

## RESEARCH ARTICLE

## Just in Time Transformers

AHMED ALA EDDINE BENALI<sup>1</sup>, (Graduate Student Member, IEEE),  
MASSIMO CAFARO<sup>1</sup>, (Senior Member, IEEE), ITALO EPICOCO<sup>1</sup>, MARCO PULIMENO<sup>1</sup>,  
AND ENRICO JUNIOR SCHIOPPA<sup>2</sup>

<sup>1</sup>Department of Engineering for Innovation, University of Salento, 73100 Lecce, Italy

<sup>2</sup>Research and Development, Inmatica S.p.A, 00144 Rome, Italy

Corresponding author: Massimo Cafaro (massimo.cafaro@unisalento.it)

**ABSTRACT** Precise energy load forecasting in residential households is crucial for mitigating carbon emissions and enhancing energy efficiency; indeed, accurate forecasting enables utility companies and policymakers, who advocate sustainable energy practices, to optimize resource utilization. Moreover, smart meters provide valuable information by allowing for granular insights into consumption patterns. Building upon available smart meter data, our study aims to cluster consumers into distinct groups according to their energy usage behaviours, effectively capturing a diverse spectrum of consumption patterns. Next, we design JITrans (Just In Time transformer), a novel transformer deep learning model that significantly improves energy consumption forecasting accuracy, with respect to traditional forecasting methods. Extensive experimental results validate our claims using proprietary smart meter data. Our findings highlight the potential of advanced predictive technologies to revolutionize energy management and advance sustainable power systems: the development of efficient and eco-friendly energy solutions critically depends on such technologies.

**INDEX TERMS** Clustering, energy consumption patterns, energy load forecasting, smart meters, transformers.

## I. INTRODUCTION

Electrical energy is one of the main drivers of modern society, powering everything from our homes and industries to the various nuances of our daily routines. With rapid economic growth, rising populations, and a shift toward renewable energy sources, accurately predicting energy consumption has become essential for minimizing carbon emissions and addressing climate change. Representing around 27% of global electricity demand, the residential sector plays a critical role in this pursuit [1]. In response to these dynamics, the International Energy Agency (IEA) released its “Electricity 2024” report,<sup>1</sup> providing insights into recent trends and future forecasts for global electricity demand. According to the report, global electricity demand grew by 2.2% in 2023, slightly lower than the 2.4% observed in 2022, primarily due to reduced consumption in advanced economies. However,

demand is expected to accelerate to an average of 3.4% per year from 2024 to 2026. This expected increase highlights the need for efficient energy management strategies that rely significantly on advancements in forecasting techniques.

Advancements in time-series forecasting has been and continues to be crucial in addressing the complexities of the energy sector. The journey began with the adoption of statistical models (e.g., ARIMA, the Autoregressive Integrated Moving Average) in the 1970s, valued for their simplicity and effectiveness in handling linear time-series data [2]. However, as the complexity and volume of data increased, more sophisticated models became necessary to improve forecasting accuracy. Support Vector Machines (SVM) and Random Forests (RF) were among the first machine learning algorithms used in the late 20th century, enabling the identification of non-linear patterns beyond the reach of traditional models [3], [4]. The development of Recurrent Neural Networks (RNNs) marked a significant advancement, enabling enhanced processing of sequential data by maintaining memory across data sequences, which

The associate editor coordinating the review of this manuscript and approving it for publication was Zhengmao Li<sup>1</sup>.

<sup>1</sup><https://www.iea.org/reports/electricity-2024/executive-summary>

was critical for forecasting tasks [5]. The introduction of Long Short-Term Memory (LSTM) networks addressed the limitations of earlier RNNs by capturing long-term dependencies within datasets [6]. Convolutional Neural Networks (CNNs), primarily used for image processing, have also been successfully adapted for time-series forecasting, recognizing local patterns within the data [7].

More recently, the adaptation of Transformer models, originally designed for Natural Language Processing (NLP) tasks, has revolutionized time-series forecasting. These models employ self-attention mechanisms that adeptly capture both short- and long-term dependencies in data, significantly enhancing forecasting precision compared to earlier models [8]. Given these advancements, effective energy management can lead to substantial cost savings for consumers, enhance the reliability of power systems, and contribute to a more sustainable future [9], [10], [11], [12].

The landscape of residential energy consumption is rapidly evolving with the integration of technologies such as electric vehicles, solar panels, and wind turbines. These additions introduce considerable variability and complexity into consumption patterns, challenging traditional forecasting methods [12], [13]. For instance, electric vehicles require substantial power for charging, often during off-peak hours [11], while households with solar panels may unpredictably feed energy back to the grid, depending on weather conditions. Similarly, wind turbines add another layer of variability, influenced by factors like wind patterns and geographic location [9]. These technologies not only increase fluctuations in energy usage but also create bidirectional energy flows as households become both consumers and potential energy suppliers [13]. Such shifts demand forecasting models capable of handling real-time, high-frequency data and adapting to rapid changes. Furthermore, smart meters provide continuous streams of granular data, capturing real-time consumption fluctuations. While valuable, this data intensifies the need for adaptive models capable of processing large-scale, complex datasets [10], [14]. Behavioral and demographic factors, including household size, energy awareness, and appliance usage, further influence consumption patterns, creating diverse profiles that require models to go beyond traditional methods to capture the unique nuances of individual households effectively.

To address the aforementioned challenges in residential energy consumption prediction, we propose the JITrans (Just In Time Transformer) model, a novel approach leveraging smart meter data to cluster residential energy consumers based on their usage patterns. By forming distinct clusters, we capture diverse consumption trends reflective of various household behaviors and lifestyles. Analyzing a dataset spanning 28 months from an Italian power company, we deploy JITrans, a transformer model designed to predict energy consumption within these clusters. Moreover, by leveraging a modular design where each transformer in the ensemble is specialized for a different forecast horizon, JITrans

achieves robust generalization for each cluster, minimizing cumulative errors over time. This modular and adaptive approach ensures that JITrans not only enhances predictive accuracy but also addresses the diverse and evolving energy usage dynamics found in residential settings. Furthermore, extensive experimental results reveal that our JITrans model outperforms established deep learning models such as LSTM, CNN-LSTM, GRU, and even the basic transformer model. Our contributions are three-fold:

- We provide a comprehensive review of related work, categorized into key areas that reflect the evolution and current trends in electricity consumption forecasting;
- We introduce a highly effective approach to improve the predictive accuracy of energy load forecasting, based on clustering the users and a novel transformer model we designed, JITrans;
- We evaluate our model generalization ability across clusters featuring varied behavioural patterns in residential energy consumption and provide an extensive comparative analysis of our proposed approach against traditional deep learning models.

The remainder of this paper is structured as follows. Section II recalls related work in clustering and energy load prediction. Section III details the problem statement, highlighting specific challenges our research aims to address. Section IV presents the overall architecture and the JITrans model. We discuss the experimental results in Section V. Finally, we draw our conclusions in Section VI, summarizing our findings and suggesting possible future research directions.

## II. RELATED WORK

A wealth of energy consumption data has become available with the advent of smart meters, facilitating significant advancements in prediction models. This section recalls related works in the field, categorizing them into key areas that reflect the evolution and current trends in electricity consumption forecasting. The reviewed works span from early developments in utilizing smart meter data to advanced machine learning models, behavioural analytics, and innovative hybrid approaches, showcasing the collective contributions and emerging trends in this domain.

### A. EARLY DEVELOPMENTS IN SMART METER DATA UTILIZATION

Gajowniczek and Ząbkowski [15] were among the first to leverage time series data from smart meters to enhance forecasting models. By employing Artificial Neural Networks (ANNs) and Support Vector Machines (SVMs), they demonstrated the potential of using individual smart meter data for short-term electricity use predictions. This work laid the ground for more sophisticated analyses of energy consumption patterns. Similarly, Ilić et al. [16] investigated the impact of smart meter data aggregation on forecasting accuracy. Their study emphasized the potential of fine-grained smart meter data to enhance forecast precision by

identifying highly predictable consumer groups, which is desirable for efficient energy management and Virtual Power Plants (VPPs). It also highlighted the challenges of predicting individual consumer behaviour and the limitations of random aggregation methodologies.

### B. BEHAVIORAL AND SOCIETAL IMPACT ANALYSIS

Starting in 2016, Zhou and Yang [17] investigated the use of big data analytics from smart meters to improve electricity consumption forecasting. Their work integrated insights from energy social science, social informatics, and energy informatics, focusing on household energy use behaviours. They also addressed interdisciplinary methodologies and data privacy challenges, laying a thorough understanding of how household behaviours influence energy usage patterns. In 2019 and 2020 many studies, such as [18] and [19], used predictive modelling methods that focused on using an ANN based model to learn how different customers would react to changes in electricity prices. These studies further advanced the understanding of consumer behaviour in energy consumption, emphasizing the responsiveness to economic incentives. By 2021, Brugger et al. [20] expanded the scope by investigating the impact of emerging societal trends like digitalization, the sharing economy, and increased consumer awareness on future energy demands in European countries. Through their research, they came up with four different energy demand scenarios for 2050. These showed how important policymaking is in shaping these trends and what those changes mean for predicting energy use. This gave them a full picture of how societal factors and technological advances affect predicting energy use.

### C. CLUSTERING TECHNIQUES FOR ENHANCED FORECASTING

This critical section of research focuses on refining forecasts by identifying typical daily consumption profiles through clustering techniques and integrating behavioral insights. Shahzadeh et al. [21] improved load forecast accuracy by clustering consumers into distinct groups based on detailed smart meter data, training individual neural networks for each cluster, and highlighting the critical role of feature selection in clustering for enhanced forecast accuracy. Similarly, Yildiz et al. [22] leveraged historical smart meter data to refine household electricity load forecasts through a dual method approach that combined the Smart Meter Based Model (SMBM) for short-term prediction with the Cluster-Classify-Forecast (CCF) method, which identifies typical daily consumption profiles. Wen et al. [23] further developed this approach by introducing a clustering method relying on an improved  $k$ -means algorithm integrated with Principal Component Analysis (PCA) for dimensionality reduction. They employed Dynamic Time Warping (DTW) for effectively grouping similar electricity consumption patterns, thus providing valuable insights for Demand Side

Management (DSM) and personalized energy-saving strategies. Extending these methodologies, Cz  t  ny et al. [24], and Michalakopoulos et al. [25] employed a range of clustering techniques, including  $k$ -means, fuzzy  $k$ -means, Hierarchical Agglomerative Clustering (HAC), and Density-based Spatial Clustering of Applications with noise (DBSCAN), to develop detailed consumption profiles that categorize nearly a thousand households. Their work emphasizes the impact of factors like meter type, day, seasonality, and housing type on consumption patterns, which is crucial for tailoring Demand Response (DR) programs and improving forecasting accuracy. Both sets of studies demonstrate how important it is to (i) include human behaviour in forecasting models, (ii) identify the best consumer groups, and (iii) use clustering techniques in a broader way to understand how people use energy. Studies like [26] and [27] have also used profile clustering analysis to look into how customers behave when they use energy. This study shows that behavioural insights are important to improve load forecasting and energy management strategies. These efforts highlight the critical role of clustering and profiling in understanding consumer behaviour, forecasting electricity consumption accurately, and optimizing energy utilization strategies.

