University of Maribor Press
DOI: https://doi.org/10.18690/um.fov.5.2026.3
ISBN: 978-961-299-160-9

Supervised Machine Learning for Renewable Energy

Dijana Oreški,
Vjeran Strahonja, and
Marija Pokos Lukinec

Abstract
Accurate renewable energy forecasting is important for optimizing grid integration and advancing environmental sustainability. This chapter develops predictive models based on supervised machine learning for solar energy consumption using historical data from solar power plants, integrating various data sources: historical energy consumption, actual weather conditions (including temperature, insolation, and wind speed), and historical weather forecasts. Advanced artificial intelligence and machine learning algorithms including deep learning were trained on a multi-source dataset to identify complex temporal patterns and weather-energy patterns. The models achieved high precision, demonstrating robustness against meteorological variability. Accurate predictive models enable utilities to reduce fossil-fuel-based reserve capacity, minimize grid inefficiencies, and enhance renewable energy utilization. For environmental sustainability, these models directly support decarbonization goals by enabling larger solar integration, reducing associated carbon emissions from backup generation, and promoting resource-efficient energy planning. By facilitating the reliable and efficient integration of solar power, this approach represents a small step toward achieving zero net emissions in the energy sector.

Keywords: renewable energy; power plant; machine learning; forecasting; environmental sustainability

Dijana Oreški
University of Zagreb, Faculty of Organization and Informatics, Pavlinska 2, 42000 Varaždin, Croatia
e-mail: dijana.oreski@foi.unizg.hr
Vjeran Strahonja
University of Zagreb, Faculty of Organization and Informatics, Pavlinska 2, 42000 Varaždin, Croatia
e-mail: vjeran.strahonja@foi.unizg.hr
Marija Pokos Lukinec
University of Zagreb, Faculty of Organization and Informatics, Pavlinska 2, 42000 Varaždin, Croatia
e-mail: mapokos@foi.hr

1 Introduction

The rapid global transition toward renewable energy sources is not only a strategic imperative for mitigating climate change and reducing greenhouse gas emissions but also a cornerstone for achieving long-term environmental and economic sustainability. Solar, wind, and hydro power technologies provide clean energy, yet their variability and uncertainty introduce operational challenges for modern power systems. Solar energy, in particular, has emerged as one of the fastest-growing and most promising resources due to its modularity, scalability, and declining costs. However, its inherently changeable nature - caused by meteorological fluctuations, and seasonal variability - complicates reliable grid integration and real-time operational planning. These challenges are further amplified by the accelerating global demand for electricity, driven not only by population growth and electrification of transport and industry but also by the proliferation of energy-intensive digital technologies such as large language models (LLMs) and AI-driven data centers. As the share of variable renewable energy sources rises, the ability to accurately forecast solar generation or consumption and optimize system flexibility becomes increasingly important for reducing reliance on fossil-fuel-based backup capacity, minimizing grid inefficiencies, and supporting the broader goal of deep decarbonization.

Accurate forecasting of solar consumption is therefore critical to ensure stable electricity supply, reduce reliance on fossil-fuel-based backup systems, and optimize resource allocation. Short-term and mid-term forecasts allow grid operators to better schedule generation, minimize balancing costs, and improve the efficiency of energy markets. However, accurate prediction is challenged by the fact that renewable energy sources are inherently dependent on external, often unpredictable, factors such as weather conditions. Similarly, energy consumption patterns are influenced by user behavior, time of day, seasonality, and socio-economic factors. Traditional forecasting methods often struggle to capture these nonlinear and dynamic relationships.

Recent advances in artificial intelligence (AI) and machine learning (ML) have opened new opportunities for developing highly accurate and adaptive forecasting models in the energy sector. Unlike traditional rule-based or statistical methods, which rely on fixed assumptions and limited analytical flexibility, ML algorithms can autonomously learn complex patterns from large volumes of historical data and continuously adapt to emerging trends. This makes them suitable for both energy generation and consumption prediction, where nonlinear interactions and temporal dependencies play a critical role. Deep learning architectures, for example, can capture intrinsic relationships among diverse variables such as solar irradiance, ambient temperature, wind speed, and historical consumption profiles, providing predictive capabilities far beyond those of conventional approaches [1]. Furthermore, the integration of multi-source datasets - combining measured solar output, real-time meteorological data, and weather forecast information - significantly improves model robustness against meteorological variability. By enabling data-driven decision-making for grid stability, load balancing, and energy efficiency, ML-based forecasting methods provide a foundation for optimizing renewable energy integration and advancing sustainability objectives.

