Journal of Posthumanism
2026
Volume: 6, No: 6, pp. 72-87
ISSN: 2634-3576 (Print) | ISSN 2634-3584 (Online)
posthumanism.co.uk

DOI: https://doi.org/10.63332/joph.v6i6.4224

Artificial Intelligence for Energy Sustainability: Predicting Photovoltaic Generation in P2P Networks

Kandel L. Yandar[1], Oscar Revelo Sánchez[2], Manuel Bolaños Gonzales[3]

Abstract

The incorporation of renewable sources into electrical systems is essential for advancing toward sustainable energy models. However, intermittency and meteorological variability introduce uncertainty into supply operations and planning, necessitating reliable predictive tools. In this context, artificial intelligence (AI) offers alternatives for modeling generation behavior and supporting decision-making. This paper presents an AI-based knowledge discovery (KDD) model for predicting electricity generation in a peer-to-peer (P2P) network. The methodology was structured in four phases: exploration, analysis, construction, and validation. LSTM–CNN and Transformer deep learning architectures were implemented and trained on multivariate time series that integrate electrical, meteorological, and weather variables. The data come from five institutions in the department of Nariño, Colombia. The results show robust performance in both architectures ($R^2 > 0.83$), demonstrating the strategic potential of AI in distributed energy systems.

Keywords: *Artificial intelligence; Deep learning; Photovoltaic forecasting; Peer-to-peer energy networks; Knowledge discovery (KDD)*

Introduction

Electricity is a fundamental component of economic and social development, sustaining productive, industrial, and domestic activities (Laloux & Rivier, 2013). However, historical dependence on fossil fuels as the main source of power generation has had significant environmental impacts and highlighted the need to transition to more sustainable and resilient energy models.

Against this backdrop, the integration of renewable sources has taken on a leading role in the planning of modern electrical systems. Resources such as solar and wind energy offer clear environmental advantages; however, their large-scale use poses significant challenges. This process requires not only modifying generation methods but also transforming distribution and management models, promoting decentralized schemes capable of responding to the variability inherent in resources such as solar radiation or wind (*REN21*, 2024).

In this context, distributed generation emerges as an alternative that favors the decentralization of the electrical system and promotes the active participation of multiple actors. Based on this approach, peer-to-peer (P2P) energy networks allow users to act as prosumers, that is, as simultaneous generators and consumers of energy, facilitating direct energy exchanges and reducing dependence on centralized infrastructure (Akila et al., 2024). However, the efficient

[1] Universidad de Nariño, Pasto – Colombia; Email: klyandar23B@udenar.edu.co, 0000-0001-6106-3900
[2] Universidad de Nariño, Pasto – Colombia; Email: orevelo@udenar.edu.co, 0000-0003-2882-5779
[3] Universidad de Nariño, Pasto – Colombia; Email: mbolanos@udenar.edu.co, 0000-0002-3077-415X

operation of this type of network requires models capable of anticipating local generation with sufficient accuracy and advance notice, especially in scenarios where renewable production is highly variable.

Predicting electricity generation thus becomes a key element in managing P2P networks, as it allows for optimizing energy exchange, reducing losses, and improving system stability. In recent years, artificial intelligence has shown great potential for addressing this problem, particularly through deep learning techniques that enable the modeling of complex nonlinear relationships in electrical and meteorological time series. In many cases, these techniques have outperformed traditional statistical methods, especially in scenarios characterized by high variability and large data volumes.

In this context, this paper proposes an AI-based knowledge discovery model to predict electricity generation in a P2P network. The proposal is developed within the framework of the project "Multiple Agent Transactional Model of Unconventional Energy for the Department of Nariño (MTE)"(*PROYECTO SIGP CON CÓDIGO DE REGISTRO 89530, DESARROLLO DE UN MODELO TRANSACCIONAL DE ENERGÍA NO CONVENCIONAL DE MÚLTIPLES AGENTES PARA EL DEPARTAMENTO DE NARIÑO*, 2022), whose objective is the implementation of a decentralized energy exchange pilot system in southern Colombia. The methodological approach adopted is based on the KDD process and the application of deep learning architectures, specifically LSTM–CNN and Transformer models.

The article is organized as follows: Section II describes the materials and methods used; Section III presents the results obtained; Section IV discusses the most relevant findings; and Section V presents the conclusions of the study.

Materials and Methods

The methodology adopted in this study was designed with the purpose of building and evaluating an artificial intelligence-based knowledge discovery model for predicting electricity generation in a P2P network. To this end, a structured approach was followed in four main phases: exploration, analysis, construction, and validation. This methodological organization allowed us to progressively address data processing, the design of predictive models, and the evaluation of their performance in a real distributed generation environment.

Phase 1. Exploration

The objective of the exploration phase was to identify the most relevant approaches, techniques, and variables used in the literature for predicting electricity generation. To this end, a Systematic Literature Review (SLR) was conducted, previously developed by the authors (Yandar et al., 2024), which served as the conceptual and methodological basis for the present study.

The activities carried out were:

1. **Planning:** formulation of research questions, definition of selection and exclusion criteria, and choice of scientific databases.
2. **Execution:** search, removal of duplicates, initial filtering, and methodological evaluation of relevant studies.
3. **Synthesis:** extraction and comparison of the most significant findings on predictive methods applied to renewable energies.

Phase 2. Analysis

Based on the results of the RSL, the most representative techniques for predicting renewable generation were identified and classified according to the stages of the KDD process, according to Fayyad et al.(Fayyad et al., 1996): selection, preprocessing, transformation, data mining, and evaluation. This analysis made it possible to establish the methodological guidelines for the design of the proposed model.