### D. ADVANCEMENTS IN DEEP LEARNING MODELS

The advancements in deep learning models for electricity consumption forecasting emphasize the integration of novel techniques and algorithms to enhance prediction accuracy. Li et al. [28] proposed an optimized ANN model refined with an improved Particle Swarm Optimization (PSO) algorithm and PCA, aiming to improve short-term electricity consumption predictions. In parallel, Chou and Tran [29] highlighted the superior accuracy of hybrid Machine Learning (ML) models that combine forecasting with optimization techniques, addressing the challenges of data quality in residential energy consumption patterns.

The application of sequence-to-sequence architectures using both Gated Recurrent Unit (GRU) and LSTM models by Sehovac et al. [30] as well as the integration of CNN with LSTM by Lotfipoor et al. [31] and Abraham et al. [32], marked advancements in handling time series data effectively for short-term energy forecasting. Additionally, Mahjoub et al. [33] who explored LSTM models in comparison to GRU and Drop-GRU configurations, further highlight the versatility of these models in adapting to diverse temporal patterns in energy forecasting. These hybrid models have shown notable improvements over traditional techniques such as ARIMA by achieving better error metrics like Mean Absolute Error (MAE) and Mean Squared Error (MSE).

Furthermore, the development of novel Transformer-based models by Neo Wu et al. [34], including the Temporal Fusion Transformer (TFT) by Bryan Lim et al. [35] and the Informer model by Haoyi Zhou et al. [36], has introduced self-attention mechanisms that capture complex patterns and long-term dependencies effectively, enhancing both univariate and

multivariate forecasting accuracy. These models emphasize the versatility and robustness necessary for managing large datasets from smart meters and improving long sequence predictions through innovative techniques such as ProbSparse self-attention and self-attention distilling.

Additionally, the work by Fekri et al. [37] on an Online Adaptive RNN and Mahjoub et al. use of LSTM models [33] have demonstrated the adaptability of deep learning approaches to continuously learn and adapt to new data in real-time, addressing challenges such as concept drift and changing consumption patterns. Finally, recent methodologies like the novel discrete deep learning-based methodology for energy consumption classification by Khashei et al. [38], the Channel Aligned Robust Blend Transformer (CARD) model by Wang et al. [39], and the Pathformer by Chen et al. [40] have each introduced discrete adjustments and multi-scale architectures that improve time series forecasting. These models utilize channel-aligned attention mechanisms, token blend modules, and adaptive pathways to achieve high accuracy and robustness against overfitting and variable dependencies in energy consumption datasets.

Collectively, these advancements highlight a significant shift towards sophisticated, data-driven deep learning frameworks that not only refine the accuracy of short and long-term forecasts, but also incorporate comprehensive behavioural analyses and adaptability to evolving energy consumption patterns.

### E. BAYESIAN AND PROBABILISTIC MODELS

Advancements in Bayesian and probabilistic models have significantly improved electricity consumption forecasting using smart meter data. Taieb et al. [41] developed an additive quantile regression model for forecasting the distribution of future electricity demand at the household level, addressing the volatility of smart meter data. Concurrently, Wang et al. [42] enhanced hourly energy predictions in educational buildings using a Random Forest algorithm, which outperformed traditional models like Regression Tree and Support Vector Regression. Furthering this, another study by Wang et al. [43] introduced Constrained Quantile Regression Averaging (CQRA) to optimally combine probabilistic forecasts and minimize pinball loss, demonstrating effectiveness with datasets from ISO New England and Irish smart meters. In [44], Weeraddana et al. used an approach based on a stacked nonparametric Bayesian model with Gaussian Processes to tackle non-linear and non-stationary consumption patterns, surpassing traditional methods like ARIMA and RF in handling short time series data. Additionally, Xu et al. [45] explored monthly residential electricity consumption using a clustering-based probability distribution model, further illustrating the potential of probabilistic models in robust electricity consumption forecasting. These developments highlight a move towards forecasting models that integrate diverse data sources to enhance prediction accuracy and robustness.

### F. HYBRID AND ENSEMBLE MODELS

The development of hybrid and ensemble models represents a novel shift. Researchers have embraced these approaches, combining advanced machine learning techniques to significantly enhance prediction accuracy and effectively address the complexities inherent in data characteristics. Tian et al. [46] introduced a pioneering method using Generative Adversarial Network (GAN) to forecast energy consumption, which improves prediction by combining original and synthetically generated data, thus addressing data scarcity and enhancing forecast accuracy. Similarly, Alhussein et al. [47] developed a hybrid model that combines CNN for feature extraction and noise filtering with LSTM for capturing temporal dependencies. This model demonstrated superior accuracy over traditional LSTM models by achieving lower Mean Absolute Percentage Error (MAPE), particularly in predicting individual household loads.

Furthering the use of hybrid approaches, Bian [48] enhanced short-term electricity consumption forecasting by integrating  $k$ -means clustering and Fuzzy  $C$ -Means (FCM) with a neural network. This method clusters users based on consumption patterns before forecasting. Similarly, Khan et al. [49] developed an ensemble forecasting model that combines spatial-temporal clustering with LSTM and GRU techniques, utilizing high-resolution smart meter data to enhance prediction accuracy in residential settings.

Moreover, Dai and Meng [50] introduced a model that integrates what they call online model-based functional clustering with functional deep learning, employing a dynamic clustering mechanism and an adaptive functional deep neural network to adapt to time-varying consumption patterns and improve daily load demand forecasts. Adding to the diversity of hybrid models, Saoud et al. [51] proposed a forecasting model that merges the Stationary Wavelet Transform (SWT) with transformers. This method decomposes consumption data into high and low-frequency components, allowing to effectively learn and predict complex patterns. Demonstrated to surpass existing models in prediction accuracy and robustness, their approach also exhibits resilience to noise and disturbances.

### G. EMERGING TRENDS AND BIG DATA INTEGRATION

In exploring the intersection of emerging trends and big data integration in energy consumption forecasting, researchers are harnessing sophisticated data analytics to enhance accuracy and manage complex datasets effectively. Kalksma et al. [52] utilized modified Support-Pruned Markov Models (SPMM) to predict appliance usage and energy footprints with high accuracy, showcasing the impact of sequential pattern mining in optimizing energy management despite its data-intensive demands. Similarly, Oprea and Băra [53] leveraged scalable big data frameworks integrated with machine learning algorithms for effective short-term load forecasting in residential settings. Their approach not only improves forecast accuracy but also adapts to changing

conditions by processing data from smart meters and weather sensors using a NoSQL database.

These examples illustrate a broader trend in the field: the integration of Artificial Intelligence of Things (AIoT) for real-time data processing [54], the application of edge computing to reduce latency, the use of digital twins for scenario simulation, and blockchain for secure energy transactions. Additionally, developments in machine learning models, including deep and reinforcement learning, and integrative data platforms are enhancing the predictive capabilities of energy systems, while predictive maintenance and regulatory analytics ensure optimal and compliant operations. These advancements highlight a significant shift towards sophisticated hybrid and ensemble frameworks that leverage diverse data sources and modelling techniques, refining the accuracy and adaptability of forecasts in smart grid systems.

### H. COLLECTIVE CONTRIBUTIONS AND EMERGING TRENDS

All of this related work showcases notable progress in the domain of energy consumption prediction through the utilization of data from smart meters. Initial research demonstrated the viability of utilizing smart meter data for near-term forecasts, but following investigations used clustering methods, other machine learning models, and hybrid approaches to improve precision and dependability. Current developments exploit the availability of high-resolution data by employing innovative techniques, while also highlighting the flexibility of models to adjust to different circumstances. These contributions highlight the continuous development and improvement of forecasting methods, which provide more effective and environmentally friendly energy management practices.

## III. PROBLEM STATEMENT

### A. CHALLENGES

Accurately forecasting electricity consumption at the household level has become increasingly challenging, owing to the dynamic and complex nature of modern energy usage. Based on the findings from the literature review, traditional forecasting methodologies are becoming less effective as they struggle to incorporate the impact of new technologies like electric vehicles, which often charge overnight, and residential solar panels, which introduce significant variability in energy production and consumption. These challenges are further complicated by a myriad of factors including, but not being limited to, household demographics, appliance usage patterns, meteorological conditions and economic variabilities, each adding layers of complexity to energy management.

Moreover, existing forecasting models whilst proficient at aggregating data for broad community-level analysis, consistently underperform when tasked with applying these insights to individual households. This shortfall primarily stems from two interconnected issues: the models' inability to

accurately manage the abrupt daily variations in energy usage and their failure to develop tailored forecasting solutions for each consumer, given the considerable diversity in behaviour and consumption patterns across households. Noise plays a crucial role, since measurements of individual households may have unpredictable behaviour, which averages out when considering, instead, broad-community level data.

These limitations are exacerbated as the granularity of data collection increases through advanced metering technologies, which unmask the inadequacies of traditional models in capturing the nuanced and varied energy consumption dynamics unique to each household.

### B. METHODOLOGICAL INNOVATIONS

To address these deficiencies, this study introduces a novel forecasting methodology that leverages  $k$ -means clustering to segment the consumer base into distinct groups based on their energy consumption behaviours observed throughout the years. This method does not merely aggregate consumers based on superficial similarities but delves deeper into consistent behavioural patterns, identifying clusters of users who exhibit virtually identical energy usage profiles over extended periods. By pinpointing these distinct clusters, our study leverages the capabilities of transformer models, renowned for their proficiency in handling sequential time-series data. These models are adept at capturing the complex temporal dynamics within each cluster, thus enabling us to apply a specific forecasting approach tailored to the nuanced patterns of each group. This method significantly enhances the accuracy of our forecasts. Our approach is detailed in Figure 1, which outlines the progression from data loading to the application of the JITrans model.

Our innovative approach not only improves the granularity and accuracy of energy consumption forecasts but also supports utility providers in managing demand surges and integrating renewable energy sources more effectively. By refining forecasting accuracy for specific clusters of households, this research contributes significantly to the development of adaptive, efficient, and sustainable energy systems.

This study fills a gap in the existing literature by providing a scalable and practical solution to a pressing global challenge, marking an advancement in the field of energy forecasting. The following section, details the technical specifications and operational framework of the JITrans model utilized in our forecasting methodology.

