AI for Sustainable Energy

Kostas Vlachodimitropoulos[†]
Local AI
Kalamata, Greece

George Vlachodimitropoulos
University of Aegean
Samos, Greece

Charalambos Skianis
University of Aegean
Samos, Greece

ABSTRACT

The continuation of the substitution of fossil fuel generated electricity to renewable sources generated electricity is now at a critical point. Since these sources have a varying and uncontrollable pattern, they cannot be aligned to that of the consumption. Besides the expensive storage, the new trend is to influence the consumption side with flexible use of electricity pushing its pattern to match the green production one. There is also a European Framework promoting Smart Energy Appliances, appliances that communicate with the Grid regarding the usage patterns.

This introduces the need of modeling and forecasting the end-user's consumption and the operation of intelligent agents on the consumption side making hundreds of decisions throughout the day that affect the usage of energy coming from the Grid. AI is the key enabler of these operations, and in our work, which is funded by the EU Horizon 2020 program, we use Reinforcement Learning, LSTM, RCNN, etc. for precision forecasting usage in EV Charging Stations, Households, Radio Base Stations, etc.

We developed an AI-based intelligent agent to make on-the-fly decisions for energy consumption for prosumer households. We train the agent to minimize the energy cost of a household by controlling the battery given the energy market prices, the energy production by the PV, and the user preferences. We do that using state-of-the-art Reinforcement Learning algorithms such as Deep Inventory Management, Deep Q Networks, and others. The agent learns by making actions in an environment, receiving rewards/penalties for those actions, and modifying its action pattern (policy) accordingly. We use the Gym framework, which is offered by OpenAI, to create a simulation environment for our agent and for that we will need to model the states, the possible actions, and their outcomes.

The result is a win-win case, with Financial Savings for the household and positive impact on Electricity Production Carbon Emissions.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

WOODSTOCK'18, June, 2018, El Paso, Texas USA

© 2018 Copyright held by the owner/author(s). 978-1-4503-0000-0/18/06...$15.00

https://doi.org/10.1145/1234567890

KEYWORDS

Neural networks, cybersecurity, classifiers of IoT flows, Lightweight attack detection

1 Introduction - Challenge

Based on the global effort to reduce carbon emissions European Nations including Greece have gone a long way to convert their Electricity Production from the use of fossil fuels to green renewable sources. This conversion varies among the developed countries but in many of them has reached more than 50%.

The remaining 50% is more challenging to achieve because the renewable sources have their own generation pattern that does not coincide with the consumption side. There is a great effort worldwide to deploy energy storage for this reason, but this has many limitations of its own, beyond the extremely high investments involved. The European Union has issued directives and recommendations to achieve a level of flexibility on the consumption profile, aligning it as much as possible to the green production one.

This alignment involves smart energy agents acting from the consumer side managing the production (if any), the local storage and the actual consumption of the new smart energy appliances that emerge. AI provides powerful mechanisms that can help at this mechanism, contributing accurate forecasts and real time decision making targeted to optimize carbon footprint of the energy consumed by the user, such as a EV owner, a household, a corporate building etc.

1.1 Previous Work – AI4CS

Funded by the Horizon 2020 Interconnect project, we have created an AI powered solution that helps EV owners to find the place and the time to charge their EV, targeting both the user customer experience and the usage of Green Electricity and the reduction of stress on the Grid.

The Interconnect project has developed SIF, a Semantics Interoperability Framework that enables the smooth information exchange regarding the electricity production, consumption, carbon intensity etc. of European Energy Grids. Our platform AI4CS (AI for Charging Stations) from one side used AI to forecast availability of Charging points and combined it with information from SIF to produce the optimal suggestions for the place and the time of the charging.

.

1.2 Current Electricity Market trends – Storage Opportunity

The electricity production outlook [20], [24] indicates the high percentage of green electricity, driven by the continuously increasing production from solar and wind power plants.

Current studies as the one recently published by Aurora Energy Research, [17], show that there is a high potential in the usage of energy storage, since it can help increase the usage of cheaper green electricity by deferring the consumption from the grid from expensive carbon intense timeslots to greener ones.

The study of the University of Bonn [21] shows clearly the disparity between the solar/wind production profile and the consumption of a typical household.