Phase 3. Construction

In this phase, the knowledge discovery model was implemented by systematically applying the stages of the KDD process:

1. **Data Selection:** collection of electrical and meteorological records from the generator nodes of the P2P network.
2. **Preprocessing:** treatment of missing and atypical values, time synchronization, and standardization of variables.
3. **Transformation:** application of normalization techniques and creation of derived variables to improve the learning capacity of the models.
4. **Modeling:** implementation and training of LSTM–CNN and Transformer architectures, optimizing their hyperparameters through cross-validation.
5. **Evaluation:** performance estimation using statistical metrics such as MAE, RMSE, R^2, MAPE, and adjusted MAPE to measure the accuracy and robustness of each model.

Phase 4. Validation

The model was validated in two ways:

1. **Temporal validation:** using an independent data set corresponding to June, maintaining the chronological sequence to evaluate the capacity for temporal generalization.
2. **Computational performance validation:** evaluation of computational performance during the inference process in the two execution environments.

Results

This section presents the results obtained throughout the phases that make up the proposed methodology. The development is organized sequentially, following the structure defined in the materials and methods section, to present the findings corresponding to the stages of exploration, analysis, construction, and validation of the model in an orderly manner.

Exploration Phase

In this phase, an RSL was carried out with the aim of identifying the main artificial intelligence techniques used in predicting electricity generation, as well as the most relevant factors considered in the development of such models. The RSL was developed following the guidelines proposed by Kitchenham and Charters(Klumpner et al., 2006) and the methodology adapted from Zapata and Barón (Barón & Zapata, 2016).

The search was designed to answer two central questions:

RQ1: What artificial intelligence techniques are used in predicting electricity generation?

RQ2: What are the key factors used in predicting electricity generation?

To answer these questions, clear inclusion and exclusion criteria were established. The inclusion criteria were publications between 2013 and 2024, studies in English or Spanish, works that applied AI techniques to the prediction of electrical variables, and that provided a sufficient methodological description. Publications in other languages, duplicate records, studies not related to the energy domain, and secondary reviews were excluded.

The literature search was conducted in the SCOPUS, ACM Digital Library, IEEE Xplore, and Google Scholar databases, using combinations of terms related to AI and electricity generation prediction.

A total of 216 records were identified, of which 47 primary studies met the criteria defined after screening and quality assessment. Figure 1 shows the distribution of the selected articles by source.

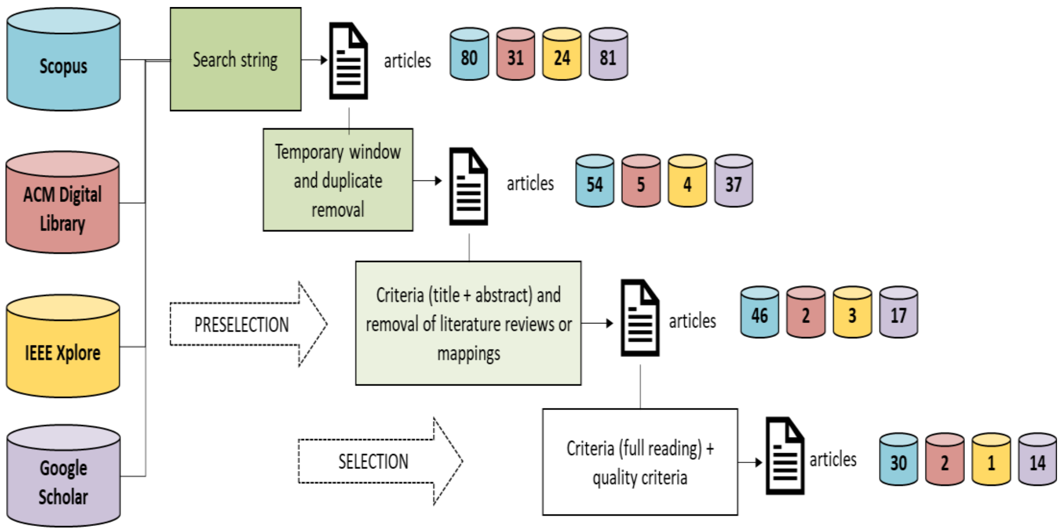

Figure 1. Study source results
Source: adapted from (Revelo Sánchez et al., 2018)

Analysis of the selected studies identified a clear predominance in the use of artificial neural networks, particularly recurrent models such as LSTM, for predicting electricity generation. Likewise, the use of techniques based on Support Vector Machines and, to a lesser extent, recent hybrid models such as LSTM–CNN, SARIMA–LSSVR, and EEMD–Informer was also evident. The literature reviewed agrees that the quality and temporal resolution of the data, together with rigorous preprocessing and adequate segmentation of the training, validation, and test sets, are determining factors for the performance of predictive models.

Taken together, these results confirm the central role of AI in energy forecasting and show a trend toward deep and hybrid architectures, which are particularly well-suited to renewable generation scenarios.

Analysis Phase

The KDD process proposed by Fayyad et al. (Fayyad et al., 1996), was used as a reference framework to organize and synthesize the findings obtained in the Systematic Literature Review. Based on this approach, the studies analyzed were classified according to the five stages of the KDD process: data selection, preprocessing, transformation, data mining, and evaluation. The main results identified in each of these stages are described below.