This chapter develops and evaluates advanced predictive models based on machine and deep learning (DL) for solar energy consumption with the aim of improving grid reliability and supporting environmental sustainability objectives. By leveraging DL and other state-of-the-art ML techniques on a rich, multi-source data set, we demonstrate how powerful AI algorithms in combination with quality data can develop accurate predictive models.

This chapter is organized as follows. Section 2 reviews the relevant literature and theoretical background and presents the methodology, detailing the Long Short-Term Memory (LSTM) and Time Series Transformer (TST) models along with a description of the heterogeneous data sources. Section 3 discusses the research results, while Section 4 highlights their implications for integration into battery optimization systems, along with a discussion of financial and environmental impacts. Finally, Section 5 concludes the research and outlines potential directions for future work.

2 Materials and methods

This section outlines the study objectives, summarizes related work to position our approach within the existing literature, and describes the methodology and datasets used to develop and evaluate the predictive models.

2.1 Research objectives

The research problem is motivated by a real-life problem faced by the solar energy company which installs solar power plants. These plants rely on sunlight, which is highly variable and influenced by local weather patterns. The company's goal is to maximize electricity generation while ensuring a reliable supply to the power grid. One major challenge is accurately predicting how much energy will be produced each day or hour. Bad forecasting can lead to overproduction or underproduction, which affects grid stability and profitability. The company needs accurate short-term and long-term forecasts to plan operations and schedule maintenance effectively. Because solar power is not always available when demand peaks, companies also rely on battery storage systems. They must decide when to store energy and when to sell it back to the grid - in a dynamic pricing environment.

Motivated by this real-life problem, our research aims to:

1. Develop accurate models for predicting renewable energy consumption. The first objective is to build machine learning models that can forecast energy consumption from renewable sources. These models need to account for variables like weather, season, location, and historical output. We develop next-day and multi-period energy-consumption forecasts by training on plant-level historical datasets that integrate three information streams: (i) historical consumption profiles, (ii) realized meteorological conditions: air temperature, solar irradiance/insolation, and wind speed and

(iii) historical weather forecasts for the same variables. Incorporating both realized weather and time-aligned forecast fields enables the models to learn from operationally available signals and to better approximate real-world forecasting conditions.
2. Evaluate and compare different machine learning algorithms. A variety of models can be used and this objective involves testing two algorithms to see which offers the best accuracy and generalization.

Ultimately, predictions should be actionable. The goal of such models is to use ML outputs to inform real-world decisions, such as when to charge/discharge batteries or how to best feed energy into the grid to reduce costs and improve system stability. This is discussed in discussion section.

2.2 Related works

The rapid growth of AI and ML applications in predictive modeling has resulted in an extensive and diverse body of research. Understanding the current state of knowledge requires examining not only the technical methods used but also the types of data on which these methods rely and the performance of the resulting models. This literature review is organized into three main areas. First, it investigates what data sources have been used to develop predictive models. Second, it examines what AI and ML algorithms have been applied, evaluating trends in algorithm selection and implementation. Finally, it analyzes the quality of the resulting predictive models, considering evaluation metrics. Together, these three perspectives provide an overview of how predictive modeling approaches have evolved and where research opportunities for improvement remain.

Recent literature highlights that diverse and multi-layered datasets are critical for developing accurate predictive models for solar energy forecasting and battery storage optimization [2]. Measurements such as irradiance, temperature, and energy production have shown the best performance. Models trained on these datasets achieve a low MAPE and high R^2. Additional categories of data include weather sensor readings, smart grid data, and energy consumption/load profiles, which are often integrated into models for battery energy storage system (BESS) optimization [3, 4]. Existing studies, however, typically rely on a single type of data or, at best, integrate two distinct data sources. No work has been identified that combines more than two heterogeneous sources of information to build predictive models in this domain. This limitation reduces the potential to capture complex relationships and richer contextual factors that could enhance predictive accuracy and robustness. Addressing this research gap represents one of the motivations for the present study.

Several studies have demonstrated that machine learning techniques can effectively model photovoltaic (PV) energy generation by leveraging meteorological variables. For example, [5] explores how algorithms such as artificial neural networks, support vector machines, and regression models can predict PV output based on inputs like temperature, irradiance, humidity, and wind speed. The authors highlight that incorporating multiple weather parameters improves the accuracy of energy forecasting compared to traditional empirical methods, supporting the adoption of data-driven approaches for optimizing PV system performance and grid integration. Additional literature demonstrates a wide range of machine learning, with recent papers emphasizing deep learning approaches superiority (e.g. LSTM and CNN-LSTM in [6]). Although a variety of machine learning algorithms have been applied, there remains a lack of comprehensive comparative analyses as well as insufficient exploration of the potential of generative models.

The quality of ML and DL models is typically assessed using both classical and advanced performance metrics such as RMSE, MAE, MAPE, and R^2.