This environment shows an opportunity to bring AI as the technology to maximize the benefits of energy storage introduction by making the right decisions in real-time switching between the usage of grid and stored electricity, charging and discharging.

2 AI for Smart Buildings

2.1 LTSM timeseries forecasting

The immense capabilities of neural networks to extract complex patterns from given data intuitively seems a great tool to use for modeling the charging activity in the EV charging networks. LSTMs are renowned for applications of handling multivariate time series [6] and in general cases where the data intrinsically show some temporal dependencies. LSTM belongs to the category of Recurrent neural networks and has been successful in time series forecast, especially when seasonality and cycles need to be accounted for. We have leveraged previous works that employ different types of recurrent neural networks (LSTMs, Random Neural Networks, GRUs) for anomaly detection in IoT systems via time series forecasting for the problem of congestion/utilisation forecast

The Charging Station Occupancy Prediction is getting more and more attention as predicting the occupancy of public EV charging stations is crucial for developing smart charging strategies. It helps reduce inconvenience for both EV operators and users.

Long Short-Term Memory (LSTM) neural networks have been successfully applied to predict energy consumption and occupancy patterns at electric vehicle (EV) charging stations. Let’s delve into how LSTM models are used for this purpose:

Data Collection: Historical data on charging station occupancy is collected. This data typically includes information such as recorded charging sessions including timestamps, current occupancy levels, day of the week, and any other relevant features.

Data Preprocessing: The collected data is preprocessed to make it suitable for training the LSTM model. This may involve steps such as normalization, handling missing values, and feature engineering to extract meaningful features.

Sequence Formation: The preprocessed data is divided into sequences of fixed length. Each sequence contains historical information on charging station occupancy and other relevant features. The target variable for each sequence is the occupancy level at the next time step.

Model Training: The LSTM model is trained using the prepared sequences of data. During training, the model learns to capture temporal patterns and dependencies in the data, enabling it to make predictions about future occupancy levels based on past observations.

Validation and Testing: The trained LSTM model is validated using a separate validation dataset to ensure that it generalizes well to unseen data. Additionally, the model's performance is evaluated on a test dataset to assess its predictive accuracy.

Prediction: Once the LSTM model has been trained and validated, it can be used to make predictions about future charging station occupancy levels. Given historical data as input, the model generates predictions for future occupancy levels at each time step.

Monitoring and Adaptation: The LSTM model's predictions are monitored in real-time, and the model may be retrained periodically using updated data to improve its accuracy and adapt to changing patterns in occupancy.

Long Short-Term Memory (LSTM) networks, as a special structure of Recurrent Neural Networks, have proven to be stable and powerful for modeling long-range dependencies in general-purpose sequence modeling ([9]-[11]). In LSTMs, each node in the hidden layer is replaced by a memory cell, instead of a single neuron [9]. The structure of a single memory cell is depicted in the figure below.

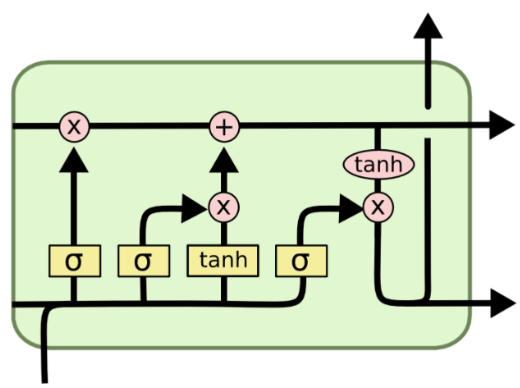

Figure 1: Structure of an LSTM Memory Cell

The memory cell contains the following components: the forget gate, the input node, the input gate, and the output gate. Each component applies a non-linear relation on the inner product between the input vectors and respective weights (altered iteratively through a training process). Some of the components have the sigmoid function, $\sigma(\cdot)$ and others the $\tanh(\cdot)$

As discussed in [6] Recurrent neural networks and LSTMs in particular, have shown great success in predicting time series online. Especially in [9] LSTMs have been used to tested, particularly on predicting traffic flows.