Data Selection

The reviewed studies show a wide variety of data acquisition strategies, which allow for the prediction of electricity generation to be addressed, considering different levels of temporal granularity and different technological contexts. Among the most representative approaches identified in the literature are the following:

One of the most common practices is the extraction of records from institutional systems, particularly in studies conducted by universities, research centers, and meteorological agencies. These sources typically provide historical databases that include variables such as solar irradiance, ambient temperature, wind speed, and power generated, which are used as the main

input for training and validating renewable generation prediction models (Alsaidan et al., 2022); (Abdul Baseer et al., 2023).
On the other hand, in cases where the availability of actual records is limited, some authors resort to computer simulation using tools such as MATLAB/Simulink (Bose, 2017); (Yuce et al., 2016). Additionally, several studies report the use of long-term historical records, covering periods of several months or years, obtained from photovoltaic plants, experimental laboratories, and wind farms. This type of information is particularly valuable for analyzing phenomena such as seasonality, equipment degradation, and variations in energy resources over different time horizons (Chaouachi et al., 2013), (Olatomiwa Lanre, 2016); (El Bakali et al., 2023), (Zahraoui et al., 2024).
Together, these methods reflect a trend toward integrating meteorological and operational sources to characterize the behavior of electricity generation.

Preprocessing

The literature shows that the multivariable nature of the generation datasets that include solar irradiance, wind speed, temperature, power generated, among others, requires the application of systematic cleaning and imputation processes, adjusted to the type of source and the temporal resolution of the data.
One of the most relevant aspects identified is the detection and correction of anomalous values, a stage in which inconsistent or physically impossible records are eliminated, such as zero power during solar hours or excessive values derived from calibration errors in sensors and inverters. In (El Abbadi & El Youbi, 2023) et (Abdul Baseer et al., 2023).
Various imputation techniques were also used, including temporal interpolation, used in (El Abbadi & El Youbi, 2023) and (Zahraoui et al., 2024), which allows missing values to be estimated from the nearest points, maintaining the natural trend of meteorological variables such as irradiance or wind speed. Imputation by averaging, applied in (Hajjaji et al., 2021), proved adequate in scenarios where the system maintains operational stability and the missing data appear between regular measurements. On the other hand, zero imputation, reported in (Martinez, 2023), was applied when the absences actually corresponded to periods without generation, as occurs in photovoltaic systems during nighttime hours.
These preprocessing strategies enable cleaner, more continuous, and representative data to be obtained, ensuring that the inputs to AI models reflect actual power generation conditions.

Transformation

The literature reviewed indicates that data transformation techniques aim to reduce the complexity of information, highlight relevant patterns, and ensure that variables are suitable for processing using machine learning algorithms.
Among the most commonly used transformations is normalization, applied using methods such as Min–Max and Z-score, to bring variables to comparable scales and promote the stability of model training, especially neural networks (Alsaidan et al., 2022), (El Bakali et al., 2023), (Zahraoui et al., 2024), (El Abbadi & El Youbi, 2023), and (Abdul Baseer et al., 2023). These techniques adjust variables to comparable scales, stabilize neural network training, and reduce the influence of extreme values.
Temporal aggregation is also widely used to capture trends and smooth out fluctuations in irradiance or wind, using moving averages (Olatomiwa Lanre, 2016); (El Bakali et al., 2023) or hourly groupings (Chaouachi et al., 2013), (El Abbadi & El Youbi, 2023), which allows the daily cycles of the renewable resource to be highlighted.
In relation to feature selection, irradiance, temperature, wind speed, and temporal variables are

consistently identified as the most influential factors in prediction, while other meteorological variables have a limited contribution (Abdul Baseer et al., 2023) and (Zahraoui et al., 2024).
On the other hand, the construction of derived characteristics improves the physical representation of the generation phenomenon. Among these, the sky index stands out, used in (El Bakali et al., 2023) and (Zahraoui et al., 2024) to correct the effects of cloudiness, and the cyclic encodings of temporal variables using sine-cosine functions, applied in (Olatomiwa Lanre, 2016) and (El Abbadi & El Youbi, 2023) to capture the natural periodicity of solar and wind resources.
Data Mining
The SRL shows evidence of the use of a wide variety of artificial intelligence techniques for predicting electricity generation, ranging from traditional neural models to hybrid architectures and advanced deep learning approaches. Among the most widely used techniques are Artificial Neural Networks (ANN), noted for their ability to model nonlinear relationships between meteorological and electrical variables (Alhebshi et al., 2019); (Bose, 2017); and (El Abbadi & El Youbi, 2023). Similarly, LSTM models stand out for their effectiveness in capturing long-term temporal dependencies in energy time series (Zahraoui et al., 2024) and (Abdul Baseer et al., 2023).
Additionally, some studies incorporate ANFIS systems, which combine neural networks and fuzzy logic to improve model interpretability and uncertainty management in energy scenarios (Alsaidan et al., 2022) and (Mutombo, 2017). Hybrid models, such as CNN–LSTM, ANN–LSTM, and ARIMA–ANN, integrate statistical approaches and deep learning techniques to increase the robustness and accuracy of predictions (Yuce et al., 2016) and (Zahraoui et al., 2024).
Evaluation and Interpretation
Analysis of the studies included in the Systematic Literature Review shows that the evaluation of the performance of electricity generation prediction models is based on a diverse set of metrics aimed at measuring both the accuracy and robustness of the algorithms. Among the most commonly used indicators is the coefficient of determination (R^2), widely used for its ability to express the proportion of variability explained by the model (Olatomiwa Lanre, 2016); (El Abbadi & El Youbi, 2023); (Abdul Baseer et al., 2023) and (Alhebshi et al., 2019).
Complementarily, absolute error metrics such as MAE and RMSE, widely used in (El Bakali et al., 2023), (El Abbadi & El Youbi, 2023), (Abdul Baseer et al., 2023), (Harrou et al., 2023), (Hussein, 2018), (Mutombo, 2017) and (Zahraoui et al., 2024), allow the magnitude and stability of predictions to be quantified, with RMSE being particularly useful due to its sensitivity to large errors. Likewise, relative metrics such as MAPE, referenced in (Alsaidan et al., 2022), (Chaouachi et al., 2013), (Olatomiwa Lanre, 2016), (Zahraoui et al., 2024) and (Abdul Baseer et al., 2023), facilitate the percentage comparison of error between different scales, although their interpretation can be distorted when real values are close to zero.
Construction Phase
In this phase, the predictive model for energy generation was constructed, systematically following the stages of the KDD process, from data collection to evaluation of the final model. Figure 2 illustrates the stages applied to the knowledge discovery process.