2.3 Data description

This study utilizes a comprehensive dataset comprising solar energy consumption data from multiple photovoltaic power plants equipped with different monitoring systems. The primary data sources include power plants from three major manufacturers: SolarEdge (12 power plants), Fusion Solar/Huawei (25 power plants), and iSolarCloud (3 power plants). Each monitoring system presents unique characteristics in terms of data accessibility, measurement intervals, and available metrics, necessitating a systematic approach to data standardization.

The dataset encompasses three main categories of variables: (1) historical energy consumption and production data from solar power plants, (2) actual meteorological conditions, and (3) weather forecast data. As such, this dataset integrates power plant operational data with relevant meteorological and temporal information. All power plant data, along with corresponding auxiliary datasets, are recorded and utilized as standardized 15-minute interval values. Within the database, each time interval is stored in the format YYYY-MM-DD HH:MM:SS (e.g., 2025-02-26 12:15:00). The temporal intervals range from 00:00:00 (the beginning of the day, representing the first interval for a given date) to 23:45:00 (the final interval of that date). Each timestamp is preceded by the corresponding calendar date, ensuring temporal consistency across all records.

The dataset comprises the following groups of features:

- Consumption – Represents the historical energy consumption aggregated from all available power plants.
- Actual weather conditions – Include measured meteorological parameters such as temperature, insolation, and wind speed.
- Weather forecasts – Encompass past, present, and future forecasted meteorological data. The observed parameters correspond to those in the actual weather conditions dataset, namely forecastTemperature, forecastInsolation, and forecastWindSpeed.
- Holiday – A binary indicator (0 or 1) specifying whether a particular day is a public holiday. A value of 0 denotes a regular (non-holiday) day, while 1 indicates a public holiday.
- TimeOfDay – Refers to 15-minute data retrieval intervals (timestamps), represented numerically and appropriately scaled or transformed.

- DayOfWeek – Numerical representation of the day of the week, suitably scaled or transformed.
- DayOfYear – Numerical representation of the day within the year, also scaled or transformed.
- Correction – A binary indicator specifying the validity of the consumption data obtained from power plants. If irregularities occur (e.g., missing or NULL values), such data are considered invalid. For these intervals, the correction flag is set to 1, and the consumption value is recorded as 0 instead of NULL. Conversely, if the consumption data are valid, the actual measured values are used, and the correction flag remains 0 for those intervals.

The process of preparing the dataset for model training was carried out in two main phases: (1) generation of the base dataset and (2) final data transformation and scaling.

In the initial phase, the raw data obtained from power plants and accompanying meteorological sources were preprocessed and merged into a unified base dataset. This phase included temporal partitioning of the data into consistent 15-minute intervals, ensuring complete temporal coverage and alignment of all features. Subsequently, feature engineering was applied to enhance the informativeness of the dataset. Based on the previously described raw variables, the script automatically generated a set of derived features designed to facilitate pattern recognition by the model. These newly created features capture temporal, seasonal, and weather-related dynamics, thereby improving the model's ability to generalize and detect relevant correlations. Before model training, additional transformations were applied to ensure numerical stability and comparability across features. A logarithmic transformation was performed to reduce skewness in variables with large variance and to stabilize the data distribution. Finally, a global scaling procedure was implemented to normalize all input variables to a common range suitable for the selected learning algorithms. After model inference, predicted values were transformed back to their original scale to allow for accurate interpretation and evaluation against real-world measurements.

2.4 Deep learning algorithms

Our investigation covered a broad spectrum of approaches, including: (i) machine learning algorithms, such as Random Forests, to establish baseline performance, (ii) convolutional neural networks (CNN, RCNN) for capturing local temporal patterns and feature dependencies, (iii) recurrent neural networks (RNN), particularly Long Short-Term Memory (LSTM) and Gated Recurrent Units (GRU), for modeling sequential dependencies and (iv) transformer-based architectures, including the Time Series Transformer (TST), Informer, and specialized variants such as PatchTST with residual connections for enhanced long-range dependency modeling.

Among these, two algorithms (LSTM and TST), showed high level of performance and those algorithms are described below and applied in this research.

Long Short-Term Memory Networks (LSTM) are a class of recurrent neural networks specifically designed to overcome the limitations of traditional RNNs [7], particularly the vanishing gradient problems during long sequence training. By introducing gated mechanisms - input, forget, and output gates - LSTMs regulate the flow of information through a cell state, enabling them to capture both short- and long-term dependencies in sequential data. This makes LSTMs well-suited for time series forecasting tasks, where patterns may span multiple temporal scales [8]. In this study, LSTMs were employed with varying numbers of layers and hidden units, as well as different regularization schemes, to assess their ability to generalize across complex temporal dynamics.

