We got the power: predicting available capacity for vehicle-to-grid services using a deep recurrent neural network

Rob Shipman*, Department of Architecture & Built Environment, Faculty of Engineering, University Park, University of Nottingham, NG7 2RD, UK. rob.shipman@nottingham.ac.uk

Rebecca Roberts, Kearney Limited, Adelphi, 12th Floor, 1-11 John Adam Street, London, WC2N 6HT, UK. rebecca.roberts@kearney.com

Julie Waldron, Department of Architecture & Built Environment, Faculty of Engineering, University Park, University of Nottingham, NG7 2RD, UK. julie.waldron@nottingham.ac.uk

Sophie Naylor, Department of Architecture & Built Environment, Faculty of Engineering, University Park, University of Nottingham, NG7 2RD, UK. sophie.naylor@nottingham.ac.uk

James Pinchin, Horizon Digital Economy Research, Nottingham Geospatial Building, University of Nottingham Innovation Park, NG7 2RD, UK. james.pinchin@nottingham.ac.uk

Lucelia Rodrigues, Department of Architecture & Built Environment, Faculty of Engineering, University Park, University of Nottingham, NG7 2RD, UK. lucelia.rodrigues@nottingham.ac.uk

Mark Gillott, Department of Architecture & Built Environment, Faculty of Engineering, University Park, University of Nottingham, NG7 2RD, UK. mark.gillott@nottingham.ac.uk

* Corresponding Author

Abstract

Vehicle-to-grid (V2G) services utilise a population of electric vehicle batteries to provide the aggregated capacity required to participate in power and energy markets. Such participation relies on the prediction of available capacity to support the reliable delivery of agreed reserves at a future time. In this work real historical trip data from a fleet of vehicles belonging to the University of Nottingham was used and a simulation developed to show how battery state-of-charge and available capacity would vary if these trips were taken in electric vehicles that were charged at simulated charging station locations. A time series forecasting neural network was developed to predict aggregated available capacity for the next 24-hour period given input data from the previous 24 hours and its increased predictive capability over a regression model trained using automated machine learning was demonstrated. The simulations were then extended to include delivery of reserves to satisfy the needs of simulated market events and the ability of the model to successfully adapt its predictions to such events was demonstrated. The authors conclude that this ability is of critical importance to the viability and success of future V2G services by supporting trading and vehicle utilisation decisions for multiple market events.

Keywords

Vehicle-to-grid; V2G; deep learning; CNN-LSTM network; machine learning; neural networks.

1. Introduction

Vehicle-to-grid (V2G) is an emerging technology that allows the aggregation of a population of distributed energy stores (DES) in the form of electric vehicle batteries to help decarbonise transport, integrate renewable energy into the grid and participate in revenue-generating markets [1–6]. V2G services aim to capitalise on these market opportunities by optimising the charging and discharging schedules of vehicles while respecting the requirement for the vehicles to continue to fulfil their requirements as a means of transport [7,8]. The target markets are often complex and vary from country to country, however, they can typically be broadly classified into three types: the capacity market, the energy market and the balancing services market.

The capacity market is a long-term instrument used by the system operator to ensure that electricity supply can continue to meet demand and can cope with times of system stress. Agreements are often made several years in advance of the delivery date and regular payments may be made to participants to support required investments [9]. Delivery of the agreed capacity however may be required at short notice, for example four hours ahead of the required period of delivery. In the energy market, generators sell electricity to suppliers for short and specific periods of time, called settlement periods that are 30-minutes long in the UK market for example. Trading may take place well in advance of the settlement period; however, the majority often takes place in the day-ahead market with further refinement of trading positions in the intraday market up to a cut-off point, called gate closure, which may run up to one hour before the settlement period [10]. The balancing or ancillary services market has the role of balancing supply and demand in real-time, building on the foundation of the other markets. Numerous mechanisms are employed to achieve this end including: firm frequency response (FFR), which requires participants to respond to a fluctuation in supply frequency by providing the agreed power within several seconds; demand side

response (DSR) which requires participants to turn up, turn down or shift demand in real-time; and reserve services that require participants to provide or consume additional power at short notice [11].

Such market opportunities have the potential to generate substantial revenue for a V2G service that could be used in-part to incentivise participants. However, overestimating available capacity and therefore entering into trading agreements that cannot be fulfilled would incur significant financial penalties. Equally, underestimating available capacity and therefore missing viable opportunities would reduce the revenue potential of the service. The viability and success of a V2G service is therefore dependent on prediction of available capacity from the fleet of vehicles. This work focusses on the near-term predictions that are required to satisfy the needs of energy trading in the dominant day-ahead and intraday markets and to assess the impact of delivery to the capacity markets and certain ancillary services that may be called upon at short notice.

This is a problem that has been addressed previously in the literature. In [12], an aggregation of EVs was modelled using a queueing network to allow available capacities for both charge and discharge to be estimated however required charge and discharge durations to obey specific distributions that may not have been feasible in a real-life scenario. An aggregate model was also introduced in [13], which relaxed this constraint and was shown to produce accurate predictions for day-ahead, hour-ahead and real-time horizons using simulated charging demand scenarios. However, the authors concluded that future work should consider the uncertain reserve utilization introduced by the V2G fleet being called upon to deliver its reserves during real-time operation.

In [14] it was noted that a live V2G system would require the real-time processing and analysis of very large datasets and therefore available capacity prediction should be treated as a big data

analysis problem for which traditional techniques are not well suited. The parallel gradient boosted decision tree was used to predict available capacity using real data from over 500 EVs for timescales up to 24 hours and compared to several other machine learning techniques with some success in terms of both accuracy and training time. In [15] it was noted that the prediction model must consider not only regular behaviour but short-term uncertainty as well. The long short-term memory (LSTM) neural network was introduced as a deep learning technique well suited to learning such patterns in time series data. It was trained to accurately predict available capacity for individual vehicles over a 4.5-hour prediction horizon using real data to support potential participation in the FFR market. However, the impact of the delivery of available reserves during real-time operation on the prediction and longer prediction horizons were not considered.

An extension of the LSTM is an encoder-decoder model in which a convolutional neural network (CNN) is used as the encoder and a LSTM network as the decoder. Although initially developed for natural language processing [16] and visual recognition tasks [17], such CNN-LSTM networks have been used with considerable success for time series forecasting problems in the literature. For example, residential energy consumption prediction [18], travel time prediction [19] and a similar network for precipitation nowcasting [20]. In this work, a CNN-LSTM model was developed to predict the aggregated available capacity for the next 24-hour period from a fleet of vehicles belonging to the University of Nottingham. Real historical trip data was used as input to a simulation that generated variations in available capacity as if these trips were taken in electric vehicles charged at simulated charging station locations. The execution of market trades was also simulated to investigate the impact of the delivery of available reserves on the predictions made by the model, which has received little attention in the literature to date.

The specific contributions made in the paper are as follows; it was demonstrated that the CNN-LSTM model was able to generate accurate near-term aggregate capacity forecasts for the fleet, which is of value in informing trading and vehicle utilisation decisions made by the V2G service; the ability of the model to adapt to the impact of energy exported by the V2G service was demonstrated, which is important in supporting multiple trading events within a short timeframe and the importance of a multivariate input and the "signposting" of upcoming market events to achieve this was highlighted.

The remainder of the paper is structured as follows; in the following section the methods are described including development of the dataset, state-of-charge simulations and the architecture of the model. In the following section, the results with and without simulated market events are presented, which are then discussed before the paper concludes.

2. Methods

The research framework used in this work is shown in Figure 1. Historical data from the University of Nottingham's (UoN) fleet of vehicles was used as input together with the simulated locations of Electric Vehicle Supply Equipment (EVSE). This data was used to progressively build a dataset that simulated aggregate available capacity from the fleet that was used as input to 3 different predictive models; a naïve persistence model, an automated machine learning (AutoML) regression model and a CNN-LSTM time series forecasting model. The operation of the V2G service itself was also simulated through market events during which energy was exported from the vehicles. The time series forecasting model alone was used to predict available capacity in this scenario.