The goal of the forget gate is to decide what information should be discarded out of the memory cell [10]. The output, denoted as f(n) ranges between 0 and 1, according to the sigmoid activation function. The forget gate learns whether a previous or future vector state is necessary for the estimation of the current value state. The input node performs the same operation with that of a hidden neuron of a typical recurrent regression model. The goal of this node is to estimate the way in which each latent state variable contributes to the final model.

As far as the input gate is concerned, its role is to regulate whether the respective hidden state is sufficiently important. It has the sigmoid function, therefore its response ranges between 0 and 1. This gate addresses problems related to the vanishing of the gradient slope of a tanh(·) operator. Finally, the output gate regulates whether the response of the current memory cell is sufficiently significant to contribute to the next cell. Therefore, this gate actually models the long-range dependency together with the forget gate.

The recurrent nature of the LSTM presents many intricacies in terms of the iterative training process for adjusting the weights of the multiple gates. The adaptation of the backpropagation algorithm for accommodating the LSTM training is called Backpropagation Through Time [11]. The backpropagation variation for training recurrent neural network architectures presents the problem of vanishing or exploding gradients. So the number of time steps that the gradient is propagated is another hyperparameter of training that needs to be monitored. This adaptation is called truncated backpropagation through time and is thoroughly explained in [12].

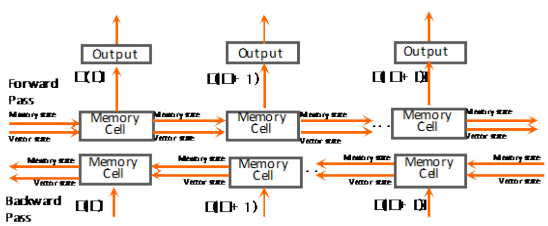

Figure 1: LSTM Architecture

2.2 Deep Reinforcement Learning Agent

Deep Reinforcement Learning (DRL) is a subfield of machine learning that combines reinforcement learning (RL) with deep learning techniques. It's particularly powerful for solving complex decision-making problems in environments with high-dimensional state and action spaces, such as playing video games, robotic control, and autonomous driving and in our case household energy management.

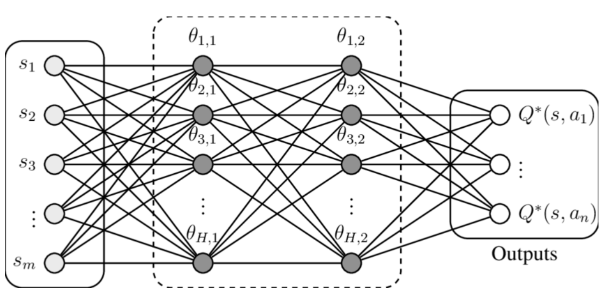

Figure 2: Deep Reinforcement Learning architecture.

A household energy management agent using reinforcement learning (RL) is an intelligent system designed to optimize energy consumption and generation within a home environment. Here's how such an agent might operate:

Agent: The agent, often implemented as a software program or embedded system, is responsible for making decisions related to energy management within the household. It interacts with various devices, appliances, and energy sources to achieve specific goals, such as minimizing electricity bills, reducing carbon footprint, or ensuring uninterrupted power supply.

Environment: The environment represents the household energy system, including energy-consuming devices (e.g., lights, HVAC systems, appliances), energy storage systems (e.g., batteries), renewable energy sources (e.g., solar panels, wind turbines), and the electrical grid. The environment provides feedback to the agent based on its actions and the state of the household energy system.

State: At each time step, the environment is in a particular state, which includes information about energy demand, energy prices, weather conditions, battery state of charge, and the status of appliances. The agent perceives this state and decides on actions based on it.

Action: The agent selects actions that affect energy consumption and generation within the household. Actions may include scheduling appliance usage, adjusting thermostat settings, controlling the charging/discharging of energy storage systems, and managing interactions with the electrical grid (e.g., buying or selling electricity).

Reward: After taking an action in a specific state, the agent receives feedback from the environment in the form of a reward. The reward reflects how well the action aligns with the agent's objectives, such as minimizing energy costs, maximizing self-consumption of renewable energy, or maintaining comfort levels within the home.