The Time Series Transformer (TST) adapts the self-attention mechanism of Transformer architectures [9], originally developed for natural language processing, to handle time series data [10]. Unlike recurrent models, which process data sequentially, TSTs operate in parallel, allowing efficient modeling of long-range dependencies without the constraints of recursion. Temporal positional encodings were incorporated to preserve the ordering of observations, while multi-head attention layers learned which past time steps were most relevant for future predictions. Residual connections and normalization techniques were applied to improve stability and convergence during training. The model's flexibility and scalability make it a strong candidate for high-dimensional and multivariate time series forecasting.

3 Research results

This section presents the results of developing and evaluating prediction models for energy consumption. The focus is on assessing predictive performance, analyzing the impact of model architecture and hyperparameter optimization, and validating results through quantitative metrics. The following subsections describe the model design, training process, optimization strategy, and key findings, highlighting the performance of the predictive models.

3.1 Model architecture

The best prediction model is based on the TST, which employs only the encoder component of the original Transformer architecture. This approach proved most effective for capturing complex temporal dependencies in the data. The data flow through the model is structured as follows:

- Input Projection – The prepared input vectors (16 input features for prediction models) are first passed through a linear layer. This layer projects the input vectors from their original lower-dimensional space into a higher-dimensional latent space (d_model), enabling the model to operate in a richer representational domain.
- Positional Encoding – Since the Transformer architecture is inherently permutation-invariant, explicit temporal order must be introduced. Sinusoidal and cosine positional encodings are added to the input vectors, assigning a unique temporal signa-

ture to each point in the 7-day input sequence. This allows the model to understand the sequential structure of the data.

- Transformer Encoder – The core of the model consists of a stack of identical encoder layers (e.g., three layers for consumption forecasting). Each layer includes two main components:
  - Multi-Head Self-Attention – This mechanism enables the model to examine the entire 7-day input sequence simultaneously and determine which time steps are most relevant to each other. Multiple attention heads operate in parallel, with each head capturing different temporal patterns (e.g., daily vs. weekly dependencies).
  - Feed-Forward Network – Following the attention step, each time step is independently processed through a small fully connected neural network to enrich its representation. Residual connections and layer normalization are applied throughout the encoder to ensure stable and efficient training of the deep architecture.
- Information Aggregation and MLP Head – After passing through all encoder layers, only the output corresponding to the last time step of the input sequence is retained. This vector effectively summarizes all relevant information from the preceding 7 days. It is then fed into a multi-layer perceptron (MLP) head consisting of one hidden layer with a ReLU activation. The MLP generates the final forecast for the next 2 days (192 time steps). Compared to a single linear output layer used in earlier versions of the model, the MLP head enables a more expressive nonlinear mapping between the learned temporal representations and the final prediction.

3.2 Hyperparameter optimization

The model training process was guided by optimization tools and hyperparameter configurations to ensure robust convergence and generalization. The AdamW optimizer was employed as a modern and robust variant of Adam, incorporating weight decay directly into the update rule. This improves regularization and often results in better overall performance compared to the standard Adam optimizer. Instead of a fixed learning rate, the OneCycleLR scheduler dynamically adjusts the learning rate throughout training. It begins with a low value, gradually increases to a predefined maximum, and then decreases toward the end of training. This cyclical schedule accelerates convergence and helps the model escape suboptimal local minima. SmoothL1Loss (Huber Loss) was used to balance sensitivity and stability. The loss behaves quadratically for small errors (similar to MSE) and linearly for large errors (similar to MAE), making it less sensitive to sudden spikes and outliers. This encourages the model to handle abrupt changes in consumption without being destabilized by excessive penalty values. Key architectural and regularization parameters were tuned for the consumption models, including the number of encoder layers (NUM_ENCODER_LAYERS), model dimensionality (D_MODEL), number of attention heads (N_HEADS), and regularization factors (DROPOUT_RATE, WEIGHT_DECAY). This hyperparameter optimization strategy ensured that each model variant was appropriately tailored to the specific forecasting challenge, contributing to improved stability and predictive performance.

3.3 Performance evaluation and comparative analysis

Model training was conducted over a predefined number of epochs, with progress carefully monitored to ensure both convergence and generalization. Three key metrics were evaluated during each epoch, separately for the training and validation sets:

- Loss – the value of the loss function, serving as the primary signal for the optimizer.
- MAE (Mean Absolute Error) – the main evaluation metric, representing the average absolute error in kWh and directly indicating how much the model deviates from actual values on average.
- wMAPE (weighted Mean Absolute Percentage Error) – a relative error metric expressed as a percentage of the true value, allowing for performance comparisons across power plants of varying sizes.