## IV. PROPOSED MODEL ARCHITECTURE AND FORECASTING STRATEGY

The main contribution of this work lies in the development of a modular forecasting system using a set of transformer models, each optimized for different forecast horizons. This approach differs from traditional methods that use a single model to forecast over a fixed or variable horizon. Instead, each model within the ensemble is trained with a distinct length of the decoder input, specifically designed to predict

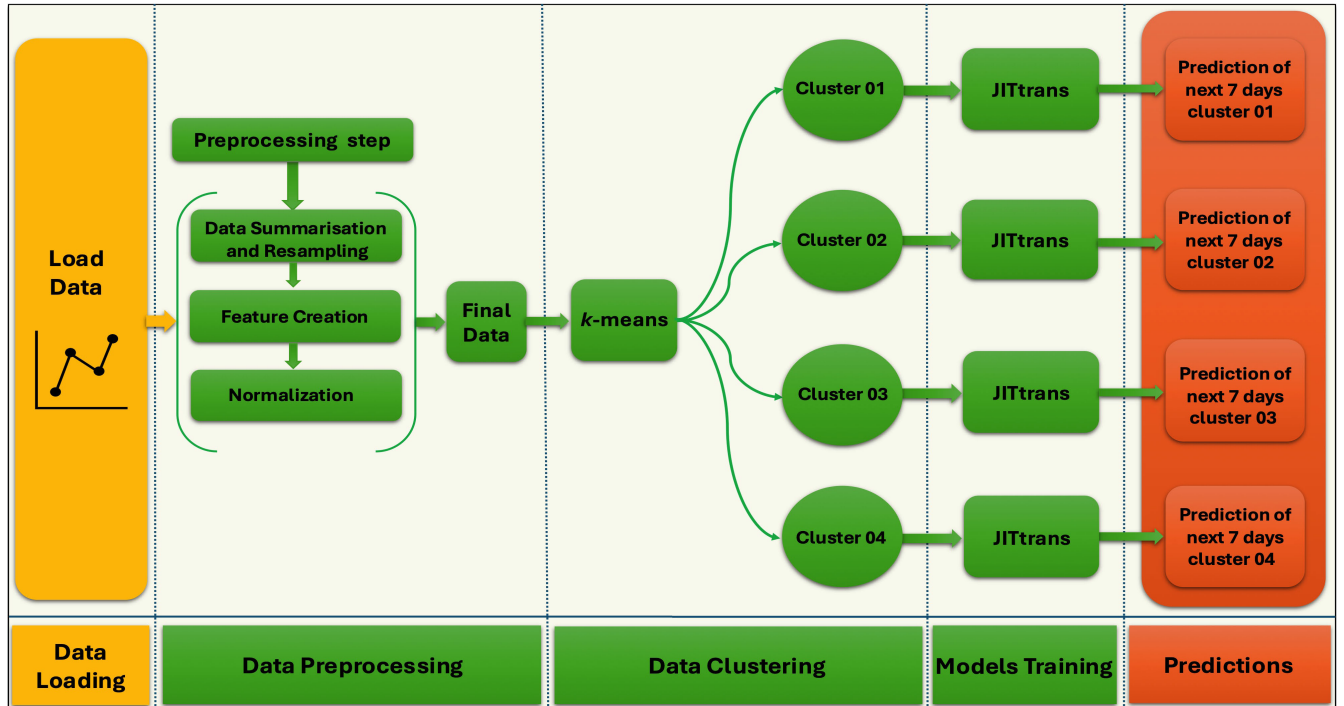


FIGURE 1. Overview of the proposed forecasting methodology.

the energy consumption at different lead time. This allows the system to specialize in capturing temporal dynamics unique to different forecast ranges, thereby improving accuracy and robustness compared to a one-size-fits-all model.

During the prediction phase, the JITrans model leverages a sequential and accumulative prediction strategy. Starting from the shortest horizon model, predictions are generated and used as inputs to the next model in the sequence. This cumulative approach refines the forecast at each step by integrating the predictions from all previous models, thus enhancing the overall forecast as the horizon extends.

We begin by recalling the technical details related to  $k$ -means and transformers, then we introduce our modular forecasting system.

### A. K-MEANS

Clustering is a technique used to group a set of objects in such a way that objects in the same group (called a cluster) are more similar to each other than to those in other groups. It's a method of unsupervised learning, and it's particularly useful in identifying distinct patterns from large datasets without prior labelling of the data. In the context of electricity consumption, clustering algorithms can identify and group similar usage patterns across different consumers. This is particularly useful for utility companies to tailor customer engagement strategies and optimize the generation and distribution of power.

Representative-based clustering seeks to partition a dataset  $D$ , consisting of  $n$  points  $\mathbf{x}_i$ , into  $k$  clusters ( $C = \{C_1, C_2, \dots, C_k\}$ ) within a  $d$ -dimensional space. Each

cluster  $C_i$  is characterized by a representative point, a.k.a. the centroid  $\mu_i$ , which is the mean of all points in the cluster.

$$\mu_i = \frac{1}{n_i} \sum_{\mathbf{x} \in C_i} \mathbf{x}, \tag{1}$$

where  $n_i = |C_i|$  denotes the number of points in cluster  $C_i$ .

The objective of clustering is to minimize the within-cluster variance, formulated as:

$$\mathcal{L} = \sum_{i=1}^k \sum_{\mathbf{x} \in C_i} \|\mathbf{x} - \mu_i\|^2, \tag{2}$$

A brute-force approach to determine an optimal clustering involves generating all possible partitions of the  $n$  points into  $k$  clusters, evaluating their optimization scores, and selecting the partition with the best score. The exact number of ways to partition  $n$  points into  $k$  non-empty, disjoint subsets is given by the Stirling numbers of the second kind, which gives a total number of unique clustering in the order of  $O\left(\frac{k^n}{k!}\right)$ . Clearly, counting and evaluating all these clusterings is impractical. Therefore, more efficient algorithms such as  $k$ -means and expectation-maximization are used in practice for representative-based clustering.

$k$ -means [55] is an iterative algorithm performing two main steps:

- 1) **Assignment Step:** Each data point is assigned to the nearest cluster centroid based on the Euclidean distance;
- 2) **Update Step:** The centroids of the clusters are updated (they are computed as the mean of all data points assigned to their respective clusters).

These steps are repeated until convergence, which is typically achieved when the centroids no longer change significantly between iterations. The algorithm may converge to a local minimum, and hence it is often run multiple times with different initializations to find the best clustering [56].

$k$ -means is computationally efficient and suitable for large datasets, although the quality of the clustering depends on the initial placement of the centroids. The time complexity of the algorithm is  $O(nkd)$ , where  $n$  is the number of data points,  $k$  is the number of clusters,  $d$  is the number of dimensions, and  $t$  is the number of iterations [57].

## B. TRANSFORMERS

Transformers are neural network architectures designed to handle sequential data efficiently across various applications. Unlike traditional models that process sequences of data in sequential order, transformers utilize the positional embedding in order to process all *tokens* in the sequence simultaneously; tokens are discrete units representing the items in the sequence, they can be patches of an image, words in a sentence or numerical values in a time series. This parallel processing capability significantly enhances the scalability and efficiency of the model with respect to sequence length and complexity. The transformer processes the input data through a sequence matrix  $\mathbf{X}$ , which undergoes transformations via self-attention mechanisms. The aim of self-attention mechanism is to obtain a finer representation of the current token by means of a weighted average of the most relevant (i.e. the most correlated with the current one) tokens in the sequence. The mechanism is based on the use of so-called *queries*  $\mathbf{Q}$ , *keys*  $\mathbf{K}$  and *values*  $\mathbf{V}$  matrices, which are different representations of the input tokens,  $\mathbf{X}$ . Namely, for each token in  $\mathbf{Q}$  we compute its correlation with all the tokens in  $\mathbf{K}$  as result of the dot product  $\mathbf{Q}\mathbf{K}$ ; this correlation is then used to obtain the finer representation of the current token as a weighted sum of the tokens in  $\mathbf{V}$ .  $\mathbf{Q}$ ,  $\mathbf{K}$ , and  $\mathbf{V}$  are derived from the input matrix  $\mathbf{X}$  by using different learnable weights:

$$\mathbf{Q} = \mathbf{X}\mathbf{W}^Q, \quad \mathbf{K} = \mathbf{X}\mathbf{W}^K, \quad \mathbf{V} = \mathbf{X}\mathbf{W}^V. \quad (3)$$

where  $\mathbf{W}^Q$ ,  $\mathbf{W}^K$  and  $\mathbf{W}^V$  are learnable weights specific to queries, keys, and values, respectively. These weights matrices project each input token into a representation according to its role as a key, query, or value.

The self-attention mechanism adjusts the attention each token receives based on the query-key affinity and then uses this attention to form a weighted sum of value vectors. The computation of self-attention requires determining first the attention weights between each pair of input tokens  $\mathbf{x}_n, \mathbf{x}_m$ :

$$\begin{aligned} a[\mathbf{x}_n, \mathbf{x}_m] &= \text{softmax}_m \left( \frac{\mathbf{k}_m^T \mathbf{q}_n}{\sqrt{d_k}} \right) \\ &:= \frac{\exp \left( \frac{\mathbf{k}_m^T \mathbf{q}_n}{\sqrt{d_k}} \right)}{\sum_{m'=1}^N \exp \left( \frac{\mathbf{k}_{m'}^T \mathbf{q}_n}{\sqrt{d_k}} \right)}. \end{aligned} \quad (4)$$

The softmax function ensures that the attention scores, representing the influence weights, are non-negative and sum to 1. The scaling factor  $\sqrt{d_k}$ , where  $d_k$  is the dimension of the token space, is used to help stabilizing the gradients during training. Self-attention is then computed as

$$\mathbf{S}\mathbf{a}(\mathbf{X}) = \mathbf{V} \cdot \text{softmax} \left( \frac{\mathbf{K}^T \mathbf{Q}}{\sqrt{d_k}} \right). \quad (5)$$

### 1) VANILLA TRANSFORMER

The vanilla transformer model, introduced by Vaswani et al. [8], may be based on an encoder, a decoder or both. Each of these components consists of multiple identical layers. An encoder layer includes a multi-head self-attention mechanism<sup>2</sup> followed by a position-wise feed-forward network. In contrast, each decoder layer incorporates an additional cross-attention mechanism (i.e., the decoder embeddings attend to the encoder embeddings) between the multi-head self-attention and the feed-forward network. This dual-attention mechanism allows the model to effectively capture and utilize information from the entire input sequence [8].

### 2) INPUT ENCODING AND POSITIONAL ENCODING

To retain the sequential order information, the transformer model employs positional encoding added to the input embeddings. This approach is crucial as the self-attention mechanism itself is permutation invariant, meaning it does not inherently capture the order of the sequence.

- 1) **Absolute Positional Encoding:** In the vanilla Transformer, positional encoding vectors are generated using sine and cosine functions of different frequencies for each dimension, as defined by the following equations:

$$PE_{(pos,i)} = \begin{cases} \sin \left( \frac{pos}{10000^{\frac{i}{d}}} \right) & \text{if } i \% 2 = 0 \\ \cos \left( \frac{pos}{10000^{\frac{i}{d}}} \right) & \text{if } i \% 2 = 1 \end{cases}, \quad (6)$$

where  $pos$  is the position,  $i$  is the dimension, and  $d$  is the total number of dimensions. This encoding allows the model to incorporate information about the position of each token in the sequence, enabling it to distinguish between tokens in different positions.