Policy: The agent follows a policy, which defines its strategy for selecting actions based on the current state. The goal of the agent is to learn an optimal policy that maximizes cumulative rewards over time while satisfying constraints such as comfort preferences and device operation requirements.

Exploration vs. Exploitation: The agent must balance exploration (trying new energy management strategies to discover their effectiveness) and exploitation (leveraging known effective strategies to maximize short-term rewards) to achieve optimal performance.

Learning Algorithm: The agent learns to improve its decision-making over time using an RL learning algorithm. This could include algorithms like Q-learning, Deep Q-Networks (DQN), Policy Gradient methods, or Actor-Critic methods, tailored to the specifics of household energy management.

Training: During training, the agent interacts with the environment over multiple episodes or iterations. It adjusts its policy based on the received rewards and updates its internal parameters to improve performance. Training might involve simulations based on historical data or experiments in a real household environment.

Evaluation: Once trained, the agent's performance is evaluated on unseen data or in real-world scenarios to assess its effectiveness in

optimizing household energy management and achieving the desired objectives, such as cost reduction, energy efficiency improvement, or environmental impact reduction.
By employing reinforcement learning techniques, household energy management agents can adapt to the dynamic nature of energy consumption and generation within homes, leading to more efficient energy usage, cost savings, and reduced environmental impact.

2.3 Simulation Environment

In reinforcement learning, the agent learns by making actions in an environment, receiving rewards/penalties for those actions and modifying its action pattern (policy) accordingly. One could either use the real environment (the real household and the energy markets) or simulated environment. In this project, a simulation environment was built, using the OpenAI Gym as a framework.

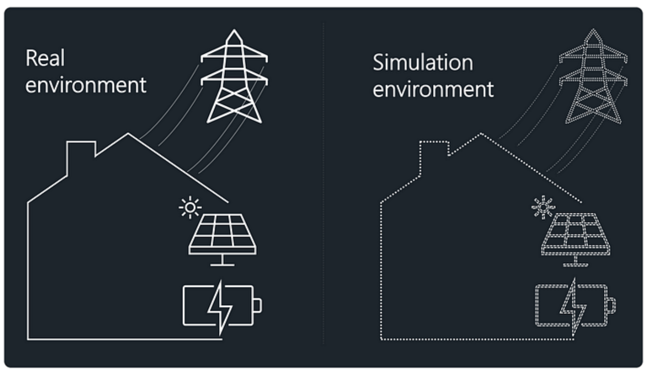

Figure 4: The simulation environment is a digital representation of the real physical environment that is used for training the reinforcement learning agent.

Developing the simulation environment for the energy decision agent involved the following steps:

- Representing the current situation (the state) of the environment numerically. The state representation needs to involve all of the information that our agent needs for training.
- Representing the possible actions (things that our agent can make with our battery) numerically. Representing and updating the current state of the home battery.
- Developing a logic of what it means to make a specific action in the environment. In other words, expressing what would happen with the household, battery and the energy cost when specific action is carried out by the agent.

3 System Architecture

The basic premise of the methodology for the Household Energy Management assistant is to model the consumption pattern for each household using time series data of the daily activity for each basic electric devices that contribute to the overall daily energy consumption. In that fashion we are going to take into account the preferences of the users. The collection of this kind of data has to be done through energy smart meters that can precisely measure the consumption with timestamps. The modeling of each household will be done by feeding the collected data into our already tested LSTM-based time series forecasting algorithm which was mentioned previously.

Next step is to design the proper simulation environment for our Deep Reinforcement Learning agent. There will be 3 versions of simulation environments as there are 3 different household setups:

a) a household that produces, stores and consumes energy (PV & battery)
b) a household that produces and consumes energy (PV)
c) a household that consumes energy

Each setup a completely different set of states, actions and logic that the agent has to follow in order to learn the best policy for the household energy management. Of course, the consumption pattern for each household that has been created will be inside the logic that the RL agent will follow.

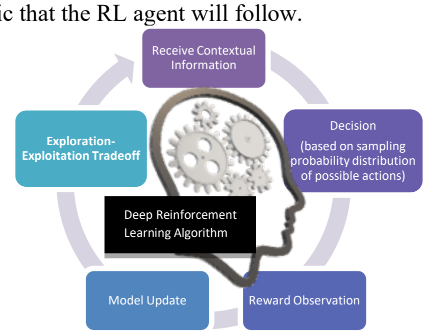

Figure 3: Workflow of the proposed method.

