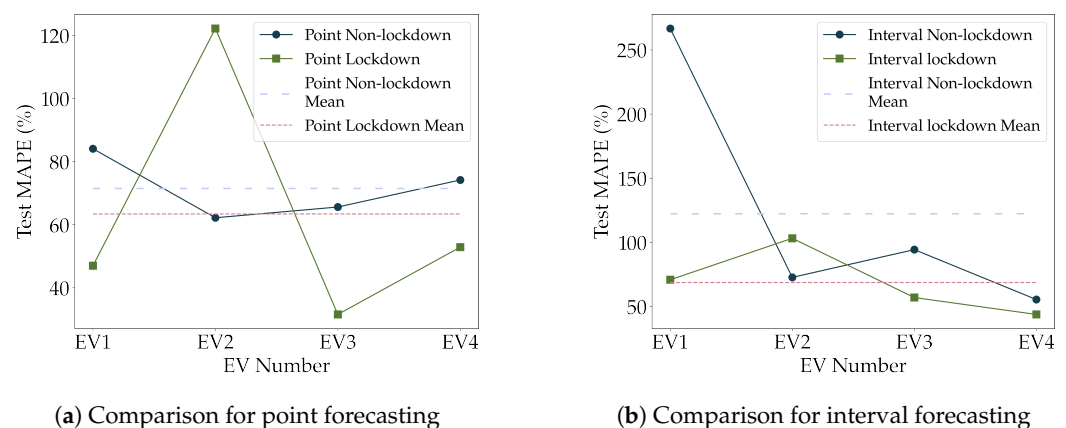


Figure 7. Comparison between lockdown and non-lockdown results for point and interval forecasts.

5.3.5. Analysis of Relation between Predictive Performance and Mann-Whitney Results

Figure 8a,b relate model performance on the test set (MAPE) and Mann-Whitney p -value between training and test data. These figures further demonstrate that there is no conclusive relationship between the Mann-Whitney p -value and model performance on the test set. This is due to Mann-Whitney comparing the datasets statistically while the prediction performance depends on many other factors such as the quality of features and randomness of data. Still, the Mann-Whitney test demonstrated that there is a significant

difference between the train and test datasets, regardless of whether lockdown data were included or not, which is a contributing factor to the higher error values observed.

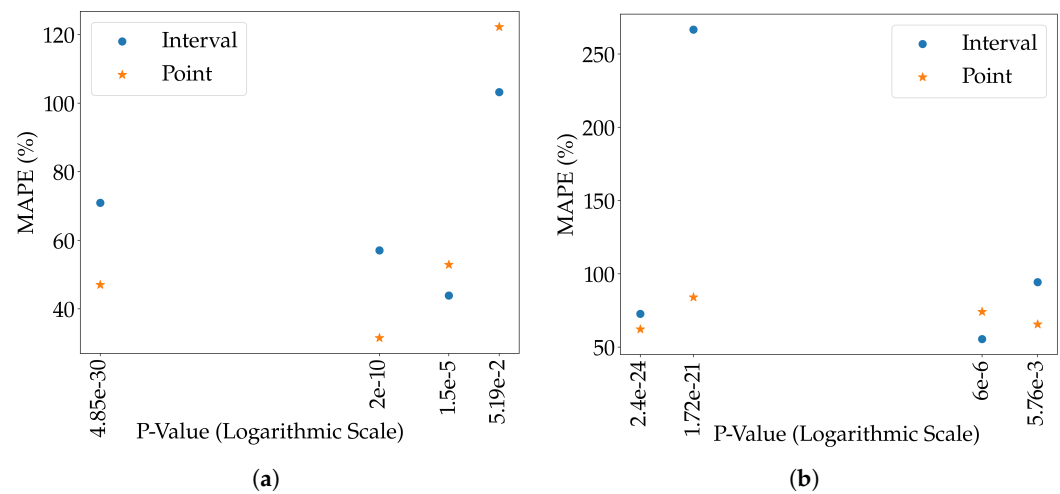


Figure 8. MAPE performance compared to Mann-Whitney p -Value significance measure. (a) Lockdown MAPE performance compared to Mann-Whitney p -Value significance measure; (b) Non-lockdown MAPE performance compared to Mann-Whitney p -Value significance measure.