- 2) **Relative Positional Encoding:** Some methods focus on pairwise positional relationships between input elements, adding learnable relative positional embeddings to the keys of the attention mechanism [58]. There are also hybrid approaches that combine both absolute and relative positional encodings [59].

### 3) MULTI-HEAD ATTENTION

The multi-head attention mechanism is a cornerstone of the transformer architecture, enhancing the model's ability to focus on different parts of the input sequence simultaneously.

<sup>2</sup>A single self-attention mechanism is not enough for a transformer to work correctly, therefore multiple attention mechanisms are used; this is called multi-head self attention.

This mechanism employs multiple attention heads, each learning to attend to different parts of the input. Multi-head attention extends the self-attention mechanism by using multiple sets of learned projections:

$$\text{MultiHeadAttn}(Q, K, V) := \text{Concat}(T_1, \dots, T_h)W^O, \quad (7)$$

where each attention head  $T_i$  is computed as:

$$T_i = \text{Attention}(QW_i^Q, KW_i^K, VW_i^V). \quad (8)$$

In this formulation,  $W_i^Q$ ,  $W_i^K$ , and  $W_i^V$  are the learnable weight matrices for the  $i$ -th head, and  $W^O$  is a matrix of learnable weights that projects the concatenated outputs of the multiple attention heads back to the original space. This structure allows the model to jointly attend to information from different representation subspaces at different positions, significantly enhancing its learning capacity.

#### 4) FEED-FORWARD AND RESIDUAL NETWORK

The feed-forward network in a transformer layer is a fully connected network applied to each position separately and identically:

$$\text{FFN}(H_0) := \text{ReLU}(H_0W_1 + b_1)W_2 + b_2, \quad (9)$$

where  $H_0$  is the output of the previous layer,  $W_1 \in \mathbb{R}^{D_m \times D_f}$ ,  $W_2 \in \mathbb{R}^{D_f \times D_m}$ ,  $b_1 \in \mathbb{R}^{D_f}$ , and  $b_2 \in \mathbb{R}^{D_m}$  are trainable parameters. Each layer in a transformer also includes residual connections and layer normalization to ensure better gradient flow and faster convergence:

$$H_0 = \text{LayerNorm}(\text{Attention}(Q, K, V) + X), \quad (10)$$

$$H = \text{LayerNorm}(\text{FFN}(H_0) + H_0), \quad (11)$$

where  $\text{LayerNorm}(\cdot)$  denotes layer normalization.

#### 5) RECENT INNOVATIONS IN TRANSFORMER MODELS

The transformer model, introduced by Vaswani et al. [8], represents a groundbreaking advancement in sequence modelling, utilizing an encoder-decoder framework. This model has achieved exceptional results across various fields, including NLP [60], speech recognition [61] and computer vision [62]. More recently, transformer-based models have become increasingly popular for time series analysis [63].

In 2024, several new transformer models have emerged, demonstrating promising results in time series forecasting. These models include:

- **CARD** [39]: Employs channel alignment to robustly enhance forecasting accuracy;
- **Pathformer** [40]: Utilizes adaptive pathways to efficiently manage multi-scale data;
- **GAFormer** [64]: Enhances time series understanding through group-aware embeddings;
- **Transformer-Modulated Diffusion Models for Probabilistic Multivariate Time Series Forecasting** [65]: Combines Transformer architectures with diffusion models for probabilistic forecasting;

- **iTransformer** [66]: Proposes an inverted Transformer architecture for better forecasting performance.

These advancements illustrate the continued evolution and adaptation of transformer models to various complex tasks in time series forecasting.

#### C. MODEL DESIGN AND FUNCTIONALITY

We now introduce the design of our JITrans transformer model. The architecture of the proposed ensemble includes multiple transformer models, each specifically trained to predict a distinct day within a 7-day forecast horizon. This setup leverages an input sequence (denoted by  $x$ ) shared among the encoders of each transformer and an input sequence (denoted by  $d$ ) for each decoder. The input sequence to the decoder is progressively extended in size moving from one transformer to the next one, adding the forecasted value obtained by the previous transformer. In our configuration, each transformer model in JITrans uses a sequence of 30 days as input to each encoder (from  $x_{t-29}$  to  $x_t$ ), providing the necessary historical context for accurate forecasting. The decoder input and the targets for each model are structured to sequentially extend the forecast range by concatenating predictions from earlier models.

Denoting by  $\hat{y}_{t+k}^{(i)}$  the forecast for day  $t+k$  of the  $i$ -th model, the first model utilizes a decoder input of 7 days from  $d_t$  to  $d_{t+6}$ , and it is trained to output days from  $d_{t+1}$  to  $d_{t+6}$  and day  $\hat{y}_{t+7}^{(1)}$ , which is a new prediction, not included in the decoder input.

Each subsequent model extends the input to its decoder by adding the day forecasted by the previous model. For example, Model 2 uses a sequence of 8 days as input to the decoder (from  $d_t$  to  $d_{t+6}$  and  $\hat{y}_{t+7}^{(1)}$ ), incorporating the forecast  $\hat{y}_{t+7}^{(1)}$  predicted by Model 1. Model 2 then makes predictions  $\hat{y}_{t+7}^{(2)}$  and  $\hat{y}_{t+8}^{(2)}$ , with  $\hat{y}_{t+8}^{(2)}$  being a new day added to the forecast and  $\hat{y}_{t+7}^{(2)}$  being an updated forecast of day  $\hat{y}_{t+7}^{(1)}$  based on additional input data.

Figure 2 provides a schematic representation of the JITrans model architecture, illustrating how data flows through the encoder and decoder inputs across the ensemble, thereby clarifying the model's operational framework.

#### D. INTEGRATION AND AVERAGING OF MODEL OUTPUTS

The ensemble uses the following strategy to integrate and refine predictions: each model contributes to the forecast of subsequent days by utilizing a combination of direct input from its decoder and the averaged predictions from previous models. For instance, Model 3, when predicting  $\hat{y}_{t+9}$ , will use the day  $\hat{y}_{t+8}$  directly from Model 2 and take an average of the predictions for  $\hat{y}_{t+7}$  from both Model 1 and Model 2. This averaging mechanism is formally expressed as:

$$\hat{y}_{t+k}^{(j)} = \frac{1}{j-k} \sum_{i=k}^{j-1} \hat{y}_{t+k}^{(i)}, \quad (12)$$



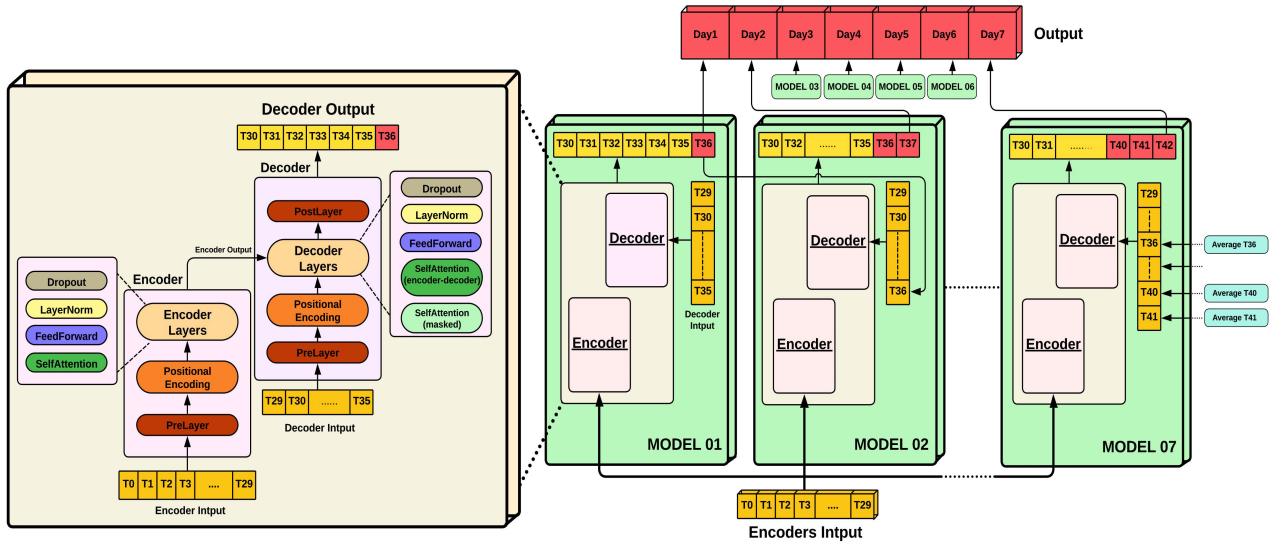


FIGURE 2. JITrans model architecture.

where  $\bar{y}_{t+k}^{(j)}$  denotes the averaged value used as input to  $j$ -th model for day  $t + k$  and  $\hat{y}_{t+k}^{(i)}$  is the forecast for day  $t + k$  of the  $i$ -th model.

This ensemble aligns with deep learning strategies by employing multiple learners to inform and enhance the decision-making process.

By leveraging the strengths of multiple specialized transformer models, this novel approach offers several advantages:

- **Enhanced Accuracy:** Averaging predictions across multiple models reduces the impact of outliers and potential errors from any single model’s forecast;
- **Increased Robustness:** The approach mitigates the risk of overfitting to recent data trends by integrating broader historical perspectives, leading to more stable and reliable predictions;
- **Scalability and Flexibility:** The architecture is designed to be scalable, allowing for the addition of more models to cover longer forecasting horizons or to adapt to changes in the underlying data patterns.

## E. DATASET OVERVIEW AND DATA PREPROCESSING

### 1) INITIAL DATASET OVERVIEW

Our study utilizes a comprehensive dataset spanning 28 months, from December 2019 to March 2022, comprising hourly electricity consumption records from 321,561 residential customers. Each record includes a timestamp, customer identification number, and the amount of electricity consumed in the last hour, tallying approximately 20,400 observations per customer.

### 2) PRELIMINARY DATA CLEANING

To ensure the data’s integrity and applicability for residential consumption analysis, we conducted extensive preliminary data cleaning:

- 1) **Inactive or Discontinued Usage:** We removed records from consumers who did not have readings for a continuous 24-hour period.
- 2) **Unusual Consumption Patterns:** Entries showing unusually low consumption were excluded, as they were probably from unused households. Similarly, entries exceeding 12 kWh per hour were removed, since these are indicative of industrial or commercial load, rather than residential usage.

This targeted cleaning refined the dataset, allowing us to focus the analysis on genuine residential consumption patterns and reducing the initial dataset to 73,802 active consumer profiles with typical residential energy usage.