LSTM architecture:

The LSTM neural network architecture is comprised by one input layer, one output layer and two hidden layers with 50 neurons each (dense formulation). The Loss function used for adapting the weights is the Mean Square Error (MSE) which is the most typical loss function used for training in regression problems [14] and the optimization scheme is the ADAM optimizer [15]. The Backpropagation Through Time (BPTT) was stopped at three consecutive steps going back so the truncated version of the Backpropagation scheme was implemented for avoiding vanishing gradients.

4 Performance Evaluation

4.1 AI4CS Pilot

The LSTM-based Charging Station Occupancy Prediction was tested in pilot which was conducted in Collaboration with Hrvatski Telekom [25], [26]. Hrvatski Telekom operates the largest EV charging station network in Croatia and provided us with rich historical data on the charging sessions for each charging station for the last 5 years. During the pilot phase we had daily access to new data to fine tune our model. Due to our participation in Interconnect Horizon Project, we had also access to energy grid information coming from the Interoperable Recommender which is an innovative tool that has direct access to energy grid information related to the energy mix produced

(Fossil fuel or renewable energy) and the stress of the energy grid in an hourly basis. The Interoperable Recommender provides a signal on request about promoting consumption when there is enough renewable energy production, and the grid is not stressed or avoiding consumption when the energy is not green and the grid is stressed. For the pilot in Croatia, we built a mobile application for the benefit of EV owners which takes into consideration their preferred time and driving distance they are willing to drive and our application makes suggestions on where to go to find a charger that is more likely to be unoccupied and the energy coming from the grid is greener. The application will also make suggestions that might be time shifted for one or two hours from the selected time if the grid is stressed and the customer is not likely to find a close charger available. Hrvatski Telekom has shown interest in our application, and we are in close contact with them for integrating some of the functionalities to their app's backend.

4.2. AI4CS Pilot Results

The pilot in Collaboration with Hrvatski Telekom had a duration of more than a month and the outcome was remarkable. In the pilot phase 112 HT subscribers participated as Friendly Users and we collected all the information from their interactions with the app such as:

- Engagement of consumers and usage of app
- Collection of App Logs – Calculation of KPIs
- Collection of User Feedback for Experience Evaluation

The application got a User Satisfaction score of 3.8 out of 5 and a NPS score of 7.5/10 which means that users were likely to suggest the app to others. The app also contributed in the grid decompression since it successfully suggested alternative timeslots when needed and also increased the green electricity usage.

Table 1: AI4CS Pilot Key Results

Demonstration Project Key Performance Indicators		
Users Engagement	112 users	480 Recommendations
User Satisfaction	MOS Score 3.8	1:bad, 2:poor, 3:OK, 4:very good, 5:excellent
NPS score	7.5	1:Not Recommend at all – 10:Strongly Recommend
Avoidance of Grid Peak Load	130 Recommendations to avoid Grid Peak Load (based on Interoperable Recommender)	2.1% Contribution Reduction
CO_2 emissions reduction	Increase of Green Electricity Usage	2% Increase (est.)

Below you can see how many times the app proposed a different timeslot and the EV owner accepted it. And the pie chart of the user satisfaction scoring.

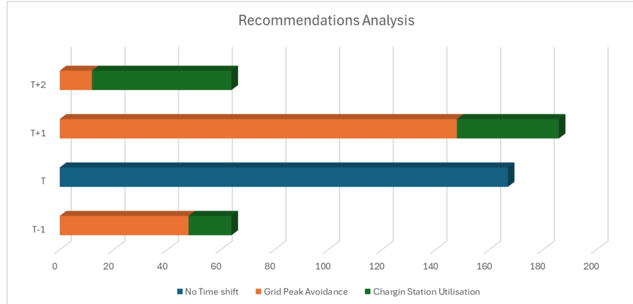

Figure 6: Shifting the time of EV charging to avoid Grid Load based on Interoperable Recommender Insights

