



HAL
open science

Time Series Analytics for Electricity Consumption Data

Adrien Petralia

► **To cite this version:**

Adrien Petralia. Time Series Analytics for Electricity Consumption Data. Proceedings of the VLDB Endowment (PVLDB), 2024. hal-04706310

HAL Id: hal-04706310

<https://hal.science/hal-04706310v1>

Submitted on 23 Sep 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Time Series Analytics for Electricity Consumption Data

Adrien Petralia

Supervised by: Themis Palpanas, Philippe Charpentier
EDF R&D - Université Paris Cité, LIPAPE
adrien.petralia@gmail.com

ABSTRACT

Rationalizing and better-managing energy consumption have become primary objectives in the global effort to prevent climate change. In this context, enhancing our understanding of electricity consumption behavior is crucial. Thus, electricity suppliers have installed millions of smart meters worldwide over the past decade, capturing time-stamped electricity consumption data of the total main power consumed in individual households. Nevertheless, suppliers face significant challenges in extracting detailed information from these aggregated signals, such as identifying which appliances the customer owns and their typical usage. This task is complicated by the reliance on low-frequency smart meter readings, which combine signals from various appliances operating simultaneously. Moreover, the scarcity of annotations and the large amount of long, variable-length consumption series collected further complicate data analysis and interpretation. In this Ph.D. work, we propose a set of new solutions to tackle the appliance detection problem and extract detailed information from smart meter data that overcome the challenges listed above. First, we propose to tackle this task as a binary time series classification (TSC) problem and subsequently describe the Appliance Detection Framework (ADF), designed to enhance classifiers' performance using long and variable consumption series. Moreover, we introduce TransApp, a deep-learning architecture that is first pretrained in a self-supervised way to enhance its performance on appliance detection tasks. Finally, we propose an interactive system based on a combination of TSC and explainable classification that enables the localization of appliance patterns without the use of strong labels.

VLDB Workshop Reference Format:

Adrien Petralia. Time Series Analytics for Electricity Consumption Data. VLDB 2024 Workshop: VLDB Ph.D. Workshop.

1 INTRODUCTION

With the urge to fight against climate change, energy saving is emerging as a major lever. One way to achieve this goal is to help consumers better understand their consumption so that they can play an active role in the energy transition. In the last decade, millions of smart meters have been installed across the globe by electricity providers [1], capturing detailed timestamped data of the total electricity consumed in individual households (called load

curve or electricity consumption time series). These data are supposed to allow electricity suppliers help individual customers understand their consumption. However, suppliers face significant challenges in extracting detailed information from them, such as identifying which appliances the customer owns and their typical usage. This information is valuable since it can help customers save by identifying over-consuming devices. To overcome these challenges, we need to directly detect appliance use from the recorded smart meter data. However, the *very low-frequency* readings commonly used worldwide (e.g., 30min in France and the UK, and 60min in Spain [13]) aggregate the power consumption of multiple appliances running simultaneously, smoothing out appliance patterns and exacerbating the analysis of these data [8]. In addition, the large number of long and variable-length consumption series collected by suppliers, coupled with the scarcity of annotations available, further complicated the task of training accurate solutions.

Detecting appliances [8] is related to Non-Intrusive Load Monitoring (NILM), a well-known and growing research field that aims to identify the power consumption, pattern, or on/off state activation of individual appliances using only the total recorded consumption signal [10]. Many solutions have been proposed in the literature, such as signature-based methods that use information related to the unique patterns of specific appliances [5]. Most of these studies rely on smart meters capable of recording data at frequencies 1Hz, or higher. However, since these devices use the smart grid systems for data transmission, the majority of data currently available to suppliers are at much lower frequencies. In addition, these solutions must be trained using individual appliance power, i.e., knowing each appliance's exact state of activation and consumption power for each timestamp. Unfortunately, gathering such data is expensive, as each appliance needs to be monitored with sensors to measure its individual consumption. Most of the data available from suppliers for training their solutions are based on customer surveys that only indicate whether or not they own a particular appliance. Another way to tackle this challenge is to cast the appliance detection problem as a binary classification task, where a time series classifier is trained to detect an appliance in a consumption series. Time series classification represents a growing field of interest, and many algorithms have been proposed in the literature [6]. However, due to the lack of publicly available labeled datasets, no method from this family has been applied to the specific task of appliance detection with very low-frequency consumption series. In addition, even though these methods can be used for detecting *if* an appliance has been used, they cannot determine *when* it was used.

To address the aforementioned problems, we propose the following research directions:

- Appliance Detection Framework (ADF) [9], a framework that enhances the performance of time series classifiers on

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment. Proceedings of the VLDB Endowment. ISSN 2150-8097.

the appliance detection task using long and variable length consumption series.

- TransApp [9], a deep learning time series classifier that is first pretrained in a self-supervised manner to enhance its performance on appliance detection tasks.
- DeviceScope [7], an interactive application that enables the detection and localization of individual appliance patterns within a given period using explainable time series classification approaches.