### 3) ADVANCED DATA PREPROCESSING

With a refined dataset of consumer profiles, further pre-processing was applied to prepare the data for the detailed analysis carried out by our research:

- 1) **Noise Reduction and Smoothing:** We employed a simple moving average (SMA) with a 7-day window to smooth the daily electricity consumption data, reducing short-term fluctuations and highlighting long-term trends;
- 2) **Handling Missing Values:** The dataset used in our analysis was well-prepared, with no missing values present. This completeness is a result of the preliminary data cleaning process, which ensured that all data used for the study was continuous and complete, eliminating the need to fill in any gaps;
- 3) **Feature Engineering:** From the basic raw data features, i.e. timestamp, customer identification number and hourly electricity consumption, we derived additional features such as day of the week, month, and year. These details were carefully extracted from the timestamps to show how electricity use changes with

the seasons and on a regular basis. This made our analysis much more useful and in-depth by explicitly adding a new context that was only latent in the raw data.

In particular, we extracted time-based features from the 'Date\_Time' column to capture temporal patterns:

- Day of the Week ('day\_of\_week'): Indicates the day of the week (0 for Monday, 6 for Sunday). This feature helps the model learn weekly patterns and trends in the data;
- Month ('month'): Captures monthly seasonal effects, which can influence the target variable due to factors like holidays or seasonal demand;
- Year ('year'): Useful for identifying long-term trends or shifts over different years.

From these features, we created the following set of context features: `context_features = ['SMA_7', 'day_of_week', 'month']`. We did not include 'year' because it was less relevant, given that our dataset spans only 28 months. 'SMA\_7' represents the consumption readings in kWh. It is included in the context features and also serves as the target variable to be predicted.

Using these three features, we created a contextual feature that represents the correlation between 'SMA\_7', 'day\_of\_week', and 'month'. To achieve this, we applied Principal Component Analysis (PCA) to reduce the dimensionality of the context features from three to one principal component. The benefits of this approach include: (i) simplification: reduces the complexity of the data, making model training more efficient; (ii) variance retention: captures the most significant variance in the data with minimal loss of information; (iii) multicollinearity reduction: mitigates issues arising from highly correlated features. The reduced context data was included as an additional input to the model, potentially enhancing its predictive capabilities by incorporating key contextual information. Our final feature set used for modeling was `['SMA_7', 'day_of_week', 'context_data_reduced']`.

- 4) **Normalization and Scaling:** The RobustScaler was used on the smoothed consumption data to minimize the influence of outliers. Unlike standard scaling methods, RobustScaler uses the median and interquartile range (IQR), making it less sensitive to outliers. We scaled our 'SMA\_7' feature, which represents the consumption reading values. Since our data are raw and include not only households but also small industries and the overall regional load, the 'SMA\_7' feature may contain outliers due to fluctuations in the time-series data. Using RobustScaler ensures that these outliers do not unduly influence the scaling process.

The StandardScaler was applied to other contextual features to standardize the data. Basically, this consists in computing a standard z-score (subtracting the mean and normalizing by dividing by the standard deviation, scaling to unit variance). We scaled the 'day\_of\_week'

feature and the context features. This ensures that the 'day\_of\_week' feature is normalized, making it easier for the model to learn patterns. The 'day\_of\_week' feature is ordinal and ranges from 0 to 6. Standardizing this feature helps in treating all days equally in terms of scale, which will improve the model's performance;

- 5) **Dimensionality Reduction:** Principal Component Analysis (PCA) was employed not just to reduce the complexity of our feature space, but also to create contextual features that capture essential variations, as previously explained. These PCA-derived features were then added to our model as new variables, enhancing the predictive capability by providing a condensed representation of critical data dimensions, which improves the effectiveness of our analysis.

#### 4) CLUSTERING FOR DETAILED ANALYSIS

The cleaned and transformed data were used to perform clustering on the selected 73,802 consumer profiles. This step involved grouping consumers into distinct categories based on their energy consumption behaviours, which is fundamental for subsequent tailored analysis and for developing specific energy conservation strategies.

#### 5) DATA MODELING

The thoroughly preprocessed and clustered data provided a robust foundation for deploying sophisticated predictive models. These steps ensured high-quality data inputs, enabling more reliable and accurate forecasts of residential energy consumption patterns.

### F. TRAINING THE MODEL

In this section, we provide details related to the training of our Just In Time transformer model, whose code is freely available.<sup>3</sup> The number of parameters to be learned is slightly more than 125 million. These have been learned by using the dataset, which includes hourly measurements related to 73,802 customers. We have aggregated the available measurements by computing the average daily consumption so that the dataset provides for each user his/her mean consumption for a period spanning 852 days.

The dataset itself has been split as follows in order to train the model: 80% for training, 10% for validation, and 10% for testing. The chosen loss function and the optimizer algorithm were, respectively, Mean Squared Error (MSE) and Adam. For the systematic and efficient selection of the optimal hyperparameters, we utilized Optuna, a hyperparameter optimization framework. Optuna employs Bayesian optimization techniques to explore the parameter space and has been integral in determining the best set of parameters for our model on the basis of the performance metrics defined in our experimental setup.

<sup>3</sup><https://github.com/cafar/JIT-Transformer>

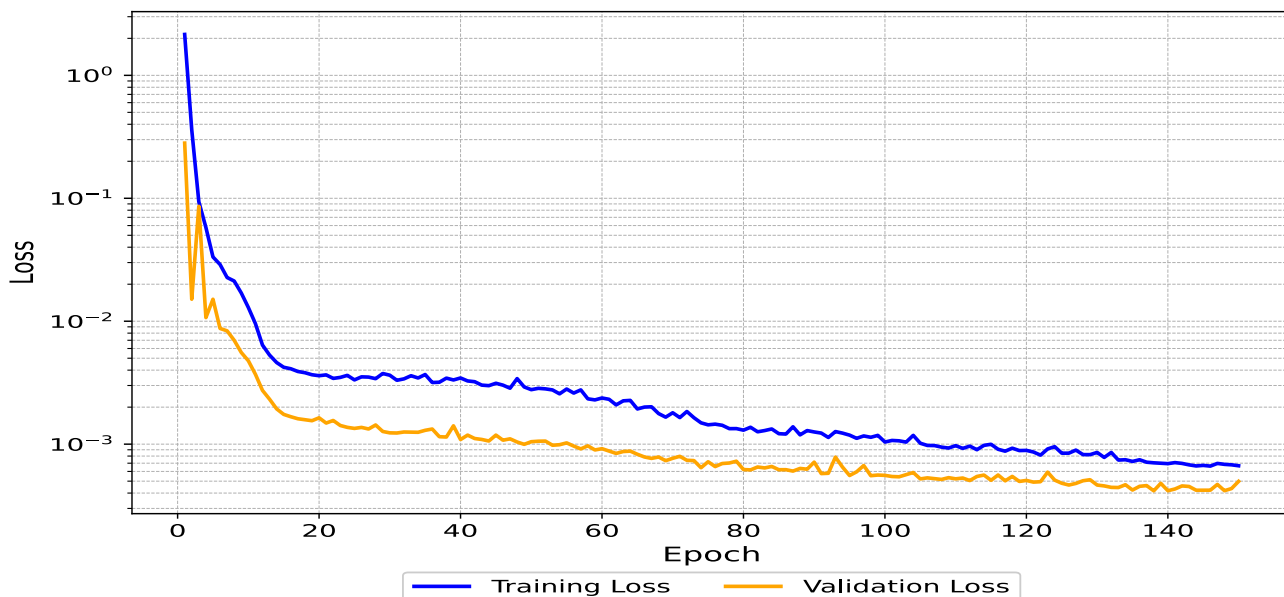

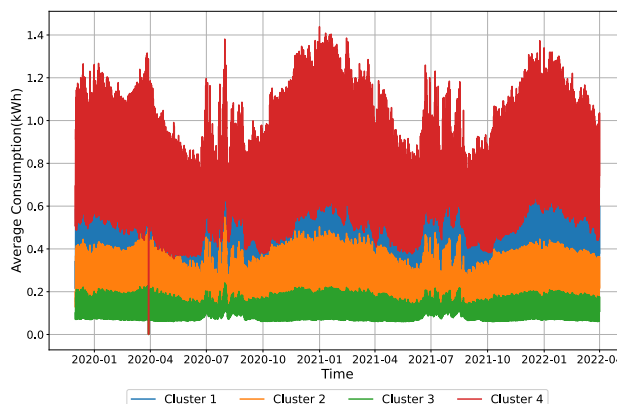


FIGURE 3. JITrans training and validation losses.



(a) Elbow method to establish the optimal number of clusters.



(b) Average consumption patterns for each cluster.

FIGURE 4. Clustering and average consumption patterns.

The model has been trained on the Italian Leonardo supercomputer<sup>4</sup> for 150 epochs with a batch size equal to 64 and a learning rate initially set to 0.0001, using one NVIDIA Ampere A100 Tensor Core GPU equipped with 64 GB HBM2e of memory and NVLink 3.0 (200 GB/s). The learning rate and other hyperparameters were optimized through Optuna to maximize the model’s performance on the validation set. The training process required about 274 seconds. Moreover, we have also verified that the model, besides on the data center NVIDIA Ampere A100 Tensor Core GPU equipped with 64 GB of RAM, can be trained on consumer grade GPUs, by training it on a NVIDIA RTX 4090 GPU with 24 GB of RAM in about 201 seconds. The NVIDIA RTX 4090 is a more recent GPU than the

Ampere A100, it is cheaper and slightly faster. Figure 3 shows the training and the validation loss.

## V. EXPERIMENTAL RESULTS

Here, we present and discuss the results of our study, including the determination of the optimal number of clusters, the visualization of the clustering results, the performance comparison across different models, and a detailed analysis of the temporal predictions.