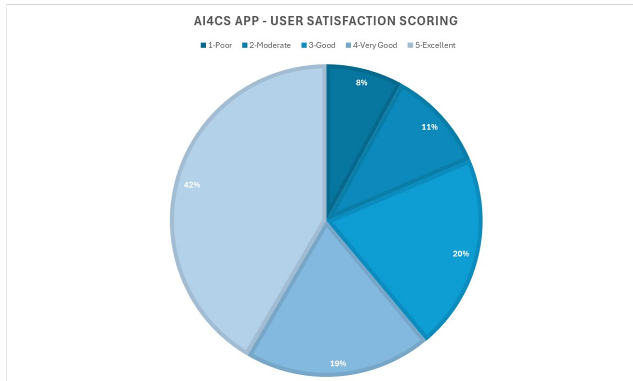

Figure 7: AI4CS APP – User Satisfaction Scoring

We have conducted experiments to: 1) validate the efficacy of the deep learning predictive model idea for Charging Station Occupancy Prediction and 2) compare the two architectures of AI in terms of accuracy, one is the LSTM architecture that we propose and the other is a classical ARIMA approach.

We train each of the formulations of AI networks (always as a regressor) with the same dataset that has been derived from the preprocessing of the historical data in charging sessions that Hrvatski Telekom provided us.

The results, as expected, shown a superior performance of the LSTM architecture compared to the classic ARIMA machine learning approach. We present the results

Table 1: Performance metrics for AI predictive models. The proposed LSTM approach outperforms the ARIMA model.

Table 2: Comparative Performance of the model

Neural Network architecture	Train MSE	Test MSE
LSTM	0.0224	0.334
ARIMA	10.55	13.89

5 Conclusion

In this paper we propose that a cooperation of two deep neural networks, one for time series forecasting and one for making recommendations on the household energy usage, plus information coming from the grid (related to the hourly energy mix and the stress level) can be used to create a smart household energy management app. managing energy consumption in a household is crucial for several reasons. First, Efficient energy use can lead to lower utility bills, saving money for the household budget. By being mindful of energy consumption, individuals can reduce wasted energy and optimize usage, thus reducing their overall expenses. Additionally, Energy production often involves the burning of fossil fuels, which releases greenhouse gases into the atmosphere. By managing energy consumption in a smart way like promoting more consumption when there is plenty of green energy production, households can decrease their carbon footprint and contribute to mitigating climate change and environmental degradation. Finally, it promotes a stable and reliable Grid. During peak demand periods, strain on the electrical grid can lead to blackouts or brownouts. By managing energy consumption, households can help alleviate pressure on the grid, contributing to overall grid stability and reliability.

ACKNOWLEDGMENTS

This research was partially supported by funding from the H2020-IOT-2016-2017 (H2020-IOT-2017) Program under Grant Agreement 780139 for the SerIoT Research and Innovation Action.

REFERENCES

[1] E. Gelenbe, "Random neural networks with negative and positive signals and product form solution,"Neural Computation, 1(4), 502–510, 1989.

[2] E. Gelenbe, "Learning in the recurrent random neural network.," Neural Computation, 5(1), 154–164, 1993

[3] E. Gelenbe and Y. Yin, "Deep learning with random neural networks," Neural Networks (IJCNN), 2016 Int. Joint Conference on, IEEE, 1633–1638, 2016.

[4] O. Brun, Y. Yin and E. Gelenbe, "Deep Learning with Dense Random Neural Networks for detecting attacks against IoT-connected home Environments," Procedia Computer Science, 134, 458-46 2018. https://doi.org/10.1016/j.procs.2018.07.183

[5] X. Yuan, C. Li and X. Li, "DeepDefense: Identifying DDoS Attack via Deep Learning," *2017 IEEE International Conference on Smart Computing (SMARTCOMP)*, Hong Kong, 2017, pp. 1-8. doi: 10.1109/SMARTCOMP.2017.7946998

[6] T. Guo, Z. Xu, X. Yao, H. Chen, K. Aberer and K. Funaya, "Robust Online Time Series Prediction with Recurrent Neural Networks," 2016 IEEE International Conference on Data Science and Advanced Analytics (DSAA), Montreal, QC, 2016, pp. 816-825.