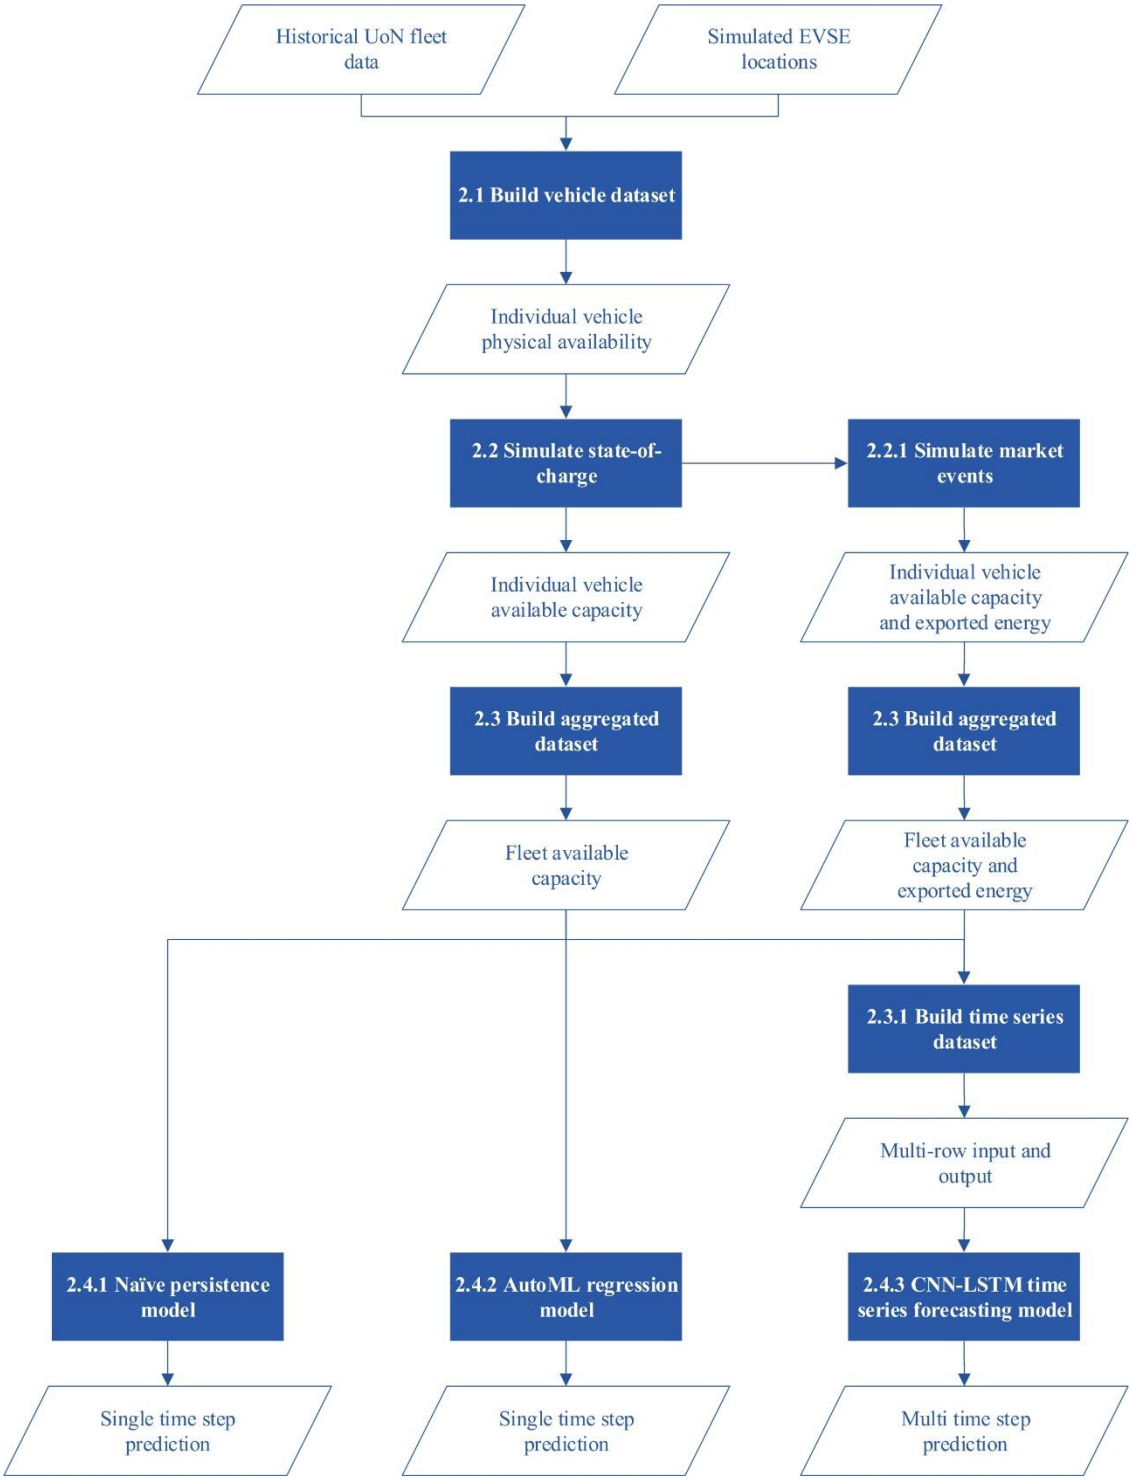

Figure 1: Research framework used in this work; each process is described in the sub-section indicated.

2.1. Vehicle Dataset

53 weeks of historical data from a fleet of 48 vehicles belonging to 6 different departments of the University of Nottingham was analysed. Each of these vehicles was equipped with Trakm8 telematics equipment [21] that provided details of each trip made by the vehicle including the time and GPS location at the start and end of the trip together with the distance travelled. Almost 190,000 trips were completed during the 53-week period. The mean number of trips for the 48 vehicles was 4,211 (s = 2,721) with a maximum of 13,268 and a minimum of 464. An anonymised excerpt is shown in Table 1.

Table 1: Anonymised excerpt from the vehicle trip dataset, where v is the vehicle ID

v	start time	start latitude	start longitude	end time	end latitude (end_lat_v)	end longitude (end_lng_v)	distance (m)
1	2019-09-06 09:55:27	51.93866	-1.17032	2019-09-06 10:01:34	51.94276	-1.17109	1109
1	2019-09-06 10:10:27	51.94274	-1.17107	2019-09-06 10:28:25	51.93907	-1.15245	4411
20	2019-10-01 08:59:08	52.84003	-1.09277	2019-10-01 09:35:28	52.85763	-1.04591	8772
20	2019-10-01 09:39:58	52.85762	-1.0459	2019-10-01 10:01:45	52.84005	-1.09298	6461

The Haversine formula, shown in Equation 1, was used to calculate the great-circle distance between the location at the end of each trip with the simulated locations of six V2G EVSE installations on three university campuses within the city of Nottingham, UK. These locations were chosen following analysis of typical parked locations, assessment of infrastructure requirements and interviews with fleet managers [22]. Where the vehicle was parked within a 100m radius of an EVSE it was defined to be close enough to be connected to the charge point. A radius of 100m was

chosen to account for small variances in GPS locations and as a suitable distance within which a connection could be made with only minor changes in behaviour e.g. parking in a different space within the same car park. The 53-week period was divided into consecutive half-hour periods. If a vehicle was stationary within 100m of an EVSE for a full half-hour period, it was assumed that it would be connected and hence available to a V2G aggregation service for charge or discharge i.e. $a_v = 1$.

Equation 1: Haversine formula to calculate great-circle distance, $dist_i$, in km between the location of a parked vehicle v (end_lat_v and end_lng_v) and EVSE i, ($evse\_lat_i$ and $evse\_lng_i$) where r is the radius of the earth (6,371km) (a); vehicle availability, a_v, based on the distance to the closest EVSE (b).

$$(a)\ dist_i = 2.r.\arcsin\left(\sqrt{\sin^2\left(\frac{evse\_lat_i - end\_lat_v}{2}\right) + \cos(end\_lat_v).\cos(evse\_lat_i).\sin^2\left(\frac{evse\_lng_i - end\_lng_v}{2}\right)}\right)$$

$$(b)\ (\min\{dist_i\}_{i=1}^{6} < 0.1 \rightarrow a_v = 1) \wedge (\min\{dist_i\}_{i=1}^{6} \geq 0.1 \rightarrow a_v = 0)$$

The total kilometres travelled for each trip was recorded for the half-hour period and day within which the trip was completed. Where more than one trip was completed by a vehicle within a half-hour period, the distance of each trip was summed. Two additional binary features were also added to the dataset that had the potential to impact vehicle availability: holiday i.e. whether the day was a public or university holiday; and weekend. Although the latter could be derived from the day, it was added as a separate feature given its potential importance. The resulting dataset contained over 850,000 rows, 57% of which represented half hour periods in which a vehicle was available. This dataset was used as input to the state-of-charge (SoC) simulations. Sample entries are shown in Table 2.

Table 2: Example entries from the vehicles dataset used as input to the SoC simulations, where v is the vehicle ID from 1 to 48, d is the department ID from 1 to 6, a_v is availability of vehicle v and km is trip distance.

v	d	day	half-hour	weekend	holiday	a_v	km
1	3	1	1	0	1	0	0.4
2	2	3	35	0	0	1	0
3	6	4	26	0	1	0	1.2
3	6	4	27	0	1	1	0
4	1	6	20	1	0	1	0

2.2. State-of-Charge Simulation

To determine battery state-of-charge for each vehicle during the sample period, a simulation was developed that calculated the SoC for each half hour period depending on a vehicle's availability and any trips that had been completed. The aim of this simulation was to support analysis of the learning models by producing a pattern of variation for state-of-charge with similar characteristics to a real-life scenario rather than to accurately simulate given vehicles and/or EVSE. The parameters used and their settings for the experiments considered in this work are defined in Table 3.

Table 3: Parameters for the SoC simulation

Parameter	Description	Setting	Units
battery_capacity	Maximum capacity of a vehicle's battery	40	kWh
rapid_rate	Rapid charging rating of a vehicle, typically using DC power	50	kW
slow_rate	Slow charging rating of a vehicle	6	kW
export_rate	Power rating of export to the grid	50	kW
vehicle_η	Energy used per km travelled by a vehicle	0.2	km/kWh

SoC_min	The minimum state-of-charge that must observed to satisfy vehicle user requirements when calculating available capacity	50	%
SoC_start	Battery charge at the start of the simulation	100	%
charging_η	Efficiency of the charging process accounting for losses	90	%

Much research has been undertaken to understand the impact of battery degradation on electric vehicles (EV) as a result of V2G [23–26]. Evaluation of the impact these different operational models might have on the outputs of the work discussed in this paper is out of scope. However, it must be acknowledged that there is a very real correlation between battery power flow rates and potential return on investment from V2G market trading; high power rates for example have been shown to significantly increase battery aging in Lithium-ion batteries [27]. Depending upon the operational cycle selected however, V2G operation can in some instances extend the life of the battery beyond that of an equivalent battery not undertaking V2G services [28].

In this work we use a simplified charging and discharging model in order to simulate the typical use expected of EVs, whilst ensuring SoC is kept within a usable limit for V2G market trading [29]. It was assumed that a vehicle would use rapid DC charging to a maximum 80% SoC during the typical working day from 07:30 to 19:00, which would allow vehicle batteries to recharge as quickly as possible to support ongoing usage while working within an efficient SoC range for DC charging. It was also assumed that slow charging would be used when vehicles were typically stationary for longer periods between 19:30 to 07:00, which would allow vehicles to charge to 100% where possible while helping to manage battery degradation.