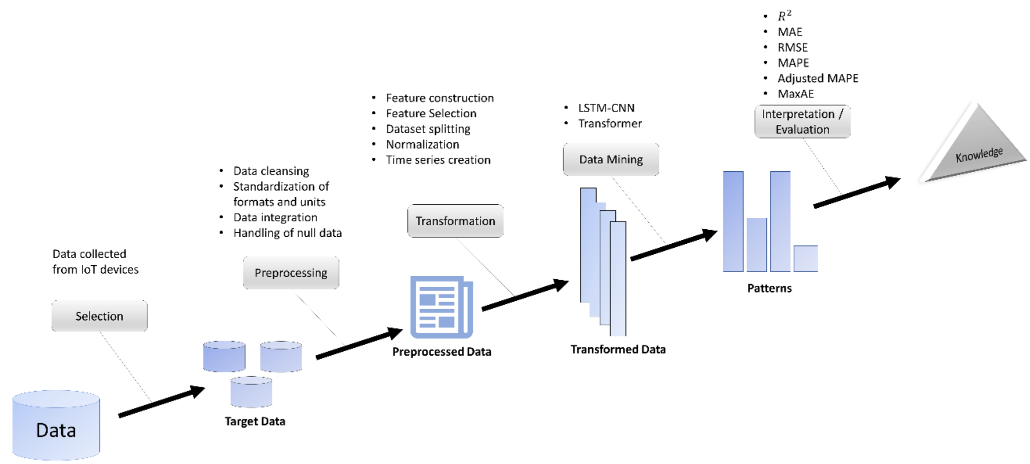

Figure 2. KDD process
Source: Adapted from (Fayyad et al., 1996)

Data Selection

The dataset used was constructed based on the infrastructure of the project "Transactional Model of Unconventional Energy with Multiple Agents for the Department of Nariño" (*PROYECTO SIGP CON CÓDIGO DE REGISTRO 89530, DESARROLLO DE UN MODELO TRANSACCIONAL DE ENERGÍA NO CONVENCIONAL DE MÚLTIPLES AGENTES PARA EL DEPARTAMENTO DE NARIÑO*, 2022), which implemented a monitoring system in five nodes of the P2P network: UDENAR, UCC, UNIMAR, UCESMAG, and HUDN.

Each node has smart meters, photovoltaic inverters, and weather stations integrated into a central platform, which allowed for the collection of multivariate time series of electrical variables (active power, voltage, current) and meteorological variables (irradiance, temperature, humidity, wind speed), as well as temporal information associated with daily and seasonal generation cycles.

The volume of data varied between institutions due to installation dates, transmission continuity, and equipment availability; UDENAR recorded the highest number of measurements, while HUDN had the lowest coverage due to its late incorporation. Even so, all nodes provided relevant and complementary information.

Differences in the density and continuity of the records were subsequently addressed in preprocessing and transformation, ensuring a consistent and adequate dataset for modeling electricity generation in the P2P network.

Preprocessing

Preprocessing was applied independently at each node to ensure quality, temporal consistency, and homogeneity before modeling. Common actions included: conversion of inverter power from W to kW, temporal synchronization between the meter, inverter, and weather station using merge_asof with a tolerance of 120 s, and treatment of missing values to preserve the continuity of the series.

Similar patterns were identified during integration: meters and stations recorded mostly continuous data, with occasional gaps in climate variables, which were imputed using temporal interpolation, maintaining the natural trend of the solar resource. Inverters showed systematic gaps during nighttime hours, which were assigned a value of zero as they corresponded to the physical absence of generation.

In nodes with shorter windows or prolonged interruptions, such as in HUDN, temporal consistency was maintained without losing relevant information. After integrating the devices, the data was consolidated through temporal concatenation, and a node identifier was added, resulting in a final dataset with more than 360,000 records and 81 variables.

The decisions made—flexible synchronization, interpolation only in meteorological variables, and the use of zero in periods without generation—ensured physical consistency, continuity, and the absence of bias due to over-imputation, providing a solid basis for the transformation and modeling stages.

Transformation

The Transformation phase brought together the operations necessary to structure the dataset and prepare it for deep learning-based electricity generation modeling.

The target variable generated_energy_kWh was constructed from the active power recorded by the inverters at two-minute intervals. To enrich the temporal context, cyclic time and day encodings (sine–cosine) and a weekend indicator were added.

The dataset was divided into training (70%), validation (15%), and testing (15%) sets, respecting the temporal order. Subsequently, normalization was applied with RobustScaler, adjusted only with the training data, to reduce the impact of outliers typical of solar generation.

Feature selection was performed using a combined approach of Pearson correlation and recursive feature elimination (RFE) with Random Forest, resulting in a final set of 34 relevant electrical, meteorological, and temporal variables.

Finally, time series were constructed using sliding windows of 120 steps (four hours), which served as direct input for the LSTM–CNN and Transformer models, allowing short- and long-term temporal dependencies to be captured.

Data Mining

The data mining phase consisted of applying artificial intelligence models to identify patterns in photovoltaic generation and produce predictions from the transformed dataset. Two deep learning architectures adapted to the multisite characteristics of the P2P network were evaluated: a hybrid LSTM–CNN model and a Transformer model.