[7] S. Basterreich and Gerardo Rubino. "Random Neural Network Model for Supervised Learning Problems: A Tutorial," Neural Network World 25.5 (2015): 457–499. Crossref. Web.

[8] T. Guo, Z. Xu, X. Yao, H. Chen, K. Aberer and K. Funaya, "Robust Online Time Series Prediction with Recurrent Neural Networks," *2016 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*, Montreal, QC, 2016, pp. 816-825

[9] Y. Tian and L. Pan, "Predicting Short-Term Traffic Flow by Long Short-Term Memory Recurrent Neural Network," *2015 IEEE International Conference on Smart City/SocialCom/SustainCom (SmartCity)*, Chengdu, 2015, pp. 153-158. doi: 10.1109/SmartCity.2015.63D.I.,.

[10] K. Greff, R. K. Srivastava, J. Koutník, B. R. Steunebrink and J. Schmidhuber, "LSTM: A Search Space Odyssey," in *IEEE Transactions on Neural Networks and Learning Systems*, vol. 28, no. 10, pp. 2222-2232, Oct. 2017. doi: 10.1109/TNNLS.2016.2582924.

[11] P. J. Werbos, "Backpropagation through time: what it does and how to do it," in *Proceedings of the IEEE*, vol. 78, no. 10, pp. 1550-1560, Oct. 1990. doi: 10.1109/5.

[12] H. Tang and J. Glass, "On Training Recurrent Networks with Truncated Backpropagation Through time in Speech Recognition," *2018 IEEE Spoken Language Technology Workshop (SLT)*, Athens, Greece, 2018, pp. 48-55. doi: 10.1109/SLT.2018.8639517

[13] doi: 10.1109/ICCSN.2011.6014810

[14] W. Zhang *et al*., "LSTM-Based Analysis of Industrial IoT Equipment," in *IEEE Access*, vol. 6, pp. 23551-23560, 2018. doi: 10.1109/ACCESS.2018.2825538

[15] Z. Zhang, "Improved Adam Optimizer for Deep Neural Networks," *2018 IEEE/ACM 26th International Symposium on Quality of Service (IWQoS)*, Banff, AB, Canada, 2018, pp. 1-2. doi: 10.1109/IWQoS.2018.8624183

[16] Lan Wang and E. Gelenbe, "Adaptive dispatching of tasks in the cloud," IEEE Transactions on Cloud Computing 6 (1), 33-45, 2018.

[17] Aurora Energy Research, "European Battery Markets Attractiveness Report" - February 2024

[18] Boston Consulting Group (BCG) and Google, "How AI Can Speed Climate Action," November 20, 2023 Reduce Carbon and Costs with the Power of AI | BCG

[19] Nikos Tsafos, Chief Energy Adviser to Greek PM, "Greek Energy Transition Outlook", Delphi Economic Forum April 2024.

[20] ADMIE-IPTO, "Greek Electricity Production Mix - Monthly Energy Report", March 2024

[21] Rainer Stamminger, Rheinische Friedrich-Wilhelms-Universität BonnUniversity of Bonn, "Synergy Potential of Smart Appliances - EU Average, typical household Study", 2023 http://www.come-on-labels.eu/download-library/synergy-potential-of-smart-appliances

[22] ENTSO-E Transparency Platform, https://transparency.entsoe.eu/

[23] Electricity Maps "Electricity consumption by source", "Hourly Csrbon Intensity" Dashboard, https://app.electricitymaps.com/map

[24] Greek Energy Exchange – enex, Hourly Production/Consumption/Price https://www.enexgroup.gr/

[25] Interconnect EU Horizon Project Report, https://interconnectproject.eu/wp-content/uploads/2023/12/Booklet.pdf

[26] AI for Charging Stations Project, InterConnect Demo Open Call, AI4CS App Demo Video

[27] *Github Repo for forecasting models of the AI for Charging Stations project,* *https://github.com/local-ai-gr/CNROpt/blob/main/Forecasting%20with%20location%20included%20final%20model.ipynb*
[28] Gymnasium open source Python library for reinforcement learning algorithms, https://gymnasium.farama.org/index.html