6. Conclusions

Accurately predicting electricity consumption is an important factor in providing an adequate and reliable energy supply. The electricity demand will continue to grow as society transitions away from using ICE vehicles to EVs that are able to be charged in residential homes. However, predicting energy consumption at the individual household level is more challenging than forecasting for office buildings, schools, or regions due to the high variability in electricity consumption patterns [19,37]. This challenge is further amplified by the need to accommodate EV charging.

This paper proposes LSTM-BNN interval load forecasting for individual households in presence of EV charging based on the LSTM deep learning model and Bayesian inference. The LSTM model incorporates the dropout layer which is active during the inference time and responsible for generating a set of point predictions for a single input sample. Then, the Bayesian technique is employed to create interval forecasts from this set of predictions. The achieved accuracy varies greatly among households due to the variability and randomness of their energy consumption patterns. Examining the performance of point and interval prediction models shows that the LSTM-BNN interval prediction model performs similarly to a standard LSTM point prediction model with the benefit of providing an interval for the prediction. Although the proposed LSTM-BNN is more complex and involves longer training time than the traditional LSTM point forecasting models, LSTM-BNN predictions quantify uncertainty and offer additional information for decision-making. This paper also examined the impact of the COVID-19 lockdown on the load forecasting for these households: results show that the proposed LSTM-BNN achieves similar results for lockdown and non-lockdown periods. We stipulate that the randomness of the EV charging patterns outweighs the impact of change due to the lockdowns.

As demonstrated in our study, EV charging is highly variable and predicting household energy consumption in presence of EV charging is difficult. For use cases such as infrastructure planning, forecasting energy consumption for a neighborhood block may be sufficient. For such scenarios, aggregating energy consumption on the block level would remove some of the randomnesses and improve forecasting accuracy. Future work will examine the results in terms of the size of the prediction interval to be able to better relate different interval forecasts. Moreover, alternative methods to Mann-Whitney will be considered to acquire better insight into the potential changes in consumption habits during different periods of time. As energy consumption patterns, including EV charging patterns,

change over time resulting in what is known as concept drift, the techniques such as online learning [18] could be integrated with the proposed approach to better capture changes over time.

Author Contributions: Conceptualization, R.S., S.M.; methodology, R.S., K.G.; software, R.S.; validation, R.S.; formal analysis, R.S., M.A.T.A.E., K.G.; investigation, R.S., M.A.T.A.E., K.G.; resources, K.G. and S.M.; writing—original draft preparation, R.S.; writing—review and editing, K.G., M.A.T.A.E., S.M.; visualization, R.S., M.A.T.A.E.; supervision, K.G.; project administration, K.G.; funding acquisition, K.G., S.M. All authors have read and agreed to the published version of the manuscript.

Funding: This research has been supported by Ontario Centre of Innovation under grant OCI #34674 and by the Natural Sciences and Engineering Research Council of Canada (NSERC) under grant ALLRP 570760-21. Computation was enabled in part by the Digital Research Alliance of Canada.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Data analyzed in this study are obtained from London Hydro and are protected under a signed non-disclosure agreement. Approvals are needed for sharing this data.