The initial battery SoC was set according to the *SoC_start* parameter and each half hour period processed in chronological order. The state-of-charge for each vehicle for a half-hour period (SoC_v^{hh}) was calculated as follows, where SoC_v^{hh-1} is the state-of-charge of vehicle *v* for the immediately preceding half-hour period:

1. For each half hour period a vehicle was available (a_v=1) during the day from 07:30 to 19:00:

$$SoC_\delta = \frac{rapid\_rate * 0.5 * charging\_\eta}{battery\_capacity}$$

$$SoC_v^{hh} = Min(SoC_v^{hh-1} + SoC_\delta,\ 80)$$

2. For each half hour period a vehicle was available (a_v=1) overnight from 19:30 to 07:00:

$$SoC_\delta = \frac{slow\_rate * 0.5 * charging\_\eta}{battery\_capacity}$$

$$SoC_v^{hh} = Min(SoC_v^{hh-1} + SoC_\delta,\ 100)$$

3. For each half hour period in which a trip was completed i.e. where km > 0:

$$SoC_\delta = \frac{vehicle\_\eta\ * km}{battery\_capacity}$$

$$SoC_v^{hh} = Max(SoC_v^{hh-1} - SoC_\delta,\ 0)$$

4. For all other half hour periods, the SoC was unchanged:

$$SoC_v^{hh} = SoC_v^{hh-1}$$

Following these updates, the potential available capacity from each vehicle within each half hour period ($available\_capacity_v^{hh}$) was calculated using the state-of-charge at the end of the previous period as follows:

$$available\_capacity_v^{hh} = Max\left(\frac{(SoC_v^{hh-1} - SoC\_min)}{100}, 0\right) * battery\_capacity * a_v$$

Thus, capacity was available whenever a vehicle was available with a SoC in excess of the minimum constraint. There are several simplifying assumptions in this simulation: the charging rate was

assumed to be linear; vehicle efficiency was assumed to be the same in all conditions and for all trip types; it was assumed no battery degradation had occurred; and it was assumed all charging stations and vehicles were of the same type. While these factors would need to be considered in a real system, this level of detail was not necessary to support the primary purpose of this study.

The first week of data was used to initialise the simulated SoC. The simulation was run with several different *SoC_start* settings and the vehicles SoC at the end of the initialisation week observed to be identical in all cases. This showed that the simulation was not sensitive to the initial SoC and the subsequent 52 weeks therefore represented valid behavioural patterns rather than nuances of the starting configuration. To illustrate the output of the simulation, Figure 2 shows changes in SoC and available capacity for two sample vehicles from a representative 7 days. Vehicle A was frequently used during the working week but consistently returned to park next to an EVSE allowing it to maintain a high SoC and enable substantial available capacity, such a vehicle would likely be valuable to an aggregator. Vehicle B however was only parked next to an EVSE on Monday and although it was stationary for much of the week this capacity was not available. If this pattern was sustained the vehicle would only occasionally be of value to an aggregator.

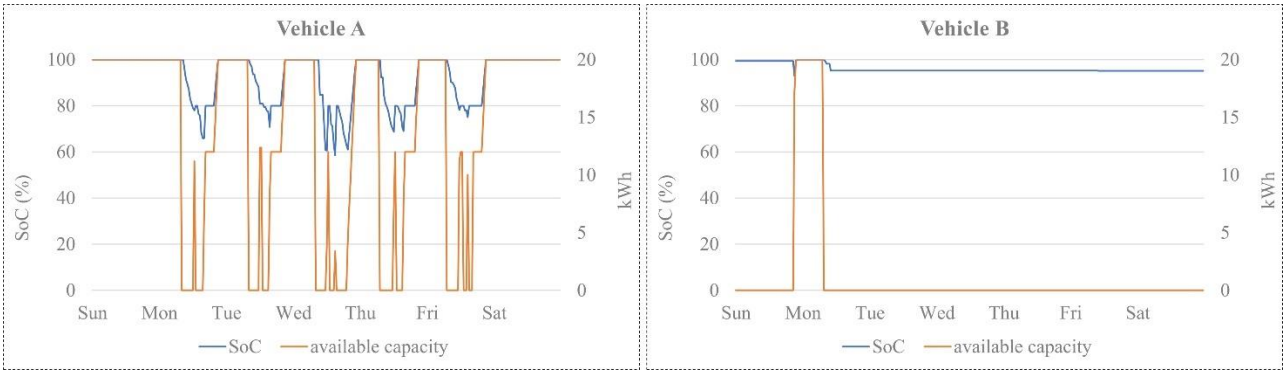

Figure 2: State-of-Charge (left axis) and available capacity (right axis) for 2 sample vehicles over a one-week period in the simulation.

For some vehicles the SoC dropped to zero, which indicated that the usage of the vehicle and/or the chosen locations of the EVSEs was not compatible with participation in a V2G service as they were not able to retain enough charge to sustain normal usage. These vehicles were thus excluded from subsequent analysis, leaving a total of 33 vehicles.

2.2.1. Market Event Simulation

To investigate the impact of exporting energy, the simulations were also extended to include market events that occurred during a specified half-hour period. It was assumed that at the time of a market event all available capacity would be exported, as may be the case where an aggregator calls upon a fleet to deliver all its available reserves to satisfy the requirements of a larger trade for example. For each half-hour period in which a market event occurred (*hh*), the following modifications to the simulation were made, where $exported\_energy_v^{hh}$ is the energy exported by vehicle *v* during half-hour period *hh*:

$$exported\_energy_v^{hh} = Min(available\_capacity_v^{hh}, export\_rate * 0.5)$$

$$SoC_\delta = \frac{exported\_energy_v^{hh}}{battery\_capacity}$$

$$SoC_v^{hh} = SoC_v^{hh-1} - SoC_\delta$$

Note that export losses were not considered in the simulation. These modifications had the effect of exporting all available capacity from each vehicle within the constraints of the maximum export rate. The half-hour period of a simulated market event was chosen randomly within a specified date range with a uniform distribution. Although market events may occur more regularly in a real-life service, this random distribution is likely to represent the most challenging scenario for the model to adapt to. For simulations in which market events were considered, a total of 52 events were generated for the 52-week period following the initialisation week.

2.3. Aggregated Dataset

The output of the state-of-charge simulations was used to generate an aggregated dataset that formed the input to the machine learning models. To determine the total available capacity during each half-hour period, available capacity from each of the 33 viable vehicles was summed. Figure 3 shows an example week of this aggregated available capacity.

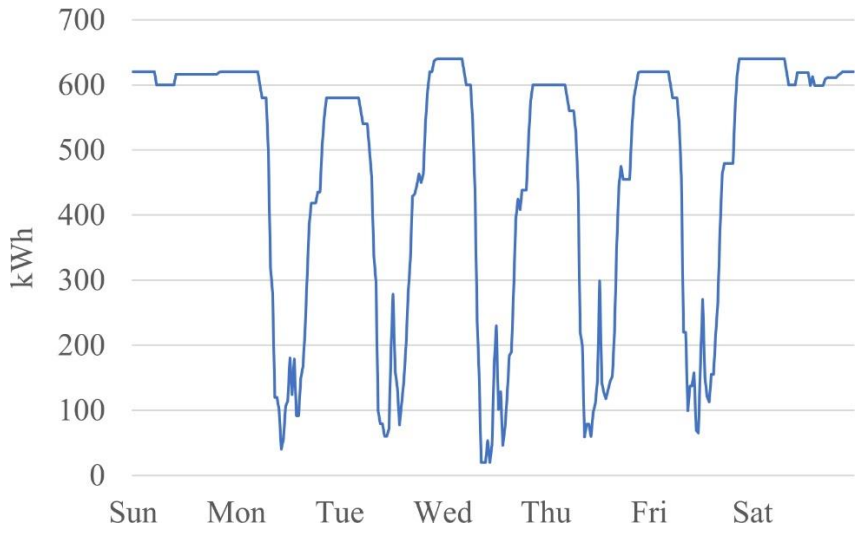

Figure 3: Aggregated available capacity from the 33 viable vehicles over an example week.

For simulations including market events, exported energy was also summed across the 33 vehicles to yield the total exported energy for each half-hour period. This resulted in an aggregated dataset consisting of 17,520 rows (365 days * 48 periods). The first 280 days (76.7%) of the dataset were used as a training dataset and the subsequent 85 days (23.3%) were used as an independent test dataset. There were fourteen holidays in total, 11 (78.6%) of which were in the training dataset and 3 (21.4%) within the test dataset.

For simulations including market events, 40 were randomly generated during the training period and 12 during the test period. Sample entries from the aggregated dataset are shown in Table 4.

Table 4: Sample entries from the aggregated dataset used as input to the machine learning models.

day	half-hour	weekend	holiday	exported energy	available capacity
2	14	0	0	0	580.0
1	29	0	0	0	110.7
6	11	1	0	640.0	640.0
0	37	0	1	0	597.5

2.3.1. Time Series Dataset

The CNN-LSTM network is a multivariate, multi-step time series forecasting model that makes use of information from the immediate past to inform the near future. In this work the previous 24 hours of data was used to predict the following 24 hours. The input data was therefore extended to include the features from the prior 48 half-hour periods and the target feature was extended to include the available capacity for 48 future half-hour periods. Thus, for a given half-hour period or timestep *t*, the model was presented with rows *t*-1 to *t*-48 from the aggregated dataset and was required to make a prediction of available capacity for timesteps *t* to *t*+47. This process is visualised in Figure 4.

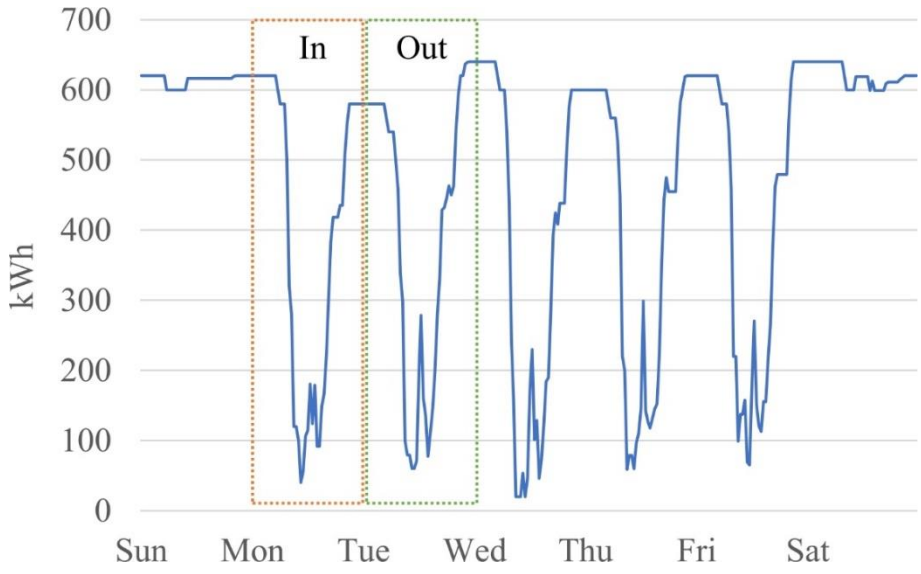

Figure 4: Input and output windows for a sample datapoint at the start of Tuesday. The previous 24-hours of data were used as the input window and predictions were required for the subsequent 24-hours.