The LSTM–CNN model combines the ability of LSTMs to capture temporal dependencies with the ability of CNNs to detect local variations. It receives sequences of 120 steps with 34 electrical and meteorological variables, as well as an embedding to identify the node. The architecture includes a recurrent branch of three LSTM layers (64–32–16 units) and another convolutional branch with two Conv1D layers (64–32 filters). The outputs are integrated into a dense layer of 64 neurons with 30% dropout and terminated in a linear layer that estimates the energy generated.

The Transformer model was implemented as a non-recurrent alternative to capture parallel temporal relationships. It uses the same input as the hybrid model and processes the sequences using an encoder based on layer normalization, multi-head attention (4 heads, key dimension of 64), and a feed-forward network of 128 neurons with residual connections and 0.3 dropout. The flattened output is projected using a linear layer to obtain the prediction in kWh.

Evaluation and Interpretation

This section presents the overall results obtained by the LSTM–CNN and Transformer models for predicting photovoltaic power generation in the P2P network, as well as their computational performance in different execution environments.

- Model performance evaluation

The quantitative analysis was performed using the R², MAE, RMSE, MAPE, adjusted MAPE, and MaxAE metrics. Table 1 presents the results of the metrics.

Metric\ Model	LSTM-CNN	*Transformer*
R^2	0.8481	0.8368
MAE	0.0157	0.01859
RMSE	0.0370	0.03837
MAPE	429.10%	493.81%
MAPE ajustado	21.26%	23.01%
MaxAE	0.3793	0.39575

Table 1. Evaluation Metric Results
Source: Own Elaboration

Globally, both models exhibited solid performance, with R² values above 0.83, reflecting a high correlation between actual and estimated values.
The LSTM–CNN model achieved the best overall performance, recording MAE = 0.0157 kWh and RMSE = 0.0370 kWh. However, the Transformer presented very similar results, MAE = 0.0186 kWh and RMSE = 0.0384 kWh, confirming its ability to model complex photovoltaic generation patterns despite not using recurrent mechanisms. Furthermore, its stability in the face of abrupt variations demonstrates the effectiveness of the attention mechanism in capturing temporal relationships and dependencies between meteorological and electrical variables.
The maximum absolute error (MaxAE), close to 0.39 kWh in both models, indicates that neither presented severe deviations in the estimation of generation peaks. In terms of relative errors, the classic MAPE showed high values due to periods without generation, so an adjusted MAPE was used that excludes values below 0.1 kWh, obtaining 21.26% for the LSTM–CNN model and 23.01% for the Transformer, showing that both models maintain adequate performance in operational ranges of energy generation.
The overall results confirm that both the LSTM–CNN hybrid model and the Transformer are capable of learning and reproducing patterns typical of electricity generation in distributed environments.

- **Computational Performance Evaluation**

To analyze the computational performance of the proposed models, their performance was evaluated in two contrasting computing environments: a high-performance server with a GPU and a personal computer without a GPU. This comparison allowed us to analyze the impact of hardware on training times and the computational behavior of the models.
The results show marked differences between the two environments. On the server, the LSTM–CNN hybrid model recorded a training time per epoch of 1.6 minutes and a total time of 38 minutes, while the Transformer was faster, with 30 seconds per epoch and a total time of 12.5 minutes. This advantage is due to the Transformer's ability to process sequences in parallel and take more efficient advantage of GPU-based hardware.
On the other hand, on the personal computer, the times increased dramatically. The LSTM–CNN required 26.65 minutes per epoch and a total of 639.6 minutes (≈10.6 hours), while the Transformer required 18.9 minutes per epoch and a cumulative time of 474.16 minutes (≈7.9 hours). These results demonstrate the strong dependence of both architectures on GPU acceleration, although the Transformer maintains a relative advantage even on the CPU.
In terms of resource usage, the server showed highly efficient execution: LSTM–CNN used 40% of the GPU and 0.1% of the CPU, while Transformer used 58.5% of the GPU and 0.3% of the

CPU. RAM consumption was also low (between 5.5% and 6.8%), reflecting the stability of processing in optimized environments.
In contrast, the personal computer showed critical resource utilization. The LSTM–CNN reached 90% RAM usage, and the Transformer reached 81%, compromising the execution of additional tasks. In addition, the CPU operated at sustained loads of 65% and 70%, respectively, due to the absence of graphics acceleration.
In summary, the results confirm that computing infrastructure has a decisive impact on the efficiency of prediction models. The Transformer is positioned as the fastest and most scalable architecture in both environments, while the LSTM–CNN, although more accurate, requires more time and resources.

Validation Phase

In order to evaluate the generalization capacity of the proposed models when faced with information not used during training, a temporary validation was carried out using the records corresponding to the month of June as an independent test set. These data were not included in any of the previous modeling stages, which allowed the performance of the models to be analyzed in a completely new scenario consistent with the sequential nature of the time series.
The validation is presented on two levels: (i) overall results, integrating data from all nodes in the P2P network, and (ii) computational performance during inference in the two execution environments evaluated. Additionally, a statistical contrast test is incorporated to determine whether the differences observed between the models are significant from an inferential perspective.