Model selection was based on the MAE calculated on the validation set. After each epoch, if the current MAE was lower than all previous values, the model was considered ”best so far,” and its weights were stored for later use. To prevent overfitting and unnecessary computation, an early stopping mechanism was applied. Training was automatically stopped if no improvement in MAE was observed for a predefined number of consecutive epochs (e.g., 10), using a patience parameter (PATIENCE_EARLY_STOP). The best-performing model for consumption prediction was the Time Series Transformer (TST), which achieved an MAE of 3.80 on the validation set.

Figure 1 illustrates an example of model predictions on the validation and test sets, where the actual values of consumption are compared with the corresponding forecasts. The results show that the model successfully captures the overall temporal patterns and trends, while certain discrepancies occur during rapid fluctuations, where peak values are either underestimated or overestimated. This visual inspection complements the quantitative error metrics and provides an intuitive understanding of the model's predictive performance.

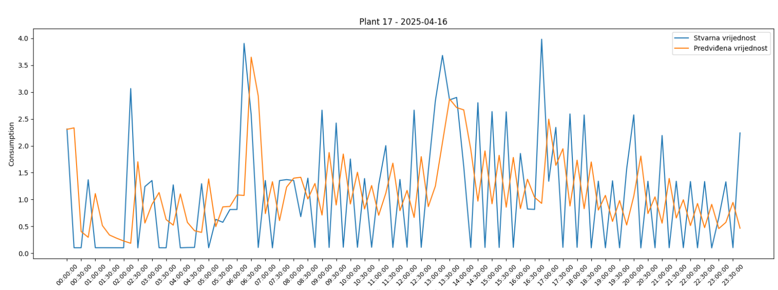

Fig. 1: Actual versus predicted values.

At the final stage of the analysis, a Permutation Feature Importance (PFI) procedure was conducted. This approach quantifies the contribution of each input feature by measuring the deterioration in model performance - expressed as an increase in the Mean

Absolute Error (MAE) - that occurs when the values of a specific feature are randomly permuted. Features causing the largest degradation in performance are interpreted as being the most influential for the model. The results of this analysis are illustrated in figure 2 for a representative example.

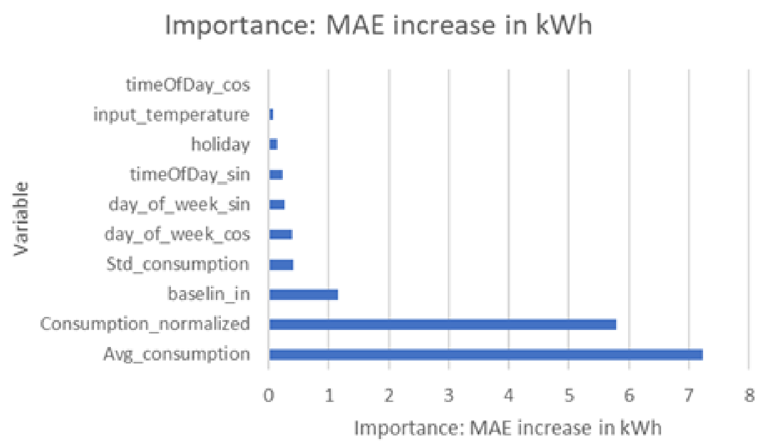

Fig. 2: Feature importance.

Figure 2 shows the permutation feature importance for the consumption prediction model, measured as the increase in MAE (kWh) when each feature is randomly shuffled. The results indicate that avg_consumption_in and consumption_normalized are by far the most influential variables, contributing 7.24 kWh and 5.79 kWh to prediction accuracy, respectively. Secondary contributors include baseline_in and std_consumption_in, while temporal features such as day-of-week and time-of-day encodings have moderate importance. Weather-related variables (e.g., temperature, wind speed, irradiation) exhibit minimal impact, suggesting that short-term consumption patterns are driven primarily by historical load profiles rather than external environmental conditions.

4 Discussion

The research has successfully developed and validated deep learning models for solar energy consumption forecasting. Key achievements include:

- Methodological Innovation: Implementation of lagged interval training and weather data augmentation techniques.
- Architectural Advancement: Development of Transformer architecture optimized for renewable energy forecasting.
- Performance Excellence: Achievement of low prediction error rates across comprehensive test datasets.

The results demonstrate that advanced machine learning architectures, when properly adapted for renewable energy applications, can achieve high forecasting accuracy that exceeds traditional approaches. This level of precision has important implications for grid management, energy trading, and sustainable energy system optimization. The high forecasting accuracy achieved by the Transformer model opens significant opportunities for integration into intelligent battery management systems for solar installations. Hereinafter, we discuss economic and environmental sustainability impact of such approaches.