For simulations not including market events, 5 input features were used in the input window for each datapoint i.e. day, half-hour, weekend, holiday and available capacity. However, both weekend and holiday features were shifted backwards by a day i.e. by 48 rows. The meaning of these features thus became whether there was a weekend or holiday at the same time tomorrow. This was information that could be more effectively used to influence the forecasted available capacity for the subsequent timesteps. Where market events were simulated, exported energy was also included in the input window and thus a total of 6 features were used. The same 48-row shifting operation was performed for the exported energy column in this case. All features were scaled using a min-max scaler between the bounds 0 and 1.

2.4. Prediction Models

2.4.1. Naïve Persistence Model

To establish a baseline for the machine learning techniques, a naïve persistence model was used. This model assumed that available aggregated capacity for the required future timestep was the same as the currently available capacity. Thus, for each row of the dataset at timestep *t*, predictions of available capacity for timesteps up to *t+47* were all made equal to the available capacity observed at *t-1*.

2.4.2. AutoML Regression Model

Figure 3 shows a relatively regular pattern of availability throughout the week with high availability during the weekend and overnight, reduced availability during the working day and a small peak during lunchtime. To provide a benchmark for the CNN-LSTM, the auto-sklearn automated machine learning (AutoML) toolkit [30] was used to fit a regression model to this general pattern. This model produced a single prediction for each row of the dataset i.e. there was no time series forecasting. The AutoML toolkit optimised pre-processing, model choice and associated hyperparameters in addition to training the model itself. Twelve regression algorithms including K-nearest neighbours, decision trees and extreme random forests were available to the optimisation process together with fourteen feature pre-processing methods and four data pre-processing methods. An ensemble model was also automatically built by auto-sklearn from the individual models generated during the run. The toolkit was run with default parameters [31] using the aggregated dataset with available capacity as the target feature. Market events were not considered and thus the exported energy column was not used. Mean absolute error (MAE) was used as the key metric to compare performance with the CNN-LSTM model as shown in Equation 2:

Equation 2: Mean absolute error (MAE), where n is the number of rows in the test dataset.

$$MAE = \frac{\sum_{i=1}^{n} |predicted_i - actual_i|}{n}$$

To compare the models, predictions were required for the subsequent 48 half-hour periods or timesteps for each row of the test dataset. However, the AutoML regressor produced a single prediction for each row rather than a time series forecast. A prediction for each of the subsequent timesteps could however be generated by presenting the input data for each individually. As each prediction was the same regardless of the timestep at which it was made, MAE was the same for each of the 48 prediction horizons i.e. MAE was the same for the set of predictions made at timestep *t, t+1* and for all subsequent timesteps.

2.4.3. CNN-LSTM Time series Forecasting

The architecture of the CNN-LSTM model is shown in Figure 5. The principal role of the CNN was to learn and encode input features and the role of the LSTM was to decode this internal representation and learn temporal dependencies between features.

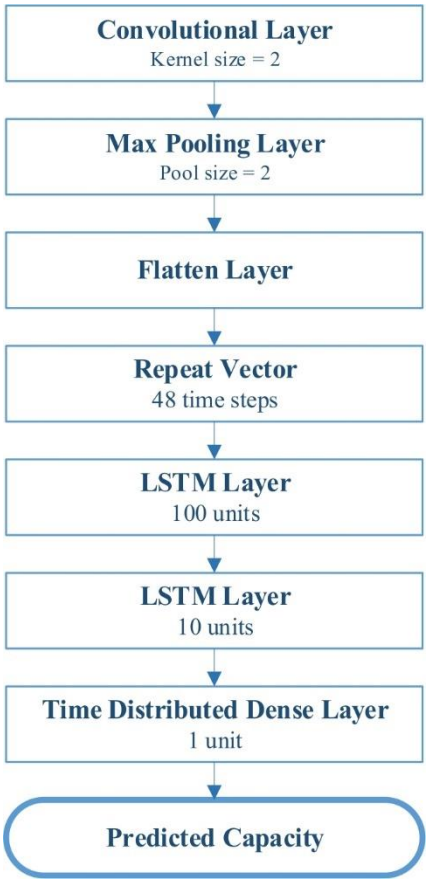

Figure 5: Architecture of the forecasting model. A convolutional neural network was used to encode the input, which was decoded using a stacked-LSTM network.

For a give timestep *t*, the network operated as follows:

1. The previous 24 hours of data from the aggregated dataset was presented as input to the convolutional layer. This resulted in 48 rows i.e. timesteps *t*-1 to *t*-48.
2. A kernel or filter with a 2-row window was used. The filter was moved along the input data one row at a time producing an output at each step via a rectified linear unit (ReLU) activation function. The input data was zero-padded thus resulting in 48 outputs for each filter. This process was repeated for 64 filters supporting the learning of multiple different features in the input data.

3. A one-dimensional max pooling layer was used to down-sample this data and help make the representation more invariant to small changes in the location of the input features. This was achieved by using a filter or pool of size 2 and iterating over the output from the convolutional layer. The maximum value of each pool was retained.
4. A flatten layer was used that concatenated the multi-dimensional output from the max pooling layer into a single vector, thus forming a representation of the features in the input from the previous day.
5. This representation was repeated to produce 48 copies to be used as input to the next layer. A copy was required for each of the subsequent 48 timesteps for which available capacity was to be predicted.
6. An LSTM layer [32] was used which consisted of 100 units, each with the potential to learn different temporal sequences and dependencies in the input data. Each recurrent unit produced 48 outputs representing the 48 timesteps to be predicted.
7. A second LSTM layer was used to support the learning of higher-level dependencies taking input from the 100 individual units in the previous layer. Each recurrent unit in this layer also produced 48 outputs representing the predicted timesteps.
8. A final fully connected dense layer was used that took input from each of the 10 LSTM units in the previous layer and produced as output the predicted available capacity using a ReLU activation function. A time-distributed layer was used meaning that each of the 48 vectors from the previous layer were presented individually. The final output was therefore a 48-element vector representing predicted capacity for the subsequent 24-hour period.

This architecture resulted in 660,083 trainable parameters when all input features were used as detailed in Table 5:

Table 5: Parameter analysis of the CNN-LSTM network

Layer	Output Shape	Trainable Parameters (no exported energy feature)	Trainable Parameters (with exported energy feature)
Convolutional	48 x 64	704	832
Max Pooling	24 x 64	0	0
Flatten	1 x 1536	0	0
Repeat Vector	48 x 1536	0	0
LSTM 1	48 x 100	654,800	654,800
LSTM 2	48 x 10	4,400	4,400
Time Distributed Dense	48 x 1	11	11
	Total:	659,955	660,083

To train the network, batches of 48 training samples were presented before network weights were updated. This batch size was chosen to represent a realistic real-life scenario using on-line learning in which data was received daily from the previous day's activity. Error was minimized using the Adam optimiser [33] and mean squared error across all predictions as the loss function. To protect against overtraining, loss on the independent training set was assessed at the end of every epoch. If this loss did not improve for 10 epochs training was stopped and the best model used to calculate the mean absolute error using the samples in the test dataset. Training was run for a maximum 100 epochs but in practice the actual number of epochs was less than this given the early-stopping feature.

To determine the accuracy of a trained model at different prediction horizons, MAE was calculated for each timestep individually. A time series of predictions for the subsequent 48 half-hour periods was produced by the model for each row of the test dataset. MAE was then calculated across the test dataset for all predictions at each timestep or horizon as shown in Equation 3, where *h* is the prediction horizon from 0 to 47, MAE_{t+h} is the mean absolute error for prediction horizon *h*, $predicted_i^{t+h}$ is the predicted available capacity at horizon *h* for row *i* of the test dataset,

$actual_i^{t+h}$ is the actual available capacity at timestep *h* for row *i* of the test dataset and *n* is the number of rows in the test dataset:

Equation 3: Mean absolute error across the test dataset for prediction horizon h

$$MAE_{t+h} = \frac{\sum_{i=1}^{n} \left| predicted_i^{t+h} - actual_i^{t+h} \right|}{n}$$

3. Results

3.1. Without simulated market events

This section presents results for the Naïve persistence forecast, AutoML regressor and CNN-LSTM model using the aggregated dataset that was produced from simulations that did not include market events. Following training, MAE was calculated for each of the 48 prediction horizons for the AutoML regressor and CNN-LSTM using Equation 2 and Equation 3 respectively. A series of 10 independent training and test runs were performed for both models and the results were averaged across these 10 runs in each case. Results are shown in Figure 6.

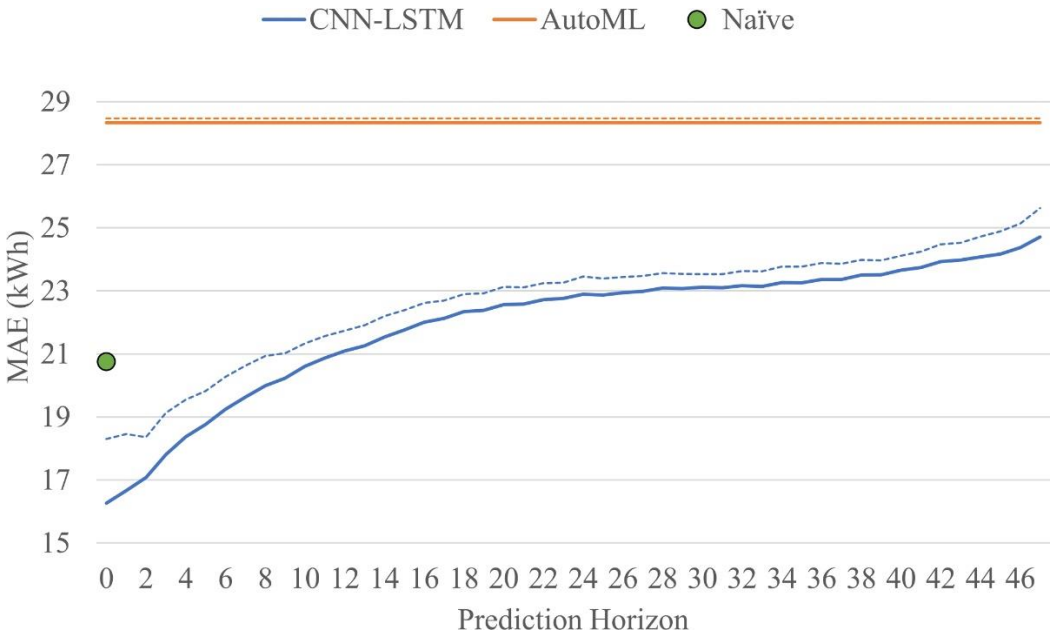

Figure 6: Average results for 10 runs of the AutoML regressor and the CNN-LSTM model for prediction horizons up to 48 timesteps (24-hours). Dashed lines show 1 standard deviation above the mean. Results at prediction horizon 0 are also shown for a Naïve persistence forecast.