Validation Results – Overall

Table 2 summarizes the results obtained in the overall validation:

Métrica\ Modelo	**LSTM-CNN**		***Transformer***	
	Consumo	**Generación**	**Consumo**	**Generación**
R^2	0.9146	0.8330	0.8991	0.8336
MAE	0.0072	0.0198	0.0088	0.0220
RMSE	0.0136	0.0430	0.0148	0.0429
MAPE	52.80%	—*	73.01%	—*
MAPE ajustado	10.78%	23.97%	12.85%	23.93%
MaxAE	0.17707	0.39596	0.17578	0.40493

Table 2. Metrics results - validation phase
Source: own elaboration
* *The classic MAPE is unstable with absolute values close to zero; therefore, only adjusted MAPE should be used for interpretation.*
The results show that both models obtained R^2 values close to 0.83, indicating similar overall performance during the June prediction. The LSTM–CNN model achieved an MAE of 0.0198 kWh and an RMSE of 0.0430 kWh, while the Transformer recorded an MAE of 0.0220 kWh and an RMSE of 0.0429 kWh. These figures reflect low and consistent absolute errors in both architectures.
The classic MAPE presented extremely high values due to the presence of records with zero generation, so the adjusted MAPE was used as the main metric. After excluding values below 0.1 kWh, similar results were obtained between the models: 23.97% for LSTM–CNN and 23.93% for Transformer.
The maximum absolute error (MaxAE) remained below 0.41 kWh in both cases, indicating that

the highest error remained within ranges during the validation period.
The graphical representations in Figures 3 and 4 show the alignment between actual generation and model predictions for the month of June 2025. The overlap of the curves shows that both architectures managed to reproduce the general shape of the time series, including the daily variations and peaks characteristic of photovoltaic behavior during this period.

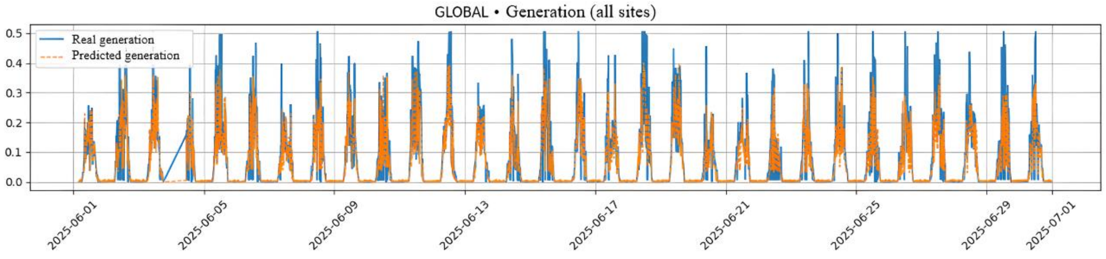

Figure 1. Power generation forecasts – LSTMCNN
Source: own elaboration

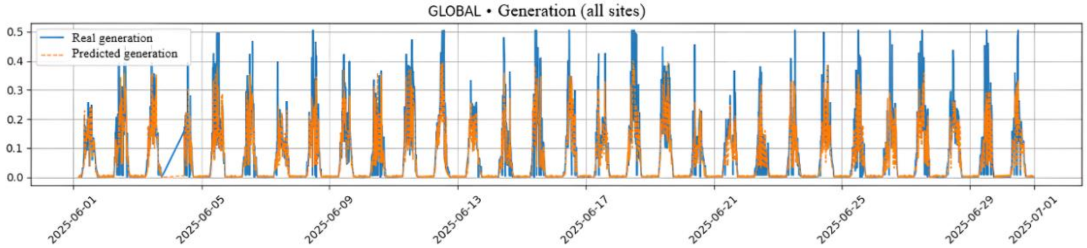

Figura 2. Energy generation predictions – Transformer
Source: own elaboration

Computational Performance Results

During the validation phase, the computational performance of both models was evaluated in two environments: a high-performance server and a personal computer without a GPU.
On the server, inference times were reduced: 36 seconds for the LSTM–CNN model and 44 seconds for the Transformer. RAM usage remained low (between 2.3% and 3.5%), while CPU load was minimal (<1%), as the GPU took on most of the processing with utilization levels between 6.5% and 8.6%. These results reflect fast, stable, and highly optimized execution on specialized infrastructure.
On a personal computer without GPU acceleration, execution times increased considerably. The LSTM–CNN completed inference in 2.9 minutes and the Transformer in 4.8 minutes. In both cases, the CPU operated at 100% capacity throughout the process, and RAM usage was higher (50% for LSTM–CNN and 20% for Transformer), indicating greater resource demand in resource-constrained environments.
It is evident that both models demonstrated stability and portability during validation, although computational efficiency depended largely on the execution environment.

Hypothesis Testing

To complement the quantitative analysis and determine whether the differences observed between the models could be attributed to chance or systematic behavior, a statistical test was performed to directly compare the performance of the LSTM–CNN and the Transformer. This analysis allows us to validate, from an inferential perspective, whether the models differ

significantly in their predictive ability.
For this purpose, we used the nonparametric Wilcoxon signed-rank test, which is suitable for small samples, in this case, the five nodes of the P2P network, and for metrics that do not require the assumption of normality. The hypothesis testing was applied to the MAE, RMSE, and adjusted MAPE metrics, considering the results obtained in both the evaluation and validation phases.
In all cases analyzed, the p-values obtained were higher than the established significance level ($p \geq 0.05$), indicating that there is insufficient statistical evidence to claim that one model significantly outperforms the other. Although LSTM–CNN presented slightly lower absolute errors in some scenarios, these differences were not statistically significant.
Overall, the results suggest that the variations observed in the performance of the models are mainly associated with the characteristics of the data and not with the intrinsic superiority of any of the architectures. Consequently, the choice between LSTM–CNN and Transformer can be based on practical criteria, such as the availability of computational resources, inference speed, or scalability requirements, given that both models are solid alternatives for predicting distributed generation in smart energy networks.