### A. CLUSTER DETERMINATION AND VISUALIZATION

The Elbow method was employed to determine the optimal number of clusters for *k*-means. Figure 4a shows the elbow plot, which indicates that four clusters are optimal, as the point where adding more clusters does not significantly improve the variance explained by the model. The clustered

<sup>4</sup><https://leonardo-supercomputer.cineca.eu>

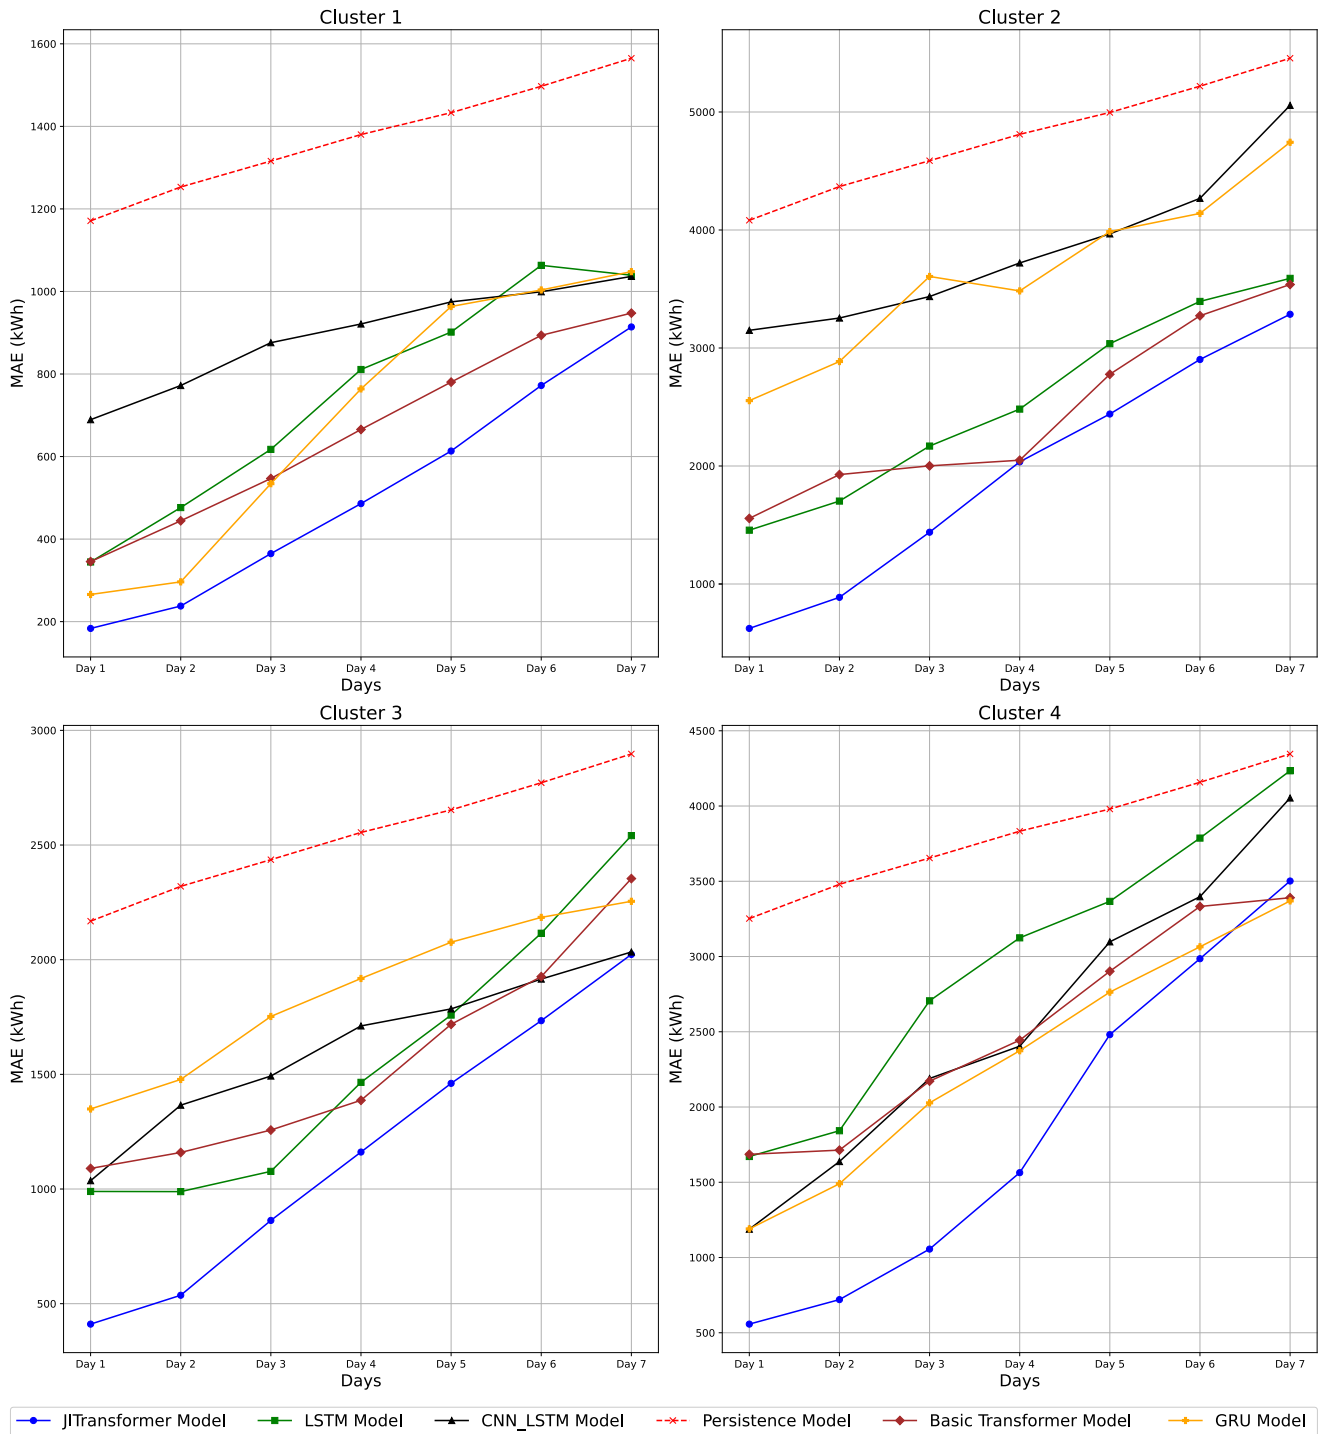



FIGURE 5. Comparison of average error per day across clusters.

data visualization shows distinct groupings of customers with similar consumption patterns. This clustering simplifies the modelling process and enhances the efficiency of the forecasts. Figure 4b clearly depicts the average consumption patterns for each cluster, demonstrating the variance in energy usage behaviors. For instance, we can observe that

the average consumption in Cluster 4 significantly differs from that in Cluster 3, indicating distinct energy usage behaviors between the two groups. This differentiation is also apparent across all clusters, each exhibiting its unique average consumption profile. In turn, the high variability discovered among the clusters directly influences the model



**FIGURE 6.** Comparison between the prediction and the ground truth varying the lead time.

design. The clustering allowed us to tailor the forecasting model to handle different consumption behaviors more effectively. This design choice ensures the model's robustness and adaptability to the diverse energy consumption patterns.

### B. EVALUATION METHODOLOGY

The predictive accuracy of the models was quantified using the average error metric. This metric, specifically the Mean Absolute Error (MAE), is computed as follows:

$$\text{MAE} = \frac{1}{n} \sum_{i=1}^n |P_i - A_i|, \quad (13)$$

where:

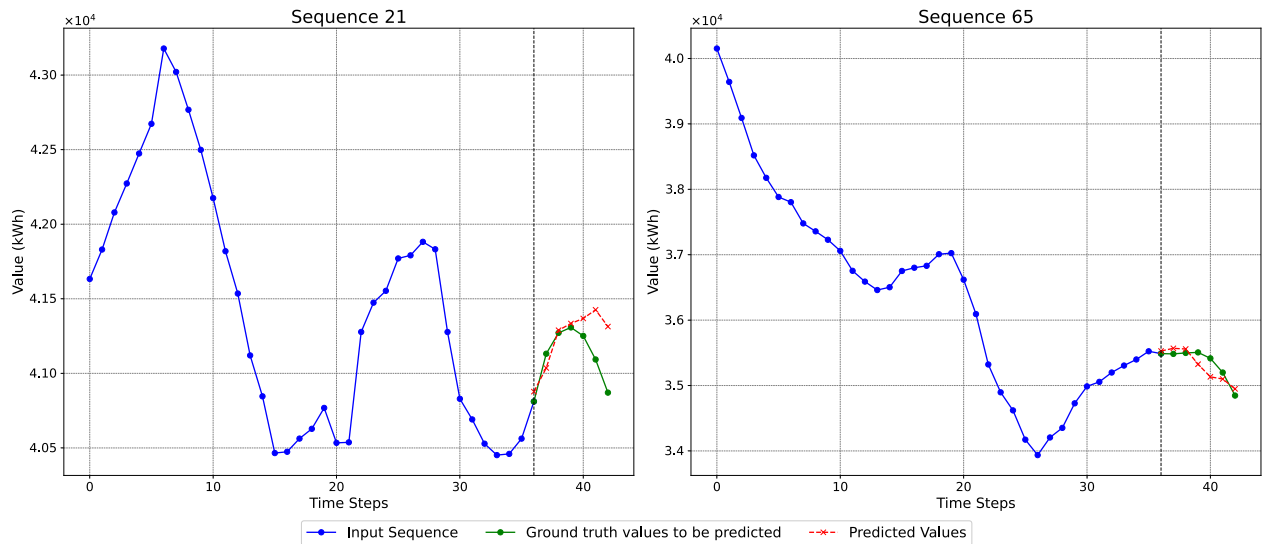
- $n$  is the number of data points;
- $P_i$  represents the predicted values of electricity consumption;
- $A_i$  represents the actual values of electricity consumption.

The absolute differences between predicted and actual values are summed up and then divided by the total number of observations, providing a straightforward measure of model accuracy in kWh.

### C. PERFORMANCE COMPARISON ACROSS CLUSTERS

The effectiveness of this model is compared against a basic transformer, a GRU model, a CNN-LSTM model, a LSTM model, and the persistence model<sup>5</sup> across the four clusters. Figure 5 provides a comparison of average errors per day for each cluster. The performance analysis clearly shows the JITrans model's superior ability to handle diverse energy usage patterns.

<sup>5</sup>The persistence model is a standard benchmark model, that works assuming that the future value of a time series is computed under the assumption that nothing changes between the current time and the forecast time.



**FIGURE 7.** Comparison between the prediction and the ground truth.

- Cluster 1: this cluster exhibits the least variability in errors, suggesting that households within this cluster have more predictable energy consumption patterns, due to similar lifestyle or appliance usage;
- Clusters 2 and 3: these clusters showed higher variability and errors, reflecting the challenges posed by more dynamic and diverse energy usage behaviours. The JITrans model was particularly effective in these clusters, reducing prediction errors and demonstrating its capacity to adapt to complex scenarios;
- Cluster 4: this cluster, which is the most heterogeneous in terms of energy usage, exhibits the highest errors across all models, emphasizing the challenges in forecasting for highly variable consumer behaviours. Despite this difficulty, our JITrans model outperformed other techniques, showcasing its robustness.

#### D. TEMPORAL ANALYSIS OF PREDICTIONS

Figures 6 and 7 depict the true vs. predicted values for some sequences and days across different clusters, highlighting the JITrans model adeptness at closely mirroring actual consumption patterns and minimizing discrepancies. Particularly on days 1, 2 and 3, the model exhibits high accuracy, which is crucial for effective load management. A slight underestimation can be noted on day 7 and only for cluster 4. In this case the GRU and the basic transformer model perform just slightly better. However, it is worth noting here that (i) the observed behaviour is expected, i.e. it is normal for the prediction accuracy to decrease as the forecasting horizon increases, and (ii) the user (in our case the stakeholder is the power company) is not really interested in accurately predicting the load on such a distant forecasting horizon: the main focus is on predicting the load for the next day.

Regarding specific sequence predictions, the examination of some randomly selected sequences - Sequence 21 and

Sequence 65 depicted in Figure 7 - showed high initial accuracy with slight deviations towards the end. Again, this is the expected behaviour. We observed stable pattern prediction with minor discrepancies, reinforcing the model's robustness and reliability.