MAE for the naïve forecast quickly escalated to 36.7 kWh and 51.3 kWh at prediction horizons 1 and 2 respectively and therefore results are only shown for horizon 0 i.e. timestep *t*. The ability for this simple approach to demonstrate an improvement over the AutoML regressor at timestep *t* did however suggest there was value in using the current available capacity to inform predictions. This hypothesis was strengthened by the CNN-LSTM model results, which consistently outperformed the AutoML regressor. Results for the latter were the same at all prediction horizons as this technique did not include any explicit time series forecasting i.e. the model provided the same prediction for a given data point independently from when that prediction was made. The MAE across all horizons was 21.95 kWh (s = 0.64 kWh) for the CNN-LSTM compared with 28.4 kWh (s = 0.13 kWh) for AutoML, which was a 22.7% improvement in accuracy. A Welch's t-test was used to confirm that this was a highly statistically significant difference ($p < 0.001$).

To highlight the differences in the performance of the two models, actual capacity and predicted capacity at a prediction horizon of 0 (i.e. the current timestep), for a sample 4-week period within the test dataset are shown in Figure 7. Given the parameters selected for the SoC simulation the maximum capacity each vehicle was able to contribute during a half-hour period was 20kWh i.e. 50% of the maximum battery capacity of 40kWh. This energy would be delivered at a rate of 40kW over the 30 minutes. With 33 viable vehicles the maximum capacity available from the fleet for any one period was therefore 660kWh.

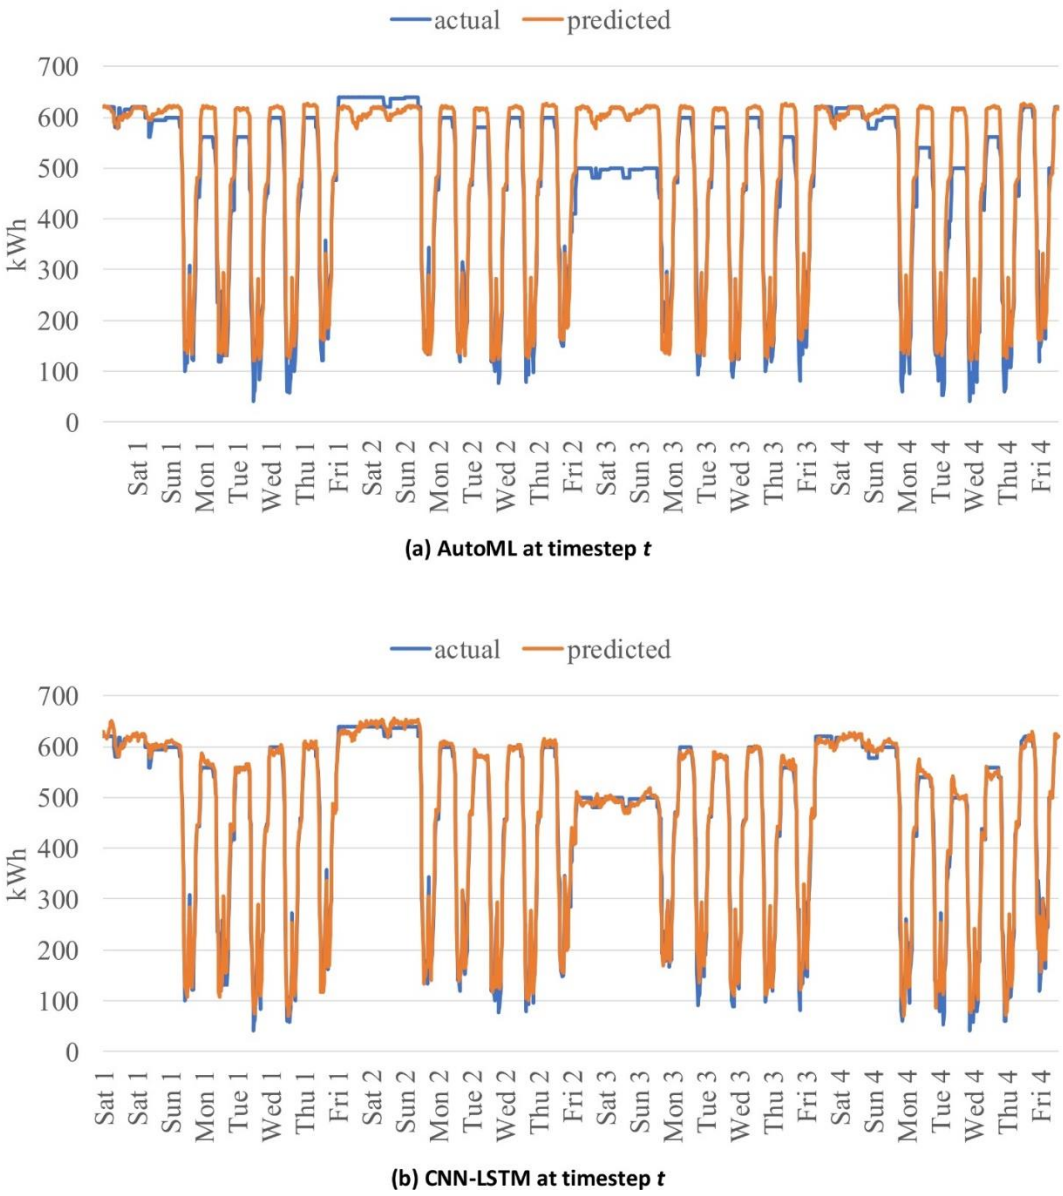

Figure 7: Actual vs predicted capacity for a sample 4-week period in the test dataset. AutoML regressor (a); CNN-LSTM model (b) at prediction horizon 0 i.e. the current timestep (t).

The figure reveals that the capacity predicted by the AutoML regressor followed a regular pattern during a calendar week and there was no response to variations from the typical pattern that occurred during the sample period. This was particularly apparent for the 3rd weekend, during which available capacity was significantly lower than usual and hence the prediction error was relatively high; such overestimation could result in financial penalties for an aggregator. The

Pearson correlation coefficient for the AutoML regressor, r, over the 4 weeks was calculated as $r = 0.969$. In contrast, the CNN-LSTM model was able to quickly adapt to deviations from the typical pattern and predicted capacity closely correlated with actual capacity with a Pearson's $r = 0.993$.

Figure 6 showed, however, that the performance of the CNN-LSTM model gradually decreased as the prediction horizon increased. To provide more insight into the difference in performance for different prediction horizons, predicted and actual capacity for timesteps 11 (6 hours) and 47 (24 hours) over the same 4-week period are shown in Figure 8.