4.1 Economic impact

The low prediction error rate provides the precision necessary for optimal battery charging and discharging decisions, which are critical for maximizing economic returns and extending battery lifecycle, in following:

- Predictive Charging Strategies: Accurate consumption forecasts enable proactive battery charging during periods of excess solar generation.
- Load Balancing Optimization: Precise demand predictions facilitate optimal energy distribution between immediate consumption, battery storage, and grid export.
- Battery Lifecycle Management: Accurate forecasting reduces unnecessary charge-discharge cycles, potentially extending battery operational life.

To fully realize the potential of AI-driven battery optimization, the current consumption forecasting model must be integrated with complementary predictive systems:

- Electricity Price Forecasting: Machine learning models predicting short-term electricity market prices with similar accuracy to consumption forecasting.
- Solar Production Forecasting: Parallel models predicting energy generation to complement consumption predictions.
- Battery State Optimization: Models determining optimal state-of-charge levels based on forecasted consumption, production, and pricing conditions.
- Grid Interaction Models: Systems optimizing the timing and magnitude of grid energy purchases and sales.

Preliminary economic analysis suggests that AI-optimized battery management systems could deliver substantial financial benefits. Besides economic benefits, there are environmental sustainability impacts, discussed in the following subsection.

4.2 Environmental sustainability impact

The implementation of high-precision solar energy consumption forecasting delivers significant environmental benefits that align with global sustainability objectives and carbon reduction targets [12]. The model's accuracy directly contributes to reduced carbon emissions through several interconnected mechanisms that optimize renewable energy utilization while minimizing environmental impact.

Optimized solar energy utilization reduces dependence on carbon-intensive grid electricity, which is typically sourced from fossil fuel power plants. This fossil fuel displacement represents one of the most direct environmental benefits of accurate consumption forecasting. Additionally, precise demand forecasting enables better grid planning and reduces the need for peaking power plants, which often operate with higher emission rates due to their rapid response requirements. The technology ensures maximum utilization of available solar generation through precise consumption predictions, effectively minimizing energy waste and maximizing the environmental value of existing renewable infrastructure.

Based on typical solar installation performance and grid emission factors, the implementation of AI-optimized energy systems delivers quantifiable environmental improvements. These benefits include significant CO_2 emission reductions, substantial grid strain mitigation, and marked improvements in renewable energy efficiency. The cascading effects extend to broader resource conservation through optimized charging patterns that reduce battery degradation, decreasing the frequency of battery replacement and associated manufacturing emissions. This approach reduces grid infrastructure stress through intelligent load management while extending equipment lifecycles, thereby reducing the need for manufacturing replacement components. The AI-driven optimization system supports circular economy principles by minimizing energy waste through precise demand-supply matching, extending the operational lifespan of solar and battery systems, and maximizing the value extracted from existing renewable energy infrastructure. This holistic approach to resource utilization represents a paradigm shift from linear consumption models to circular resource optimization strategies that prioritize longevity and efficiency.

The widespread adoption of accurate consumption forecasting systems contributes meaningfully to broader climate change mitigation efforts by enhancing the economic viability of solar installations, which accelerates renewable energy adoption rates across diverse market segments [11]. Increased distributed renewable generation supports overall grid decarbonization while simultaneously reducing dependence on fossil fuel imports, thereby enhancing energy security and reducing geopolitical risks associated with energy dependency.

The technology promotes environmental sustainability in an equitable manner by improving the economic performance of residential solar systems, making clean energy accessible to broader populations regardless of economic status [10]. Reduced grid strain benefits entire communities through improved electrical system reliability, while localized reductions in fossil fuel consumption contribute to improved air quality in urban and industrial areas where air pollution disproportionately affects vulnerable populations.

Indirect environmental benefits extend beyond immediate energy applications to include reduced mining impact through extended battery lifecycles, which decrease demand for lithium, cobalt, and other battery materials while reducing associated mining environmental impacts [13]. The technology contributes to habitat preservation by reducing the need for new power plant construction, while decreased reliance on thermal power plants reduces water consumption for cooling systems, supporting broader water conservation objectives.

The environmental benefits scale significantly across multiple implementation levels. Individual residential installations deliver measurable household carbon footprint reductions, while neighborhood-level deployments create substantial local environmental improvements that compound across communities. Large-scale regional adoption contributes meaningfully to regional and national carbon reduction targets, and technology transfer and adoption in developing countries supports global climate objectives through accessible clean energy solutions.