This portion of our analysis focuses on selected days as representative examples (owing to the impossibility of presenting detailed daily results for all four clusters and the fact that those results are essentially similar). The observed patterns, consistent across different clusters, reinforce the robustness of our modelling approach and confirm the JITrans model's capability to adapt to and accurately forecast diverse energy usage behaviours.

#### E. DISCUSSION

The JITrans model demonstrated a significant improvement over traditional forecasting models such as LSTM, CNN\_LSTM, and GRU in predicting household electricity consumption. This enhanced performance is primarily due to the transformer's sophisticated capacity to simultaneously process and analyze complex temporal dependencies and patterns within clustered data. Integrating customers with similar consumption behaviours into clusters not only streamlined the forecasting process but also significantly enhanced overall efficiency. This approach eliminated the need for creating separate models for each household, thus addressing the impracticality of such an endeavour in real-world applications. The JITrans model greatly improves the accuracy in forecasting, allowing utility companies to optimize energy distribution strategies effectively. By precisely predicting consumption patterns, these utilities can manage demand fluctuations more efficiently, seamlessly integrating renewable energy sources, and optimizing their energy supply chains. Such advancements in predictive capabilities are crucial for adapting to the changing dynamics of global

energy needs, promoting sustainable practices, and reducing wastage in energy distribution, marking a significant step forward in the development of smart energy management solutions.

## VI. CONCLUSION

In this paper we introduced a novel approach to energy load forecasting, based on our JITrans model specifically designed for accurate energy consumption forecasting. We enhance traditional transformers by implementing specific training for each day in the prediction horizon, where a dedicated model averages the prediction errors of previous models to predict each day. This design allows for more effective and accurate time series forecasting by reducing cumulative errors and leveraging information from previous predictions.

Our model demonstrated superior performance in various numerical benchmarks, outperforming state-of-the-art models in accuracy and robustness. We attribute this improvement to the model's ability to explore and utilize information within each token from previous predictions, effectively mitigating prediction errors over time.

From a theoretical perspective, future research will delve into refining this model by investigating novel error reduction strategies and optimization techniques. From a methodological standpoint, we aim to extend the application of our model to additional time series forecasting domains, testing its adaptability and effectiveness. From a practical viewpoint, we envisage the integration of external variables such as climatic conditions and economic data to further refine the forecasting capabilities of our model. Planned future developments include:

- **Real-Time Prediction:** Implement the enhanced model in a real-time energy management system to continuously forecast energy consumption;
- **Feedback Loop:** Setup a feedback loop to update the model with new data, including changes in user behaviour, to keep the predictions accurate over time;
- **Regional datasets:** Assess the behaviour of our model on several datasets characterized by different geolocations.

## ACKNOWLEDGMENT

The authors are grateful to Prof. Massimo Tavoni and Prof. Jacopo Bonan (Politecnico di Milano, Milan, Italy) for providing the dataset used to assess their Just in Time Transformer model.

## REFERENCES

- [1] Z. Guo, J. R. O'Hanley, and S. Gibson, "Predicting residential electricity consumption patterns based on smart meter and household data: A case study from the republic of Ireland," *Utilities Policy*, vol. 79, Dec. 2022, Art. no. 101446.
- [2] G. E. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung, *Time Series Analysis: Forecasting and Control*. Hoboken, NJ, USA: Wiley, 2015.
- [3] C. Cortes and V. Vapnik, "Support-vector networks," *Mach. Learn.*, vol. 20, no. 3, pp. 273–297, Sep. 1995.
- [4] L. Breiman, "Random forests," *Mach. Learn.*, vol. 45, pp. 5–32, Oct. 2001.
- [5] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning representations by back-propagating errors," *Nature*, vol. 323, no. 6088, pp. 533–536, Oct. 1986.
- [6] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997.
- [7] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, Nov. 1998.
- [8] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 30, 2017.
- [9] M. Jaramillo, W. Pavón, and L. Jaramillo, "Adaptive forecasting in energy consumption: A bibliometric analysis and review," *Data*, vol. 9, no. 1, p. 13, Jan. 2024.
- [10] R. Mathumitha, P. Rathika, and K. Manimala, "Intelligent deep learning techniques for energy consumption forecasting in smart buildings: A review," *Artif. Intell. Rev.*, vol. 57, no. 2, p. 35, Feb. 2024.
- [11] A. Barbaresi, M. Ceccarelli, G. Menichetti, D. Torreggiani, P. Tassinari, and M. Bovo, "Application of machine learning models for fast and accurate predictions of building energy need," *Energies*, vol. 15, no. 4, p. 1266, Feb. 2022.
- [12] M. AL-Ghamdi, A. A.-M. AL-Ghamdi, and M. Ragab, "A hybrid DNN multilayered LSTM model for energy consumption prediction," *Appl. Sci.*, vol. 13, no. 20, p. 11408, Oct. 2023.
- [13] A. Mishra, H. R. Lone, and A. Mishra, "DECODE: Data-driven energy consumption prediction leveraging historical data and environmental factors in buildings," *Energy Buildings*, vol. 307, Mar. 2024, Art. no. 113950.
- [14] H. Liu, J. Liang, Y. Liu, and H. Wu, "A review of data-driven building energy prediction," *Buildings*, vol. 13, no. 2, p. 532, Feb. 2023.
- [15] K. Gajowiczek and T. Ząbkowski, "Short term electricity forecasting using individual smart meter data," *Proc. Comput. Sci.*, vol. 35, pp. 589–597, Jan. 2014.
- [16] D. Ilić, P. G. da Silva, S. Karnouskos, and M. Jacobi, "Impact assessment of smart meter grouping on the accuracy of forecasting algorithms," in *Proc. 28th Annu. ACM Symp. Appl. Comput.*, Mar. 2013, pp. 673–679.
- [17] K. Zhou and S. Yang, "Understanding household energy consumption behavior: The contribution of energy big data analytics," *Renew. Sustain. Energy Rev.*, vol. 56, pp. 810–819, Apr. 2016.
- [18] T. A. Nakabi and P. Toivanen, "An ANN-based model for learning individual customer behavior in response to electricity prices," *Sustain. Energy, Grids Netw.*, vol. 18, Jun. 2019, Art. no. 100212.
- [19] H. Wang, N. K. Mahato, H. He, X. An, Z. Chen, and G. Gong, "Research on electricity consumption behavior of users based on deep learning," in *Proc. IEEE/IAS Ind. Commercial Power Syst. Asia (I&CPS Asia)*, Jul. 2020, pp. 1491–1497.
- [20] H. Brugger, W. Eichhammer, N. Mikova, and E. Dönitz, "Energy efficiency vision 2050: How will new societal trends influence future energy demand in the European countries?" *Energy Policy*, vol. 152, May 2021, Art. no. 112216.
- [21] A. Shahzadeh, A. Khosravi, and S. Nahavandi, "Improving load forecast accuracy by clustering consumers using smart meter data," in *Proc. Int. Joint Conf. Neural Netw. (IJCNN)*, Jul. 2015, pp. 1–7.
- [22] B. Yildiz, J. I. Bilbao, J. Dore, and A. Sproul, "Household electricity load forecasting using historical smart meter data with clustering and classification techniques," in *Proc. IEEE Innov. Smart Grid Technol. Asia (ISGT Asia)*, May 2018, pp. 873–879.
- [23] L. Wen, K. Zhou, and S. Yang, "A shape-based clustering method for pattern recognition of residential electricity consumption," *J. Cleaner Prod.*, vol. 212, pp. 475–488, Mar. 2019. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0959652618337697>
- [24] L. Czétányi, V. Vámos, M. Horváth, Z. Szalay, A. Mota-Babiloni, Z. Deme-Bélafi, and T. Csoknyai, "Development of electricity consumption profiles of residential buildings based on smart meter data clustering," *Energy Buildings*, vol. 252, Dec. 2021, Art. no. 111376. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0378778821006605>
- [25] V. Michalakopoulos, E. Sarmas, I. Papias, P. Skaloumpakas, V. Marinakis, and H. Doukas, "A machine learning-based framework for clustering residential electricity load profiles to enhance demand response programs," *Appl. Energy*, vol. 361, May 2024, Art. no. 122943.
- [26] J. Yang, J. Zhao, F. Wen, and Z. Dong, "A model of customizing electricity retail prices based on load profile clustering analysis," *IEEE Trans. Smart Grid*, vol. 10, no. 3, pp. 3374–3386, May 2019.