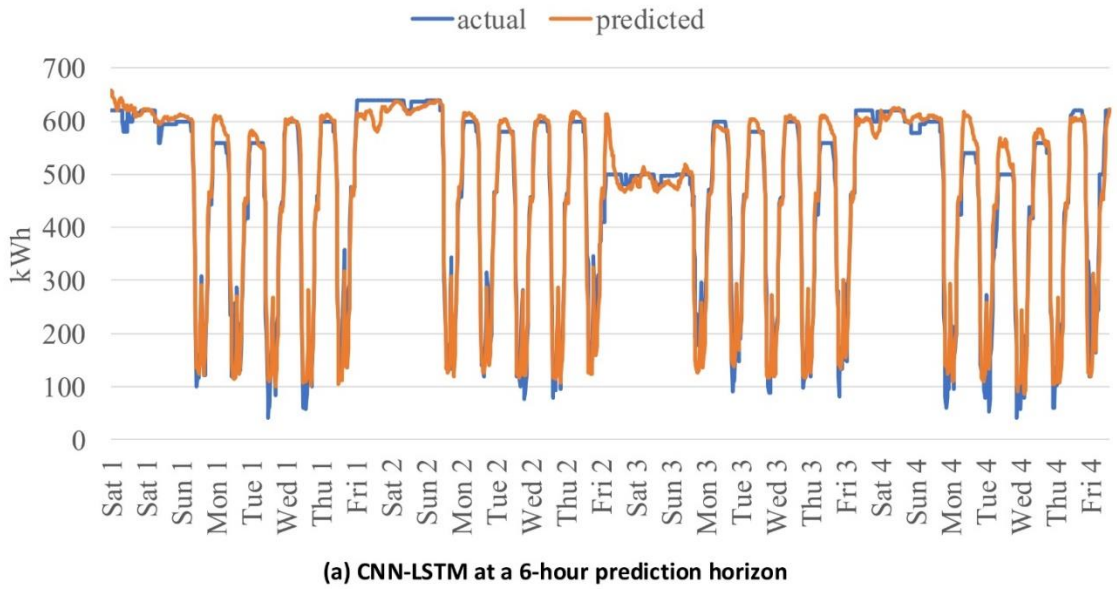

(a) CNN-LSTM at a 6-hour prediction horizon

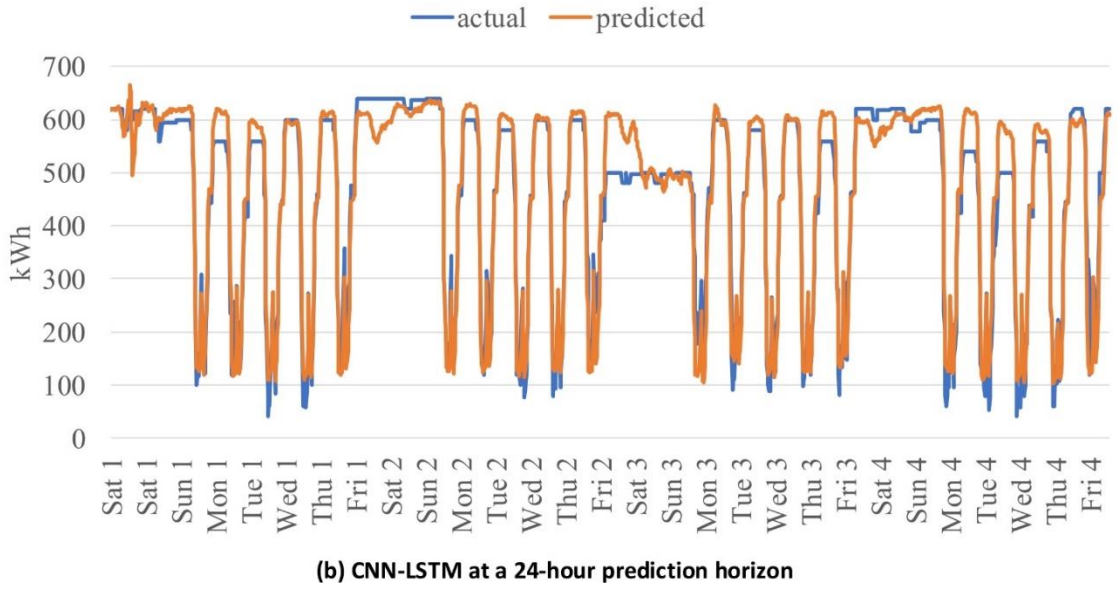

(b) CNN-LSTM at a 24-hour prediction horizon

Figure 8: Actual vs predicted capacity for the sample 4-week period for the CNN-LSTM model at timestep 11 i.e. a 6-hour prediction horizon (a); and timestep 47 i.e. a 24-hour prediction horizon (b).

At timestep 11 i.e. a 6-hour prediction horizon, delays in the model's response to deviations from the typical pattern became apparent and Pearson's r dropped to 0.985. Predicted capacity at the start of the 3rd weekend for example was initially much higher than the predicted capacity, before the model responded and more closely predicted actual capacity. The model also responded in a similar way to some of the differences in available capacity for overnight periods, for example on

the Monday and Tuesday of week 4, however these periods were typically over before the model could fully recover. As there were no features in the dataset to indicate when deviations from the typical pattern would occur the model relied on autoregression of the available capacity feature and therefore could not modify its prediction until the presence of the deviation was presented to it within its input window. This effect became particularly apparent at timestep 47 i.e. a 24-hour prediction horizon. The response to the 3rd weekend deviation did not occur until it had been present for over 24-hours and there was now no discernible response to deviations in overnight periods. Pearson's r was calculated as 0.976.

3.2. With simulated market events

This section presents results for the CNN-LSTM model using the aggregated dataset that was produced from simulations that included market events and hence the export of energy. Two different experiments were performed using this dataset:

1. *No exported energy feature*: although the market events impacted available capacity due to the export of energy, the exported energy column was not added to the input features.
2. *With exported energy feature*: the additional exported energy column was included in the aggregated dataset resulting in 6 input features.

A total of 10 independent runs were performed for each scenario. The results shown in Figure 9 show the MAE averaged over these runs. In both cases the error was higher than the experiments with no market events, which was as expected given the additional uncertainty in the simulation. The MAE across all timesteps with the exported feature was 26.90 kWh (s = 1.05 kWh) compared to 33.13 kWh (s = 1.09 kWh) without the exported feature, which was an 18.9% improvement. A

Welch's t-test was again used to confirm that this was a highly statistically significant effect ($p < 0.001$).

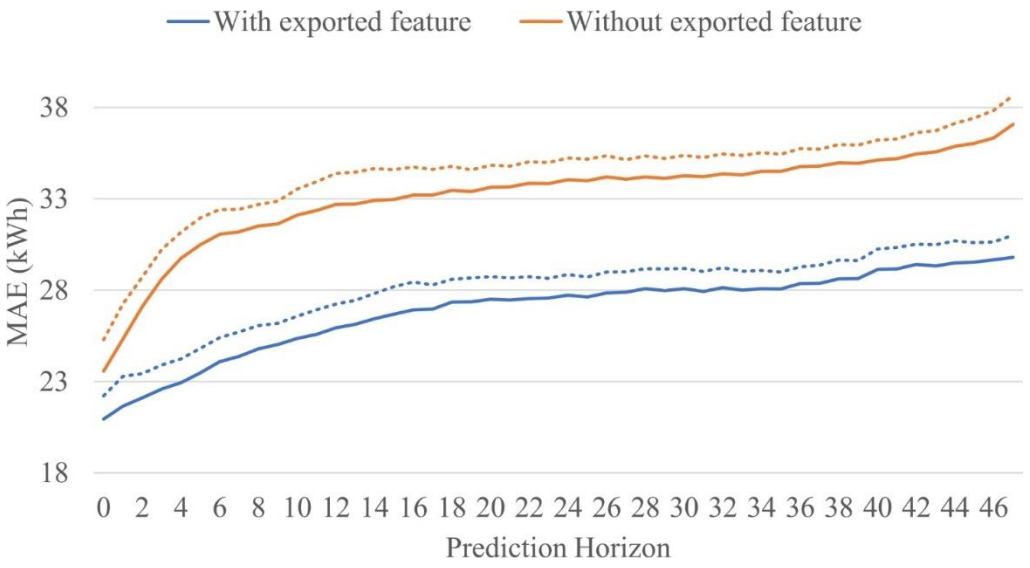

Figure 9: Average results for 10 runs of the CNN-LSTM model for a 48 timestep (24-hour) prediction horizon with and without the exported energy feature. Simulations included market events. Dashed lines show 1 standard deviation above the mean.

To explore the reasons for this discrepancy, predicted capacity for the current timestep, during an example weekend in which a market event occurred is shown in Figure 10 for a model trained with the exported energy feature (a) and without that feature (b).

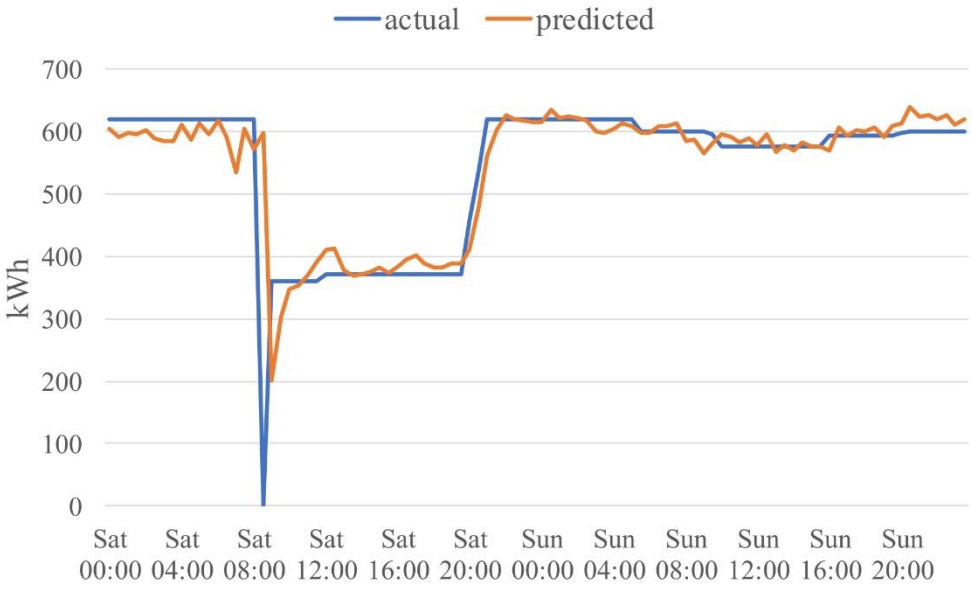

(a) Prediction at timestep *t* using the exported energy feature

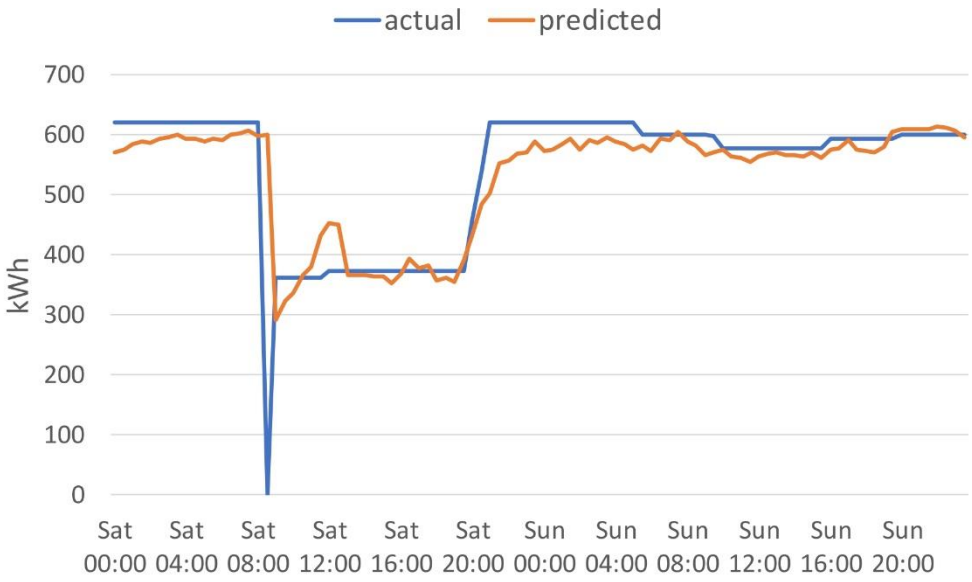

(b) Prediction at timestep *t* with no exported energy feature

Figure 10: Impact of a simulated market event at Saturday, 08:00 on predicted capacity for the current timestep for a model trained with the exported feature (a) and without that feature (b).

Typically, the pattern during the weekend was for available capacity to remain relatively consistent. However, the market event reduced available capacity to 0 before most vehicles used in that event charged back to 80% using rapid charging. During the evening, slow charging was then used to

charge back to 100%. This impact was predicted by both models, which responded to the presence of a market event by quickly reducing predicted capacity. In both cases the response did not occur until the time period after the event and hence this information was present in the input window of the model. This suggested that autoregression was used in both cases. The MAE over this period was 24.5 kWh for model (a) as opposed to 31.2 kWh for model (b), however this was not a statistically significant difference ($p > 0.05$).