Discussion

Analysis of Model Evaluation Results

The overall results show that both the LSTM–CNN and Transformer models perform well in predicting photovoltaic generation in the P2P network, with R^2 values above 0.83. This behavior indicates that both architectures are capable of capturing a significant proportion of the temporal variability associated with solar energy production, which is characterized by a high dependence on weather conditions.
The LSTM–CNN model showed a slight advantage in terms of absolute accuracy, reflected in marginally lower MAE and RMSE values. This result can be attributed to the hybrid nature of the model, which combines the ability of LSTMs to capture long-term temporal dependencies with the ability of CNNs to detect local patterns in the signal. However, the difference observed compared to the Transformer was small, showing that the latter also manages to efficiently model the temporal relationships of photovoltaic generation through its attention and parallel processing mechanism.
With regard to relative error metrics, the classic MAPE showed significant limitations due to the presence of periods with near-zero generation, a frequent condition in photovoltaic series. When applying the adjusted MAPE, relative errors were reduced to values close to 22%, providing a more representative estimate of the actual performance of the models under operating conditions.
These results indicate that both architectures exhibit stable behavior for predicting distributed generation in P2P networks, with a slight predominance of the LSTM–CNN hybrid model in prediction accuracy and competitive performance of the Transformer. This descriptive analysis forms the basis for the inferential evaluation subsequently performed using statistical tests.

Analysis of Model Validation Results

The temporal validation performed using independent data corresponding to the month of June allowed us to evaluate the generalization capacity of the models in a scenario not seen during training. The results show that both the LSTM–CNN and the Transformer maintained stable performance, reaching R^2 values close to 0.83.
In the results obtained with the absolute error, the models maintained values similar to those obtained during the evaluation phase: the MAE was around 0.02 kWh and the RMSE close to 0.043 kWh, showing that the average magnitude of the error remained under control. This implies

that both methods have an adequate generalization capacity in the face of daily variations and abrupt fluctuations in radiation.
As in the evaluation stage, the classic MAPE again presented extremely high values, a result characteristic of photovoltaic series with multiple records close to zero during nighttime hours or under conditions of minimum irradiance. To correct this distortion, the adjusted MAPE was applied, excluding values less than 0.1 kWh, which led us to obtain relative errors of around 24% for both models, which is a much more realistic representation of performance.
In terms of maximum absolute error (MaxAE), both models showed deviations of less than 0.41 kWh in their worst predictions. This result confirms that, even in the face of highly variable solar events, neither model incurred critical errors or instability that compromised the overall reliability of the estimates.
These findings show that, although the hybrid model maintained a slight advantage in average error, the performance of both architectures was comparable, demonstrating their ability to operate in real prediction scenarios within distributed energy networks.

Computational Performance Analysis

The analysis of computational performance in the evaluation and validation phases reveals marked differences between the models and the execution environments. On the server with high computational resources, both models achieved reduced times and efficient execution thanks to the use of GPU units. On the one hand, the Transformer stood out for its speed in both the training process (30 s per epoch) and the inference process (≈44 s), while the LSTM–CNN, although slightly slower (96 s per epoch and 36 s in inference), maintained stable and low resource usage. This behavior confirms that both architectures are compatible with near real-time generation prediction deployments when accelerated infrastructure is available.
In contrast, on the personal computer without a GPU, the times increased significantly since, during training, the LSTM–CNN model required more than 10 hours for training and the Transformer required about 8 hours; during inference, the times ranged from 3 to 5 minutes. In both models, the CPU reached 100% utilization, reflecting hardware saturation and considerably higher memory consumption than on the server. We conclude that, although execution in resource-limited environments is feasible, the computational cost increases significantly, reducing scalability and the frequency with which predictions can be generated.
The findings confirm that the computational performance of the models depends significantly on the availability of GPU acceleration and the parallelization capacity of the environment. In this regard, the Transformer offers an advantage in terms of speed and computational efficiency, while the hybrid LSTM–CNN model performs better in predictive accuracy. Therefore, model selection for real-world applications must consider a balance between accuracy requirements and available computational resources.

Conclusions

This study developed an artificial intelligence-based knowledge discovery model for predicting electricity generation in a P2P network, systematically applying the phases of the KDD process. The implementation and evaluation of the LSTM–CNN and Transformer architectures demonstrated that both models are capable of adequately capturing the temporal dynamics associated with photovoltaic generation at the different nodes of the system.
The results obtained show robust performance in predicting solar generation, with consistent and stable error values that confirm the models' ability to reproduce daily and seasonal variations in solar resources. These findings validate the suitability of the evaluated architectures as support tools for managing distributed renewable energy systems in P2P environments.

From a computational point of view, the study shows that the operational performance of the models is strongly conditioned by the available infrastructure. In environments with GPU units, both models exhibit reduced inference times and efficient resource usage, while in scenarios where only CPU computers are used, the computational cost increases considerably. In this sense, the choice between LSTM–CNN and Transformer must consider a balance between the required accuracy and the available computational resources, as well as the scalability needs of the system. The Wilcoxon statistical test applied to the error metrics showed that there are no significant differences ($p > 0.05$) between the performance of LSTM–CNN and Transformer. This indicates that, from a quantitative perspective, both models are statistically equivalent in terms of predictive accuracy, differing mainly in their computational behavior and the speed with which they process large volumes of data.

Acknowledgements: We would like to thank the project "Transactional Model of Unconventional Energy with Multiple Agents for the Department of Nariño (MTE)" BPIN 2021000100499, funded by the General Royalties System (SGR) – CTeI Fund of the Department of Nariño, as well as the Galeras.NET and GIIEE research groups of the Faculty of Engineering of the University of Nariño, for the technical, academic, and logistical support provided during the development of this research.

References

Abdul Baseer, M., Almunif, A., Alsaduni, I., & Tazeen, N. (2023). Electrical Power Generation Forecasting from Renewable Energy Systems Using Artificial Intelligence Techniques. *Energies*, *16*(18). https://doi.org/10.3390/en16186414

Akila, A., Bansal, R. C., & Abo-Khalil, A. (2024). Advancements in Peer-to-Peer Energy Management Systems – An Overview. *2024 7th International Conference on Electric Power and Energy Conversion Systems (EPECS)*, 167–172. https://doi.org/10.1109/EPECS62845.2024.10805521

Alhebshi, F., Alnabilsi, H., Bensenouci, A., & Brahimi, T. (2019). Using artificial intelligence techniques for solar irradiation forecasting: The case of Saudi Arabia. *Proceedings of the International Conference on Industrial Engineering and Operations Management*, *November*, 926–927.