- [27] A. A. E. Benali, M. Cafaro, I. Epicoco, M. Pulimeno, E. J. Schioppa, J. Bonan, and M. Tavoni, "Smart meters and customer consumption behavior: An exploratory analysis approach," in *Proc. Int. Conf. Extended Reality*. Cham, Switzerland: Springer, 2023, pp. 356–363.
- [28] K. Li, C. Hu, G. Liu, and W. Xue, "Building's electricity consumption prediction using optimized artificial neural networks and principal component analysis," *Energy Buildings*, vol. 108, pp. 106–113, Dec. 2015.
- [29] J.-S. Chou and D.-S. Tran, "Forecasting energy consumption time series using machine learning techniques based on usage patterns of residential householders," *Energy*, vol. 165, pp. 709–726, Dec. 2018.
- [30] L. Sehovac, C. Nesen, and K. Grolinger, "Forecasting building energy consumption with deep learning: A sequence to sequence approach," in *Proc. IEEE Int. Congr. Internet Things (ICIOT)*, Jul. 2019, pp. 108–116.
- [31] A. Lotfipoor, S. Patidar, and D. P. Jenkins, "Short-term forecasting of residential electricity demand using cnn-lstm," in *Proc. 2nd IBPSA-Scotland uSIM Conf., Building Buildings–Urban Community Energy Modeling*, 2020, pp. 118–125.
- [32] A. V. Abraham, P. Sasidharan, S. J. S. Tejas, M. Manohara, R. Muthu, and R. C. Naidu, "Predicting energy consumption using LSTM and CNN deep learning algorithm," in *Proc. 7th Int. Conf. Environ. Friendly Energies Appl. (EFEA)*, Dec. 2022, pp. 1–6.
- [33] S. Mahjoub, L. Chrifi-Alaoui, B. Marhic, and L. Delahoche, "Predicting energy consumption using LSTM, multi-layer GRU and drop-GRU neural networks," *Sensors*, vol. 22, no. 11, p. 4062, May 2022.
- [34] N. Wu, B. Green, X. Ben, and S. O'Banion, "Deep transformer models for time series forecasting: The influenza prevalence case," 2020, *arXiv:2001.08317*.
- [35] B. Lim, S. Ö. Arık, N. Loeff, and T. Pfister, "Temporal fusion transformers for interpretable multi-horizon time series forecasting," *Int. J. Forecasting*, vol. 37, no. 4, pp. 1748–1764, Oct. 2021.
- [36] H. Zhou, S. Zhang, J. Peng, S. Zhang, J. Li, H. Xiong, and W. Zhang, "Informer: Beyond efficient transformer for long sequence time-series forecasting," in *Proc. AAAI Conf. Artif. Intell.*, vol. 35, no. 12, May 2021, pp. 11106–11115.
- [37] M. N. Fekri, H. Patel, K. Grolinger, and V. Sharma, "Deep learning for load forecasting with smart meter data: Online adaptive recurrent neural network," *Appl. Energy*, vol. 282, Jan. 2021, Art. no. 116177.
- [38] M. Khashei, F. Chahkoutahi, and N. Bakhtiarvand, "A novel discrete deep learning-based intelligent methodology for energy consumption classification," *Energy Rep.*, vol. 9, pp. 4861–4871, Dec. 2023.
- [39] X. Wang, T. Zhou, Q. Wen, J. Gao, B. Ding, and R. Jin, "CARD: Channel aligned robust blend transformer for time series forecasting," in *Proc. 12th Int. Conf. Learn. Represent.*, 2024.
- [40] P. Chen, Y. Zhang, Y. Cheng, Y. Shu, Y. Wang, Q. Wen, B. Yang, and C. Guo, "Pathformer: Multi-scale transformers with adaptive pathways for time series forecasting," 2024, *arXiv:2402.05956*.
- [41] S. Ben Taieb, R. Huser, R. J. Hyndman, and M. G. Genton, "Forecasting uncertainty in electricity smart meter data by boosting additive quantile regression," *IEEE Trans. Smart Grid*, vol. 7, no. 5, pp. 2448–2455, Sep. 2016.
- [42] Z. Wang, Y. Wang, R. Zeng, R. S. Srinivasan, and S. Ahrentzen, "Random forest based hourly building energy prediction," *Energy Buildings*, vol. 171, pp. 11–25, Jul. 2018.
- [43] Y. Wang, N. Zhang, Y. Tan, T. Hong, D. S. Kirschen, and C. Kang, "Combining probabilistic load forecasts," *IEEE Trans. Smart Grid*, vol. 10, no. 4, pp. 3664–3674, Jul. 2019.
- [44] D. Weeraddana, N. L. D. Khoa, L. O'Neil, W. Wang, and C. Cai, "Energy consumption forecasting using a stacked nonparametric Bayesian approach," in *Proc. Joint Eur. Conf. Mach. Learn. Knowl. Discovery Databases*, Ghent, Belgium. Cham, Switzerland: Springer, Sep. 2020 pp. 19–35.
- [45] J. Xu, X. Kang, Z. Chen, D. Yan, S. Guo, Y. Jin, T. Hao, and R. Jia, "Clustering-based probability distribution model for monthly residential building electricity consumption analysis," in *Building Simulation*, vol. 14. Cham, Switzerland: Springer, 2021, pp. 149–164.
- [46] C. Tian, C. Li, G. Zhang, and Y. Lv, "Data driven parallel prediction of building energy consumption using generative adversarial nets," *Energy Buildings*, vol. 186, pp. 230–243, Mar. 2019.
- [47] M. Alhoussein, K. Aouragzeb, and S. I. Haider, "Hybrid CNN-LSTM model for short-term individual household load forecasting," *IEEE Access*, vol. 8, pp. 180544–180557, 2020.
- [48] H. Bian, Y. Zhong, J. Sun, and F. Shi, "Study on power consumption load forecast based on K-means clustering and FCM–BP model," *Energy Rep.*, vol. 6, pp. 693–700, Dec. 2020.
- [49] A.-N. Khan, N. Iqbal, A. Rizwan, R. Ahmad, and D.-H. Kim, "An ensemble energy consumption forecasting model based on spatial-temporal clustering analysis in residential buildings," *Energies*, vol. 14, no. 11, p. 3020, May 2021.
- [50] S. Dai and F. Meng, "Online model-based functional clustering and functional deep learning for load forecasting using smart meter data," in *Proc. Int. Conf. Smart Energy Syst. Technol. (SEST)*, Sep. 2022, pp. 1–6.
- [51] L. Saad Saoud, H. Al-Marzouqi, and R. Hussein, "Household energy consumption prediction using the stationary wavelet transform and transformers," *IEEE Access*, vol. 10, pp. 5171–5183, 2022.
- [52] M. Kalksma, B. Setz, A. R. Pratama, I. Georgievski, and M. Aiello, "Mining sequential patterns for appliance usage prediction," in *Proc. 7th Int. Conf. Smart Cities Green ICT Syst.*, 2018, pp. 23–33.
- [53] S.-V. Oprea and A. Băra, "Machine learning algorithms for short-term load forecast in residential buildings using smart meters, sensors and big data solutions," *IEEE Access*, vol. 7, pp. 177874–177889, 2019.
- [54] F. A. Alaba, U. Sani, E. G. Dada, and B. H. Mohammed, "AIoT-enabled smart grids: Advancing energy efficiency and renewable energy integration," in *Artificial Intelligence of Things for Achieving Sustainable Development Goals*. Cham, Switzerland: Springer, 2024, pp. 59–79.
- [55] A. K. Jain and R. C. Dubes, *Algorithms for Clustering Data*. Upper Saddle River, NJ, USA: Prentice-Hall, 1988.
- [56] J. A. Hartigan and M. A. Wong, "Algorithm as 136: A k-means clustering algorithm," *J. Roy. Stat. Soc. Ser. C Appl. Statist.*, vol. 28, no. 1, pp. 100–108, 1979.
- [57] M. J. Zaki and W. Meira, *Data Mining and Analysis: Fundamental Concepts and Algorithms*. Cambridge, U.K.: Cambridge Univ. Press, 2014.
- [58] P. Shaw, J. Uszkoreit, and A. Vaswani, "Self-attention with relative position representations," 2018, *arXiv:1803.02155*.
- [59] G. Ke, D. He, and T.-Y. Liu, "Rethinking positional encoding in language pre-training," 2020, *arXiv:2006.15595*.
- [60] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," 2018, *arXiv:1810.04805*.
- [61] L. Dong, S. Xu, and B. Xu, "Speech-transformer: A no-recurrence sequence-to-sequence model for speech recognition," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Apr. 2018, pp. 5884–5888.
- [62] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, "Swin transformer: Hierarchical vision transformer using shifted windows," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 9992–10002.
- [63] Q. Wen, T. Zhou, C. Zhang, W. Chen, Z. Ma, J. Yan, and L. Sun, "Transformers in time series: A survey," 2022, *arXiv:2202.07125*.
- [64] J. Xiao, R. Liu, and E. L. Dyer, "GAFoformer: Enhancing timeseries transformers through group-aware embeddings," in *Proc. 12th Int. Conf. Learn. Represent.*, 2024. [Online]. Available: <https://openreview.net/forum?id=c56TWtYp0W>
- [65] Y. Li, W. Chen, X. Hu, B. Chen, baolin sun, and M. Zhou, "Transformer-modulated diffusion models for probabilistic multivariate time series forecasting," in *Proc. 12th Int. Conf. Learn. Represent.*, 2024. [Online]. Available: <https://openreview.net/forum?id=qae04YACHs>
- [66] Y. Liu, T. Hu, H. Zhang, H. Wu, S. Wang, L. Ma, and M. Long, "ITransformer: Inverted transformers are effective for time series forecasting," 2023, *arXiv:2310.06625*.



**AHMED ALA EDDINE BENALI** (Graduate Student Member, IEEE) received the bachelor's degree in electrical engineering and the master's degree in electrical control from the Department of Electrical Engineering, Ibn Khaldoun University of Tiaret, Algeria, in 2018 and 2020, respectively. He is currently pursuing the Ph.D. degree with the Department of Engineering for Innovation, University of Salento, Italy, affiliated with the HPC Laboratory. He is a Ph.D. Intern with Inmatica S.p.A., an Italian ICT company. His research interests include the development of advanced machine learning and deep learning algorithms, with innovative applications in the energy sector.





**MASSIMO CAFARO** (Senior Member, IEEE) received the Laurea degree (M.Sc.) in computer science from the University of Salerno, and the Ph.D. degree in computer science from the University of Bari. He is currently an Associate Professor with the Department of Engineering for Innovation, University of Salento. He is the Director of the Master in Applied Data Science and the Head of the HPC Laboratory, University of Salento. He is the author of more than 110 refereed

articles and holds a patent on distributed database technologies. He focuses his research on high performance and distributed computing on both theoretical and practical aspects, in particular the design and analysis of sequential, and parallel and distributed algorithms. His research interests include parallel and distributed computing, cloud and grid computing, data mining, machine learning, deep learning, big data, security, and cryptography. He is a Senior Member of the IEEE Computer Society and ACM and the Vice Chair of the Regional Centers and Coordinator of the Technical Area on Data Intensive Computing for the IEEE Technical Committee on Scalable Computing. He serves as an Associate Editor for IEEE ACCESS and *Future Internet* (MDPI) and a moderator for the IEEE TechRxiv preprint repository (“Computing and Processing” category).



**ITALO EPICOCO** received the Ph.D. degree in innovative materials and technologies from ISUFI, University of Lecce, Italy, in 2003. He is currently an Assistant Professor with the University of Salento, Lecce, and the Director of the Advanced Scientific Computing (ASC) Division, Euro-Mediterranean Center on Climate Change Foundation (CMCC). He is the Director of the Master in Scientific Programming and the Co-Leader of the HPC Laboratory, University of

Salento. His main skills concern computer engineering and computer science. He has published more than 80 papers in refereed books, journal, and conference proceedings on parallel and grid computing. His research interests include the design of data mining algorithms on high-end computing architectures, high performance, and distributed computing. He is also involved in the optimization of numerical kernels for solving PDEs. The current application field is the Earth system models with particular focus on ocean models. His research activity also includes the benchmarking and evaluation of new emerging computational technologies based on GP-GPU, hybrid architectures, and FPGA-based computing.



**MARCO PULIMENO** received the Ph.D. degree in mathematics and computer science from the Department of Engineering for Innovation, University of Salento. He is currently a Researcher with the Department of Engineering for Innovation, University of Salento. He has published more than 40 refereed papers in journals and conference proceedings. His research interests include streaming algorithms, distributed and high performance computing, data mining, and machine learning.



**ENRICO JUNIOR SCHIOPPA** received the Ph.D. degree in physics from the University of Amsterdam, The Netherlands. He is a Nuclear and Particle Physicist and a Former Fellow with European Center for Nuclear Research (CERN), Geneva, Switzerland. He is currently a Project Manager of Inmatica S.p.A., an Italian ICT company, where he carries out research and development activities in the fields of the IoT and artificial intelligence.

In the past, he contributed to the development of AI models for the analysis of the data produced with the Large Hadron Collider, CERN. In Inmatica S.p.A., he contributes to the development of computer vision techniques for the analysis of sport events, of AI models applied to brain-machine interfaces, and of the IoT solutions based on innovative fiber optic sensors and piezoelectric wearable devices. He is the author of over 60 peer-reviewed articles in addition to more than 200 publications as the author of the ATLAS experiment. He is the author of one patent on an innovative system utilizing compact neutron accelerators and AI to look for harmful substances in parcels.

...

Open Access funding provided by ‘Università del Salento’ within the CRUI CARE Agreement