The same analysis was performed for this market event at timestep 47 i.e. a 24-hour prediction horizon. Actual and predicted capacity for both models are shown in Figure 11. In this case, it was no longer possible for either model to use autoregression as by the time the impact of the event on available capacity was present in the input data its impact on subsequent time periods had passed i.e. there was no impact of the event 24-hours after it had occurred. However, the figure shows that model (a) still responded to the event by adjusting predicted capacity. Note that the model now responded at the actual time period of the event rather than the next time period confirming that autoregression was not used as the impact on available capacity was not yet present in the input window. The impact of the market event was completely missed by model (b), which resulted in MAE over this period of 64.6 kWh as opposed to 33.6 kWh for model (a). A highly statistically significant difference ($p < 0.001$).

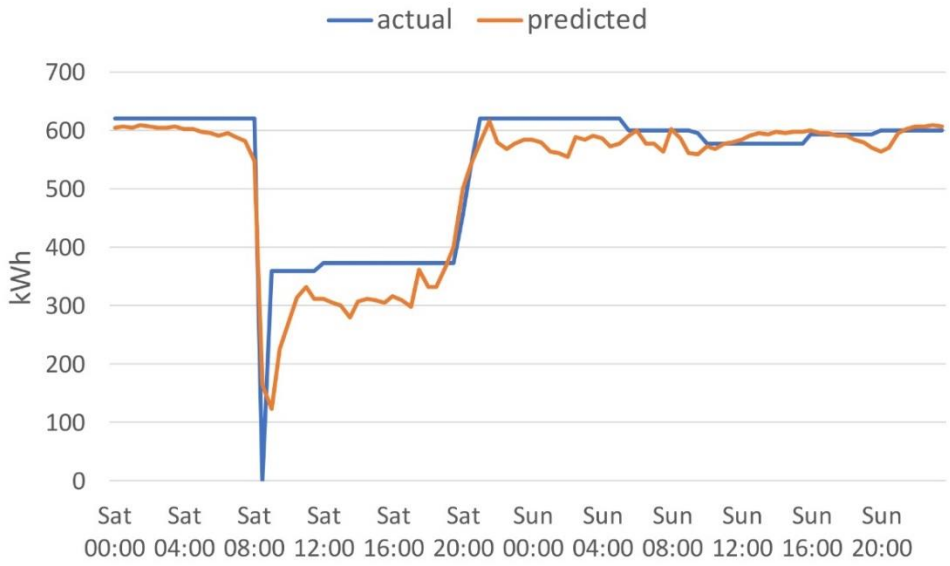

(a) 24-hour prediction horizon using the exported energy feature

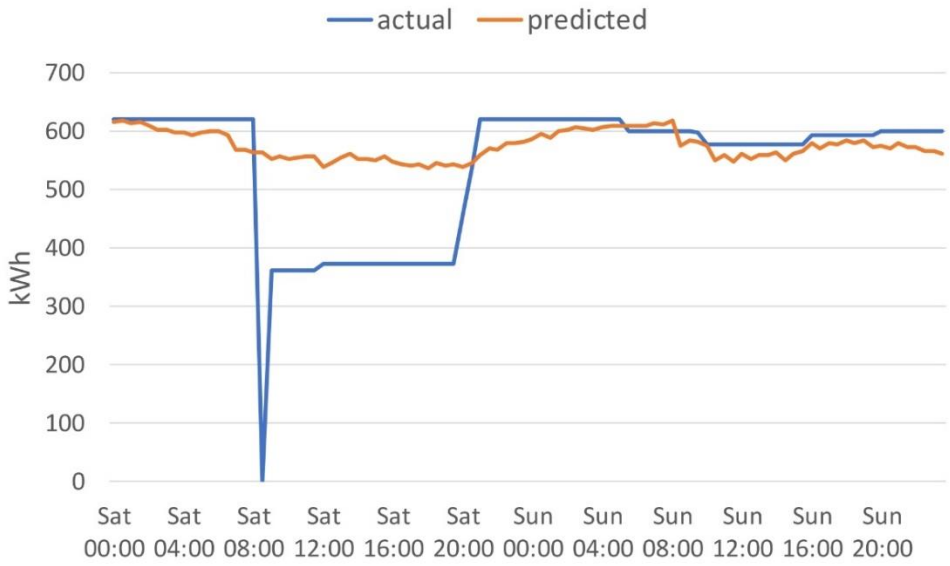

(b) 24-hour prediction horizon without the exported energy feature

Figure 11: Impact of a simulated market event at Saturday, 08:00 on predicted capacity at timestep 47 i.e. 24-hour prediction horizon for a model trained with the exported feature (a) and without that feature (b).

4. Discussion

The results presented in this paper showed that for accurate near-term predictions, consideration of the current status and recent past was important to allow inevitable variations in typical patterns of available capacity to be tracked. The CNN-LSTM model was able to use the previous day's data to more accurately predict available capacity for the next day, which would allow a V2G service to refine its position in the energy markets and assess the impact of ancillary service delivery. To do so however, knowledge of the variations in available capacity was required in the input window of the model to support use of autoregression. As an example, in the dataset used in this work, fewer vehicles than usual were available to the service on some weekends. This may have been caused by several factors, for example special events at the University or adverse weather or traffic conditions making it difficult for the vehicles to return to their usual weekend parking locations. These, or other potential factors were not explicitly represented in the training dataset however and therefore the model was only able to respond to the corresponding impact on available capacity when that impact was presented to it within its input window. For this reason, the earliest the model could respond was the length of the prediction horizon. For short prediction horizons, any errors were quickly rectified however for longer prediction horizons these errors persisted for corresponding lengths of time. For 24-hour ahead predictions any sustained variation would need to be present for at least 24-hours for the model to be able to respond through autoregression. For a variation lasting a whole weekend therefore the model was able to respond for the second day of the weekend, however for shorter durations such as for overnight periods, the variations were over before they entered the input window and thus the model was unable to respond. Prediction errors thus grew as the prediction horizon was increased.

To allow the model to respond to deviations from the typical pattern for longer prediction horizons, additional features in the multivariate model were required. This approach was taken for weekends and holidays that were shifted back a day such that they were represented in the model's input window 24-hours prior to their occurrence i.e. their meaning changed to become a weekend or holiday at the same time on the subsequent day. Another example of such a feature was the exported energy feature. In this case, the future export commitments that the V2G service had entered in to were added to the input window 24 hours in advance so that the model was able to adjust its predictions accordingly based on the impact of such exports in the training data. The results showed that the model was able to use this information to respond to randomly scheduled market events 24 hours in advance. This approach could also be extended to add features for different timescales. For example, as the V2G service refined its position closer to gate closure, additional trades could be added in a 12 hour ahead feature or a 6 hour ahead feature. These would extend the feature set used during training of the model.

In addition to predicting the impact of exports that have already been committed to, the model could also be used as an exploratory tool by the V2G service to determine the value of potential opportunities. There are likely to be multiple objectives in deciding on the best opportunities, for example: availability of the required capacity at the required time; the predicted margin of available capacity for those not using all the fleets reserves; revenue potential of the opportunity; and the ability of the fleet to recharge and support potential demands from ancillary services or additional trades in the near-future. In this way, the trained model effectively becomes part of a fitness function for a multi-objective optimisation algorithm [34,35], a technology that also has potential for optimising the design of the model [36].

The model used in this work made a prediction for 48 timesteps ahead i.e. 24 hours. This prediction horizon may not, however, be sufficient to meet all the near-term needs of the V2G service. For example, the day-ahead market typically closes at midday on the day before the trade. Predictions at a 36- or 48-hour horizon may therefore be required to maximise the opportunities in this market. The model explored in this work can however be extended to increase the prediction horizon by introducing additional timesteps in the recurrent networks at the heart of the decoder i.e. by increasing the number of repeats of the encoded features vector. This may also require increasing the number of units in the LSTM layers to learn the additional temporal dependencies. However, the results showed that the performance of the model decreased as the prediction horizon increased, a trend that would likely continue as the predictive power of the recent past diminished. One approach for improving long-horizon forecasts for LSTM networks is the inclusion of an expectation bias [37]. However, it is also likely to be of value to emphasise the contribution of a model that learns the general pattern of available capacity such as the regression model used in this work as the prediction time horizon increases. For both approaches, identification and explicit inclusion of features impacting available capacity will be important to reduce prediction errors and optimise the model.

Another important consideration is ongoing adaptation of the predictive model. While a historical data set allows training of an initial model, it must be continually updated using new data acquired during operation of the service. The Adam optimizer used in this work is an extension of stochastic gradient descent, which is well suited to such on-line learning. Training data was presented to the model in batches of 48 timesteps, which is representative of a real-life scenario in which 24 hours of data may be received at the end of each day's operation and used to continually update the model. This process would allow refinement of the model to variation in the fleet's behaviour,

however, more substantial changes may require concept drift aware algorithms that automatically tune the learning rate to support more rapid adjustment of network parameters [38].

This work has been principally concerned with the prediction of available capacity given typical vehicle behaviour. However, it should be noted that future V2G services will operate in the context of future power grids that are increasingly fed by renewable energy, which can result in large power fluctuations and frequency instabilities [39]. V2G has great potential in helping to balance out these fluctuations through participation in ancillary services for example however it will be important for the energy system as a whole to ensure that it is integrated so as to ensure it contributes to the solution rather than exacerbate the problem. In some contexts, such as for smaller scale community energy systems, it may also be necessary for the V2G service to consider prediction of renewable energy generation, which is inherently intermittent and irregular thus introducing additional challenges for the predictive models that may require consideration of additional or alternative techniques.

5. Conclusions

In this work we have developed a deep CNN-LSTM time series forecasting model to predict the available capacity from a fleet of 48 vehicles for the next 24 hours. Such forecasting is important to support trading and vehicle utilisation decisions made by emerging V2G aggregation services in the critical near-term period. The model achieved a MAE of 21.95 kWh over a 24-hour prediction horizon, which was a highly statistically significant improvement of 22.7% over a regression model trained using automated machine learning. It was demonstrated that the CNN-LSTM model was able to use knowledge of the immediate past to correct for inevitable deviations from the regular patterns of available capacity learned by the AutoML regressor. However, this ability diminished as

the prediction horizon grew thus highlighting the need to represent the sources of such deviations in the input features where possible.