Acknowledgments: The authors would like to thank London Hydro for supplying industry knowledge and data used in this study.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Zhao, H.; Guo, S. External Benefit Evaluation of Renewable Energy Power in China for Sustainability. *Sustainability* **2015**, *7*, 4783–4805. [CrossRef]
2. Grid Modernization and the Smart Grid. Available online: <https://www.energy.gov/oe/grid-modernization-and-smart-grid> (accessed on 13 August 2022).
3. Alternative Fuels Data Center: Emissions from Electric Vehicles. Available online: https://afdc.energy.gov/vehicles/electric_emissions.html (accessed on 13 August 2022).
4. Ghunem, R. Smarter, Faster and Smaller Power Grids: A Step towards a Green Economy. 2022. Available online: <https://nrc.canada.ca/en/stories/smarter-faster-smaller-power-grids-step-towards-green-economy> (accessed on 13 August 2022).
5. Yamashita, K.; Joo, S.K.; Li, J.; Zhang, P.; Liu, C.C. Analysis, control, and economic impact assessment of major blackout events. *Eur. Trans. Electr. Power* **2008**, *18*, 854–871. [CrossRef]
6. Ozcan, A.; Catal, C.; Kasif, A. Energy Load Forecasting Using a Dual-Stage Attention-Based Recurrent Neural Network. *Sensors* **2021**, *21*, 7115. [CrossRef]
7. Sehovac, L.; Grolinger, K. Deep learning for load forecasting: Sequence to sequence recurrent neural networks with attention. *IEEE Access* **2020**, *8*, 36411–36426. [CrossRef]
8. Sun, L.; Qin, H.; Przystupa, K.; Majka, M.; Kochan, O. Individualized Short-Term Electric Load Forecasting Using Data-Driven Meta-Heuristic Method Based on LSTM Network. *Sensors* **2022**, *22*, 7900. [CrossRef] [PubMed]
9. Jung, S.; Moon, J.; Park, S.; Hwang, E. An attention-based multilayer GRU model for multistep-ahead short-term load forecasting. *Sensors* **2021**, *21*, 1639. [CrossRef]
10. Al-Ogaili, A.S.; Hashim, T.J.T.; Rahmat, N.A.; Ramasamy, A.K.; Marsadek, M.B.; Faisal, M.; Hannan, M.A. Review on scheduling, clustering, and forecasting strategies for controlling electric vehicle charging: Challenges and recommendations. *IEEE Access* **2019**, *7*, 128353–128371. [CrossRef]
11. Yu, Y.; Si, X.; Hu, C.; Zhang, J. A review of recurrent neural networks: LSTM cells and network architectures. *Neural Comput.* **2019**, *31*, 1235–1270. [CrossRef]
12. Fekri, M.N.; Grolinger, K.; Mir, S. Distributed load forecasting using smart meter data: Federated learning with Recurrent Neural Networks. *Int. J. Electr. Power Energy Syst.* **2022**, *137*, 107669. [CrossRef]
13. Jagait, R.K.; Fekri, M.N.; Grolinger, K.; Mir, S. Load forecasting under concept drift: Online ensemble learning with recurrent neural network and ARIMA. *IEEE Access* **2021**, *9*, 98992–99008. [CrossRef]
14. Hastie, T.; Tibshirani, R.; Friedman, J.H.; Friedman, J.H. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*; Springer: Berlin/Heidelberg, Germany, 2009.
15. Zhang, X.; Mahadevan, S. Bayesian neural networks for flight trajectory prediction and safety assessment. *Decis. Support Syst.* **2020**, *131*, 113246. [CrossRef]
16. Wan, C.; Xu, Z.; Pinson, P.; Dong, Z.Y.; Wong, K.P. Probabilistic forecasting of wind power generation using extreme learning machine. *IEEE Trans. Power Syst.* **2013**, *29*, 1033–1044. [CrossRef]