This research directly supports multiple United Nations Sustainable Development Goals, including [14] (Goal 7: Affordable and clean energy), through enhanced solar system economics and performance that advance affordable and clean energy access. The work contributes to [15] (Goal 9: Industry, innovation and infrastructure) by demonstrating advanced AI applications in renewable energy infrastructure, supports [16] (Goal 11: Sustainable cities and communities) through improved urban energy management capabilities, and directly advances [17] (Goal 13: Climate action) through measurable contributions to carbon emission reductions and climate action initiatives.

The development of high-precision consumption forecasting represents an important technological building block for achieving a sustainable energy future. As the technology matures and scales, its environmental impact compounds through widespread adoption that creates cumulative environmental benefits across diverse applications and geographic regions. The success in solar forecasting serves as an innovation catalyst, accelerating the development of similar AI applications across renewable energy sectors while generating demonstrated environmental benefits that support policy frameworks favoring renewable energy adoption. These proven sustainability benefits attract increased investment in clean energy technologies, creating positive feedback loops that accelerate the transition to sustainable energy systems.

The environmental sustainability implications of accurate solar energy consumption forecasting extend far beyond individual installations, contributing to systemic changes in energy production, consumption, and management that collectively advance global environmental objectives while supporting the transition to a sustainable energy future.

5 Conclusions

This research demonstrates that deep learning, particularly advanced architectures such as TST, can enhance the accuracy of solar energy consumption forecasting, achieving precision levels that open new possibilities for intelligent energy management systems. The research reveals that while powerful algorithms are essential, the quality and consistency of underlying data play an equally crucial role in determining model performance. Without careful data preparation and temporal alignment across different monitoring systems, even the most advanced machine learning models cannot achieve their full potential. The evolution from LSTM networks to the final TST implementation illustrates the importance of architectural innovation tailored specifically to the unique characteristics of renewable energy data. The final model achieved high performance representing advancement over traditional forecasting approaches and positioning the system as ready for live operational testing in real-world environments.

Beyond the technical achievements, this work establishes the pathway for expanding AI-driven energy management ecosystems. The success of accurate consumption forecasting provides the foundation for more sophisticated applications, including the integration of electricity market price forecasting models and the development of comprehensive solar generation prediction systems. The convergence of consumption forecasts, production predictions, and dynamic pricing data creates opportunities for multi-objective optimization frameworks that can simultaneously optimize energy efficiency, cost-effectiveness, and environmental impact.

Several limitations must be acknowledged. The current results, while promising, are based on a relatively constrained dataset from a limited number of power plants, which may affect the generalizability of findings across diverse installation types and geographical regions. Additionally, restricted data access prevented the inclusion of all available monitoring systems. The simulation of historical weather forecasts, while a common practice in the field, may not perfectly capture the noise and uncertainty characteristics of real-world meteorological predictions.

Future research directions focus on addressing these limitations while expanding the scope and impact of the work. Dataset expansion through the inclusion of additional solar power plants and extended data collection periods will enhance model robustness and generalizability across diverse operational contexts. Integration of previously inaccessible systems will provide comprehensive coverage of available infrastructure and improve model validation across different manufacturer platforms. Most critically, the research plans point toward moving beyond prediction into optimization, with particular emphasis on intelligent battery storage management systems that leverage AI to determine optimal energy storage and release strategies.

The ultimate goal of the research includes comprehensive energy management ecosystems where accurate consumption forecasting serves as the foundation for sophisticated optimization algorithms that balance energy efficiency, cost-effectiveness, and environmental sustainability. This includes the development of real-time adaptation systems that continuously learn from changing operational conditions, standardized interfaces for utility grid interaction and compliance, and integrated platforms that simultaneously optimize multiple objectives across consumption, production, storage, and market participation. The success demonstrated in this research validates the potential for AI-driven solutions to address critical challenges in renewable energy management while contributing meaningfully to global sustainability objectives. As the technology matures and scales, its impact will extend beyond individual installations to support systemic changes in energy production, consumption, and management that collectively advance environmental sustainability and economic efficiency in the transition toward a clean energy future.

Generative artificial intelligence usage

During the preparation of this manuscript, the authors used ChatGPT, Claude, Perplexity, and Elicit for editing references, grammar and spelling checks, formatting and editing the text, and identifying relevant works in the field. The authors have reviewed and edited the output and take full responsibility for the content of this publication.