The model was also demonstrated to be capable of adapting to the operation of the V2G service by utilising knowledge of previous energy export events to predict the impact of upcoming events at randomised times. The CNN-LSTM model achieved a MAE of 26.9 kWh with simulated market events that were “signposted” in the input features 24-hours in advance, a highly statistically significant improvement of 18.9% over the same model without such signposting. This ability to adapt to the impact of the actions of the aggregator is critical to inform ongoing operation of the service and to support multiple market events within a short timescale.

This work has considered a relatively small fleet of vehicles and a live service would require a larger and more complex operation including: scaling the service with additional vehicles and fleets; dealing with inevitable heterogeneity of vehicles, batteries, driving styles and EVSE; and optimising revenue from multiple potential market opportunities. All these factors add additional complexity to the optimisation problem and will be the subject of future work.

Acknowledgements

This work is supported by the European Space Agency contract number 4000120818/17/NL/US and is part of a collaborative project with our partners Brixworth Technologies, Cenex and Kearney who hold patent number GB2566596. The analysis of potential EVSE locations was performed as part of EV-elocity (project number 104250) funded by the Office for Low Emissions Vehicles (OLEV), the Department for Business, Energy and Industrial Strategy (BEIS) and facilitated by Innovate UK.

References

[1] Kempton W, Tomić J. Vehicle-to-grid power implementation: From stabilizing the grid to supporting large-scale renewable energy. J Power Sources 2005;144:280–94. https://doi.org/10.1016/j.jpowsour.2004.12.022.

[2] Waldron J, Rodrigues L, Gillott M, Naylor S, Shipman R. Towards an electric revolution: a review on vehicle- to-grid, smart charging and user behaviour. Proc. 18th Int. Conf. Sustain. Energy Technol. (SET 2019), 2019.

[3] Shipman R, Naylor S, Pinchin J, Gough R, Gillott M. Learning capacity: predicting user decisions for vehicle-to-grid services. Energy Informatics 2019;2:37. https://doi.org/10.1186/s42162-019-0102-2.

[4] Shipman R, Waldron J, Naylor S, Pinchin J, Rodrigues L, Gillott M. Where Will You Park? Predicting Vehicle Locations for Vehicle-to-Grid. Energies 2020;13. https://doi.org/10.3390/en13081933.

[5] Kempton W, Tomić J. Vehicle-to-grid power fundamentals: Calculating capacity and net revenue. J Power Sources 2005;144:268–79. https://doi.org/10.1016/j.jpowsour.2004.12.025.

[6] Tan KM, Ramachandaramurthy VK, Yong JY. Integration of electric vehicles in smart grid: A review on vehicle to grid technologies and optimization techniques. Renew Sustain Energy Rev 2016;53. https://doi.org/10.1016/j.rser.2015.09.012.

[7] Nuvve Corporation. Nuvve V2G Technology 2020. https://nuvve.com/technology/ (accessed

May 2, 2020).

[8] A.T. Kearney. Aggregating Energy Resources. WO/2019/020632, 2019.

[9] Engie. Understanding the Capacity Market 2016. https://www.engie.co.uk/wp-content/uploads/2016/07/capacitymarketguide.pdf (accessed May 7, 2020).

[10] Nord Pool Group. The power market 2020. https://www.nordpoolgroup.com/the-power-market/ (accessed May 18, 2020).

[11] Villar J, Bessa R, Matos M. Flexibility products and markets: Literature review. Electr Power Syst Res 2018;154:329–40. https://doi.org/https://doi.org/10.1016/j.epsr.2017.09.005.

[12] Lam AYS, Leung KC, Li VOK. Capacity estimation for vehicle-to-grid frequency regulation services with smart charging mechanism. IEEE Trans Smart Grid 2016;7:156–66. https://doi.org/10.1109/TSG.2015.2436901.

[13] Zhang H, Hu Z, Xu Z, Song Y. Evaluation of Achievable Vehicle-to-Grid Capacity Using Aggregate PEV Model. IEEE Trans Power Syst 2017;32:784–94. https://doi.org/10.1109/TPWRS.2016.2561296.

[14] Mao M, Zhang S, Chang L, Hatziargyriou ND. Schedulable capacity forecasting for electric vehicles based on big data analysis. J Mod Power Syst Clean Energy 2019;7:1651–62. https://doi.org/10.1007/s40565-019-00573-3.

[15] Yang Q, Li J, Cao W, Li S, Lin J, Huo D, et al. An improved vehicle to the grid method with battery longevity management in a microgrid application. Energy 2020;198:117374. https://doi.org/https://doi.org/10.1016/j.energy.2020.117374.

[16] Sainath TN, Vinyals O, Senior A, Sak H. Convolutional, Long Short-Term Memory, fully connected Deep Neural Networks. 2015 IEEE Int. Conf. Acoust. Speech Signal Process., 2015, p. 4580–4.

[17] Donahue J, Hendricks LA, Rohrbach M, Venugopalan S, Guadarrama S, Saenko K, et al. Long-Term Recurrent Convolutional Networks for Visual Recognition and Description. IEEE Trans Pattern Anal Mach Intell 2017;39:677–91. https://doi.org/10.1109/TPAMI.2016.2599174.

[18] Kim TY, Cho SB. Predicting residential energy consumption using CNN-LSTM neural networks. Energy 2019;182:72–81. https://doi.org/10.1016/j.energy.2019.05.230.

[19] Hou Y, Edara P. Network Scale Travel Time Prediction using Deep Learning. Transp Res Rec 2018;2672:115–23. https://doi.org/10.1177/0361198118776139.

[20] Shi X, Chen Z, Wang H, Yeung DY, Wong WK, Woo WC. Convolutional LSTM network: A machine learning approach for precipitation nowcasting. Adv Neural Inf Process Syst 2015;2015-Janua:802–10.

[21] Trakm8 Limited. Telematic solutions. 2019.

[22] Waldron J, Rodrigues L, Gillott M, Naylor S, Shipman R. Decarbonising Our Transport System: User Behaviour Analysis to Assess the Transition to Electric Mobility. 35th PLEA Conf. Sustain. Archit. urban Des. (to Appear., 2020.

[23] Anseán D, Dubarry M, Devie A, Liaw BY, García VM, Viera JC, et al. Fast charging technique for high power LiFePO4 batteries: A mechanistic analysis of aging. J Power Sources 2016;321:201–9. https://doi.org/10.1016/j.jpowsour.2016.04.140.

[24] Uddin K, Dubarry M, Glick MB. The viability of vehicle-to-grid operations from a battery technology and policy perspective. Energy Policy 2018;113:342–7. https://doi.org/https://doi.org/10.1016/j.enpol.2017.11.015.

[25] Dubarry M, Devie A, McKenzie K. Durability and reliability of electric vehicle batteries under electric utility grid operations: Bidirectional charging impact analysis. J Power Sources 2017;358:39–49. https://doi.org/10.1016/j.jpowsour.2017.05.015.

[26] Uddin K, Gough R, Radcliffe J, Marco J, Jennings P. Techno-economic analysis of the viability of residential photovoltaic systems using lithium-ion batteries for energy storage in the United Kingdom. Appl Energy 2017;206:12–21. https://doi.org/10.1016/j.apenergy.2017.08.170.

[27] Somerville L, Bareño J, Trask S, Jennings P, McGordon A, Lyness C, et al. The effect of charging rate on the graphite electrode of commercial lithium-ion cells: A post-mortem study. J Power Sources 2016;335:189–96. https://doi.org/10.1016/j.jpowsour.2016.10.002.

[28] Uddin K, Jackson T, Widanage WD, Chouchelamane G, Jennings PA, Marco J. On the possibility of extending the lifetime of lithium-ion batteries through optimal V2G facilitated by an integrated vehicle and smart-grid system. Energy 2017;133:710–22. https://doi.org/10.1016/j.energy.2017.04.116.

[29] Burgess M, Harris M, Walsh C, Carroll S, Mansbridge S, King N, et al. Assessing the viability of electric vehicles in daily life: A longitudinal assessment (2008–2012). 2013 World Electr. Veh. Symp. Exhib., 2013, p. 1–5. https://doi.org/10.1109/EVS.2013.6915017.

[30] Feurer M, Klein A, Eggensperger K, Springenberg J, Blum M, Hutter F. Efficient and Robust

Automated Machine Learning. In: Cortes C, Lawrence ND, Lee DD, Sugiyama M, Garnett R, editors. Adv. Neural Inf. Process. Syst. 28, Curran Associates, Inc.; 2015, p. 2962–70.

[31] Machine Learning Lab U of F. AutoSklearn 0.6.0 documentation. Auto-Sklearn Man 2019. https://automl.github.io/auto-sklearn/master/api.html#regression (accessed April 3, 2020).

[32] Hochreiter S, Schmidhuber J. Long Short-Term Memory. Neural Comput 1997;9:1735–80. https://doi.org/10.1162/neco.1997.9.8.1735.

[33] Kingma DP, Ba J. Adam: A Method for Stochastic Optimization 2014.

[34] Merhy G, Nait-Sidi-Moh A, Moubayed N. A multi-objective optimization of electric vehicles energy flows: the charging process. Ann Oper Res 2020. https://doi.org/10.1007/s10479-020-03529-4.

[35] Knezović K, Soroudi A, Keane A, Marinelli M. Robust multi-objective PQ scheduling for electric vehicles in flexible unbalanced distribution grids. IET Gener Transm Distrib 2017;11:4031–40. https://doi.org/10.1049/iet-gtd.2017.0309.

[36] Chandra R. Multi-objective cooperative neuro-evolution of recurrent neural networks for time series prediction. 2015 IEEE Congr. Evol. Comput., 2015, p. 101–8. https://doi.org/10.1109/CEC.2015.7256880.

[37] Abdelsalam Ismail A, Wood T, Corrada Bravo H. Improving Long-Horizon Forecasts with Expectation-Biased LSTM Networks. ArXiv E-Prints 2018:arXiv:1804.06776.

[38] Miyaguchi K. Cogra : Concept-Drift-Aware Stochastic Gradient Descent for Time-Series Forecasting 2019.

[39] Schmietendorf K, Kamps O, Wolff M, Lind PG, Maass P, Peinke J. Bridging between load-flow and Kuramoto-like power grid models: A flexible approach to integrating electrical storage units. Chaos An Interdiscip J Nonlinear Sci 2019;29:103151. https://doi.org/10.1063/1.5099241.