17. Kong, W.; Dong, Z.Y.; Jia, Y.; Hill, D.J.; Xu, Y.; Zhang, Y. Short-term residential load forecasting based on LSTM recurrent neural network. *IEEE Trans. Smart Grid* **2017**, *10*, 841–851. [\[CrossRef\]](#)
18. Fekri, M.N.; Patel, H.; Grolinger, K.; Sharma, V. Deep learning for load forecasting with smart meter data: Online Adaptive Recurrent Neural Network. *Appl. Energy* **2021**, *282*, 116177. [\[CrossRef\]](#)
19. Zhang, X.M.; Grolinger, K.; Capretz, M.A.; Seewald, L. Forecasting residential energy consumption: Single household perspective. In Proceedings of the 2018 17th IEEE International Conference on Machine Learning and Applications, Orlando, FL, USA, 17–20 December 2018; pp. 110–117.
20. L'Heureux, A.; Grolinger, K.; Capretz, M.A. Transformer-Based Model for Electrical Load Forecasting. *Energies* **2022**, *15*, 4993. [\[CrossRef\]](#)
21. Tan, M.; Hu, C.; Chen, J.; Wang, L.; Li, Z. Multi-node load forecasting based on multi-task learning with modal feature extraction. *Eng. Appl. Artif. Intell.* **2022**, *112*, 104856. [\[CrossRef\]](#)
22. Ribeiro, A.M.N.; do Carmo, P.R.X.; Endo, P.T.; Rosati, P.; Lynn, T. Short-and very short-term firm-level load forecasting for warehouses: a comparison of machine learning and deep learning models. *Energies* **2022**, *15*, 750. [\[CrossRef\]](#)
23. Jiang, Y.; Gao, T.; Dai, Y.; Si, R.; Hao, J.; Zhang, J.; Gao, D.W. Very short-term residential load forecasting based on deep-autoformer. *Appl. Energy* **2022**, *328*, 120120. [\[CrossRef\]](#)
24. Amini, M.H.; Kargarian, A.; Karabasoglu, O. ARIMA-based decoupled time series forecasting of electric vehicle charging demand for stochastic power system operation. *Electr. Power Syst. Res.* **2016**, *140*, 378–390. [\[CrossRef\]](#)
25. Yi, Z.; Liu, X.C.; Wei, R.; Chen, X.; Dai, J. Electric vehicle charging demand forecasting using deep learning model. *J. Intell. Transp. Syst.* **2022**, *26*, 690–703. [\[CrossRef\]](#)
26. Koohfar, S.; Woldemariam, W.; Kumar, A. Prediction of Electric Vehicles Charging Demand: A Transformer-Based Deep Learning Approach. *Sustainability* **2023**, *15*, 2105. [\[CrossRef\]](#)
27. Aduama, P.; Zhang, Z.; Al-Sumaiti, A.S. Multi-Feature Data Fusion-Based Load Forecasting of Electric Vehicle Charging Stations Using a Deep Learning Model. *Energies* **2023**, *16*, 1309. [\[CrossRef\]](#)
28. Zheng, Y.; Shao, Z.; Zhang, Y.; Jian, L. A systematic methodology for mid-and-long term electric vehicle charging load forecasting: The case study of Shenzhen, China. *Sustain. Cities Soc.* **2020**, *56*, 102084. [\[CrossRef\]](#)
29. Arias, M.B.; Bae, S. Electric vehicle charging demand forecasting model based on big data technologies. *Appl. Energy* **2016**, *183*, 327–339. [\[CrossRef\]](#)
30. Maciejowska, K.; Nowotarski, J.; Weron, R. Probabilistic forecasting of electricity spot prices using Factor Quantile Regression Averaging. *Int. J. Forecast.* **2016**, *32*, 957–965. [\[CrossRef\]](#)
31. Shi, Z.; Liang, H.; Dinavahi, V. Direct interval forecast of uncertain wind power based on recurrent neural networks. *IEEE Trans. Sustain. Energy* **2017**, *9*, 1177–1187. [\[CrossRef\]](#)
32. Kabir, H.M.D.; Khosravi, A.; Kavousi-Fard, A.; Nahavandi, S.; Srinivasan, D. Optimal uncertainty-guided neural network training. *Appl. Soft Comput.* **2021**, *99*, 106878. [\[CrossRef\]](#)
33. Niu, Z.; Liang, H. Nuclear mass predictions based on Bayesian neural network approach with pairing and shell effects. *Phys. Lett. B* **2018**, *778*, 48–53. [\[CrossRef\]](#)
34. Mirasgedis, S.; Sarafidis, Y.; Georgopoulou, E.; Lalas, D.; Moschovits, M.; Karagiannis, F.; Papakonstantinou, D. Models for mid-term electricity demand forecasting incorporating weather influences. *Energy* **2006**, *31*, 208–227. [\[CrossRef\]](#)
35. Falkner, S.; Klein, A.; Hutter, F. BOHB: Robust and Efficient Hyperparameter Optimization at Scale. In Proceedings of the 35th International Conference on Machine Learning, Stockholm, Sweden, 10–15 July 2018; pp. 1437–1446.
36. Wang, X.; Fang, F.; Zhang, X.; Liu, Y.; Wei, L.; Shi, Y. LSTM-based Short-term Load Forecasting for Building Electricity Consumption. In Proceedings of the IEEE 28th International Symposium on Industrial Electronics, Vancouver, BC, Canada, 12–14 June 2019; pp. 1418–1423.
37. Grolinger, K.; L'Heureux, A.; Capretz, M.A.; Seewald, L. Energy forecasting for event venues: Big data and prediction accuracy. *Energy Build.* **2016**, *112*, 222–233. [\[CrossRef\]](#)

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.