References

1. Benti, N. E., Chaka, M. D., & Semie, A. G. (2023). Forecasting renewable energy generation with machine learning and deep learning: Current advances and future prospects. *Sustainability*, *15*(9), Article 7087. https://doi.org/10.3390/su15097087
2. Suanpang, P., & Jamjuntr, P. (2024). Machine learning models for solar power generation forecasting in microgrid application implications for smart cities. *Sustainability*, *16*(14), Article 6087. https://doi.org/10.3390/su16146087
3. Bouquet, P., Jackson, I., Nick, M., & Kaboli, A. (2024). AI-based forecasting for optimised solar energy management and smart-grid efficiency. *International Journal of Production Research*, *62*(13), 4623–4644. https://doi.org/10.1080/00207543.2023.2293732
4. Bassey, K. E. (2023). Hybrid renewable energy systems modeling. *Engineering Science and Technology Journal*, *4*(6), 571–588. https://doi.org/10.51594/estj.v4i6.678
5. de Campos, B. N., Maionchi, D. d. O., da Silva, J. G., Biudes, M. S., Oliveira, N. N. d., & Palácios, R. d. S. (2025). Photovoltaic energy modeling using machine learning applied to meteorological variables. *Sustainability*, *17*(16), Article 7506. https://doi.org/10.3390/su17167506
6. Al-Dahidi, S., Ayadi, O., Alrbai, M., & Adeeb, J. (2019). Ensemble approach of optimized artificial neural networks for solar photovoltaic power prediction. *IEEE Access*, *7*, 81741–81758. https://doi.org/10.1109/ACCESS.2019.2929991
7. Abbasimehr, H., & Paki, R. (2022). Improving time-series forecasting using LSTM and attention models. *Journal of Ambient Intelligence and Humanized Computing*, *13*(2), 673–691. https://doi.org/10.1007/s12652-020-02761-x
8. Ciechulski, T., & Osowski, S. (2021). High-precision LSTM model for short-time load forecasting in power systems. *Energies*, *14*(11), Article 2983. https://doi.org/10.3390/en14112983
9. Zhou, H., Zhang, S., Peng, J., Zhang, S., Li, J., Xiong, H., & Zhang, W. (2021). Informer: Beyond efficient transformer for long-sequence time-series forecasting. *Proceedings of the AAAI Conference on Artificial Intelligence*, *35*(12), 11106–11115. https://doi.org/10.1609/aaai.v35i12.17325
10. Kämäräinen, J.-K. (2025). *Minimal time series transformer* (arXiv:2503.09791). arXiv. https://doi.org/10.48550/arXiv.2503.09791
11. Pimenow, S., Pimenowa, O., & Prus, P. (2024). Challenges of artificial intelligence development in the context of energy consumption and impact on climate change. *Energies*, *17*(23), Article 5965. https://doi.org/10.3390/en17235965
12. Danish, M. S. S., & Senjyu, T. (2023). AI-enabled energy policy for a sustainable future. *Sustainability*, *15*(9), Article 7643. https://doi.org/10.3390/su15097643
13. Würtz, S., Bogenberger, K., Göhner, U., & Rupp, A. (2024). Towards efficient battery electric bus operations: A novel energy forecasting framework. *World Electric Vehicle Journal*, *15*(1), Article 27. https://doi.org/10.3390/wevj15010027
14. United Nations Department of Economic and Social Affairs. (2015). *Goal 7: Affordable and clean energy*. United Nations. https://sdgs.un.org/goals/goal7
15. United Nations Department of Economic and Social Affairs. (2015). *Goal 9: Industry, innovation and infrastructure*. United Nations. https://sdgs.un.org/goals/goal9
16. United Nations Department of Economic and Social Affairs. (2015). *Goal 11: Sustainable cities and communities*. United Nations. https://sdgs.un.org/goals/goal11
17. United Nations Department of Economic and Social Affairs. (2015). *Goal 13: Climate action*. United Nations. https://sdgs.un.org/goals/goal13

About authors

Dijana Oreški is an Associate Professor of Artificial Intelligence (AI) and Machine Learning (ML) at the University of Zagreb, Faculty of Organization and Informatics. She is the head of the Laboratory for Data Mining and Intelligent Systems (LOUISE). Her research interest lies at the intersection of artificial intelligence and social sciences, focusing on the application of AI and ML to address societal challenges and support sustainable development. She has (co)authored more than 100 scientific papers and has participated in several dozen scientific and professional projects.

Vjeran Strahonja is a Professor Emeritus of Computer Science at the University of Zagreb, Faculty of Organization and Informatics. Throughout his distinguished academic career, he has held a number of leadership and managerial positions, including serving as Vice Dean and Dean of the Faculty of Organization and Informatics. His academic and professional work focuses on information systems, software engineering, and the strategic application of information technology in organizations. His long-standing contribution to academia and institutional development has had a significant impact on the field.

Marija Pokos Lukinec is a research and teaching assistant at the University of Zagreb, Faculty of Organization and Informatics, specializing in the field of artificial intelligence. Her research interests include data analytics and the application of artificial intelligence in education. She has (co)authored several scientific papers and participates in scientific and professional projects related to artificial intelligence.