Alsaidan, I., Rizwan, M., & Alaraj, M. (2022). Solar energy forecasting using intelligent techniques: A step towards sustainable power generating system. *Journal of Intelligent & Fuzzy Systems*, *42*(2), 885–896. https://doi.org/10.3233/JIFS-189757

Barón, A., & Zapata, C. (2016). Estrategia Metodológica para la Elaboración de Síntesis Conceptuales en Ingeniería de Software: una Aplicación al Caso del Constructo Teórico de Práctica. *4th International Conference in Software Engineering Research and Innovation CONISOFT*.

Bose, B. K. (2017). Artificial Intelligence Techniques in Smart Grid and Renewable Energy Systems—Some Example Applications. *Proceedings of the IEEE*, *105*(11), 2262–2273. https://doi.org/10.1109/JPROC.2017.2756596

Chaouachi, A., Kamel, R. M., Andoulsi, R., & Nagasaka, K. (2013). Multiobjective Intelligent Energy Management for a Microgrid. *IEEE Transactions on Industrial Electronics*, *60*(4), 1688–1699. https://doi.org/10.1109/TIE.2012.2188873

El Abbadi, N., & El Youbi, M. S. (2023). ENHANCING WIND ENERGY FORECASTING THROUGH THE APPLICATION OF ARTIFICIAL INTELLIGENCE TECHNIQUES: A COMPREHENSIVE STUDY. *International Journal on Technical and Physical Problems of Engineering*, *15*(3).

El Bakali, S., Ouadi, H., & Gheouany, S. (2023). Solar Radiation Forecasting Using Artificial Intelligence Techniques for Energy Management System. In *Lecture Notes in Networks and Systems: Vol. 714 LNNS* (pp. 408–421). https://doi.org/10.1007/978-3-031-35245-4_38

Fayyad, U., Piatetsky-Shapiro, G., & Smyth, P. (1996). From data mining to knowledge discovery in databases. *AI Magazine*, *17*(3), 37–53.

Hajjaji, I., Alami, H. El, El-Fenni, M. R., & Dahmouni, H. (2021). Evaluation of Artificial Intelligence Algorithms for Predicting Power Consumption in University Campus Microgrid. *2021 International Wireless Communications and Mobile Computing (IWCMC)*, 2121–2126. https://doi.org/10.1109/IWCMC51323.2021.9498891

Harrou, F., Sun, Y., Taghezouit, B., & Dairi, A. (2023). Artificial Intelligence Techniques for Solar Irradiance and PV Modeling and Forecasting. In *Energies* (Vol. 16, Issue 18). https://doi.org/10.3390/en16186731

Hussein, H. (2018). An Optimal Design Methodology of Adaptive Neuro-Fuzzy Inference System for Energy Load Forecasting - Hail city case study (Saudi Arabia). *Proceedings of the Fourth International Conference on Engineering & MIS 2018*, 1–7. https://doi.org/10.1145/3234698.3234765

Klumpner, C., Wijekoon, T., & Wheeler, P. (2006). New methods for the active compensation of unbalanced supply voltages for two-stage direct power converters. *IEEJ Transactions on Industry Applications*, *126*(5), 589–598. https://doi.org/10.1541/ieejias.126.589

Laloux, D., & Rivier, M. (2013). *Technology and Operation of Electric Power Systems* (pp. 1–46). https://doi.org/10.1007/978-1-4471-5034-3_1

Martinez, S. B. (2023). *Energy management systems for smart homes and local energy communities based on optimization and artificial intelligence techniques*. Universidad Politécnica de Cataluña.

Mutombo, N. M.-A. (2017). *Development of neuro-fuzzy strategies for prediction and management of hybrid PV-PEMFC-battery systems.* University of KwaZulu-Natal.

Olatomiwa Lanre, J. (2016). *Optimal planning and design of hybrid renewable energy system for rural healthcare facilities / Olatomiwa Lanre Joseph.*

PROYECTO SIGP CON CÓDIGO DE REGISTRO 89530, DESARROLLO DE UN MODELO TRANSACCIONAL DE ENERGÍA NO CONVENCIONAL DE MÚLTIPLES AGENTES PARA EL DEPARTAMENTO DE NARIÑO. (2022).

REN21. (2024). Renewables Global Status Report. https://www.ren21.net/gsr-2024/modules/global_overview

Revelo Sánchez, O., Collazos Ordoñez, C. A., & Jiménez Toledo, J. A. (2018). La gamificación como estrategia didáctica para la enseñanza/aprendizaje de la programación: un mapeo sistemático de literatura. *Lámpsakos*, *19*, 31–46. https://doi.org/10.21501/21454086.2347

Yandar, K. L., Revelo Sánchez, O., & Bolaños-González, M. E. (2024). Machine learning para la predicción de energía eléctrica: una revisión sistemática de literatura. *Ingeniería y Competitividad*, *26*(2). https://doi.org/10.25100/iyc.v26i2.13875

Yuce, B., Rezgui, Y., & Mourshed, M. (2016). ANN–GA smart appliance scheduling for optimised energy management in the domestic sector. *Energy and Buildings*, *111*, 311–325. https://doi.org/10.1016/j.enbuild.2015.11.017

Zahraoui, Y., Korotko, T., Mekhilef, S., & Rosin, A. (2024). ANN-LSTM Based Tool for

Photovoltaic Power Forecasting. *4th International Conference on Smart Grid and Renewable Energy, SGRE 2024 - Proceedings*. https://doi.org/10.1109/SGRE59715.2024.10428969