2 BACKGROUND

An **electricity consumption time series** is defined as a univariate time series $X = (x_1, \dots, x_T)$ of ordered elements $x_j \in \mathcal{R}_+^1$ following (i_1, \dots, i_T) time consumption indexes (i.e., timestamps). Subsequently, the sampling frequency is defined as the time difference between two records index $\Delta_t := i_j - i_{j-1}$. Each element x_j , usually given in Watt, indicates either the actual power at time i_j or the average electric power called during the interval time Δ_t . In our work, we refer to *very low frequency* consumption series for data sampled at more than 1min.

Detecting if an appliance has been used in a period of time can be cast as a time series classification (TSC) problem [8] where a classifier is trained in a binary supervised manner to detect the presence of an appliance using only one label (0 or 1) for an entire series. The most prominent approaches to solve TSC include random convolution-based (e.g., ROCKET, Minirocket, Arsenal), deep-learning-based (ConvNet, ResNet, InceptionTime), dictionary-based, and interval-based algorithms [6]. For the rest of the paper, we define formally the Appliance Detection Problem as follows:

Definition 2.1 (Appliance Detection Problem). Given an aggregate smart meter time series X and an appliance type a , we want to know if appliance a was activated at least once in X (i.e., was in an "ON" state, regardless of the time and number of activations).

3 OUR WORK: ADF & TRANSAPP

We first conducted an extensive performance comparison of time series classifiers on the appliance detection task [8]. The results showed that deep-learning methods (convolutional-based) performed best, but novel solutions were needed to achieve the desired accuracy. Additionally, our work pointed to the need to operate on variable-length time series, which reflects the complexity of dealing with real-world client consumption data. Building on the insights of these results, we developed the Appliance Detection Framework (ADF) and TransApp [9]. ADF is our proposed framework, designed to enhance the performance of time series classifiers in detecting the presence/absence of appliances in *long* and *variable-length* consumption time series. To do so, ADF takes as input *subsequences* of an entire consumption series and outputs one prediction. TransApp is our novel deep-learning classifier. It is first trained in a self-supervised manner using non-labeled consumption time series and then fine-tuned on labeled data to detect a specific appliance. We propose this architecture to take advantage of the large amount of non-labeled data currently available to suppliers.

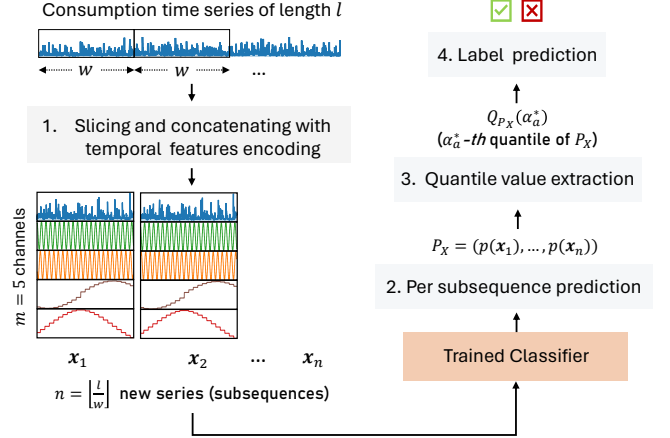


Figure 1: The Appliance Detection Framework (ADF).

3.1 ADF

The proposed Appliance Detection Framework is illustrated in Figure 1. ADF takes fragments of an entire consumption time series as input. Subsequently, the classifier used inside the framework is trained using subsequences. We note that this general framework can be used with any classifier able to predict probabilities. The proposed framework uses the following steps to detect the presence of an appliance a in a consumption time series X , as follows.

Step 1. A consumption time series X of length l is first sliced in n new non-overlapping subsequences of length w , using a tumbling window. In addition, to keep positional information about the time of the days and hours, we concatenate the sliced subsequences with temporal encoded features (see Section 3.1.a for details). This results in $n = \lfloor \frac{l}{w} \rfloor$ new multivariate time series $x^{w \times m}$. With m , the number of channels.

Step 2. Afterward, we feed each subsequence to a classifier instance previously trained to detect the specific appliance a . The model then predicts a given detection probability for each subsequence x_i , resulting in a vector of probability $P_X = (p(x_1), \dots, p(x_n))$.

Step 3. The value of the α_a^* -th quantile is then extracted from P_X . This value is determined during training to maximize an accuracy measure on a validation dataset. Unlike simple majority voting, this approach accounts for the model's confidence in each subsequence.

Step 4. At the end, the final predicted label is given by rounding the extracted value.

a) Temporal Features Encoding. We introduce additional channels to encode time to improve the model's understanding of time-related patterns and enhance detection. More precisely, we add new channels as encoded features related to the days and hours by projecting these discrete features on a sinusoidal basis as $Te_{sin}(i_t) = \sin\left(\frac{2\pi i_t}{p}\right)$ and $Te_{cos}(i_t) = \cos\left(\frac{2\pi i_t}{p}\right)$ with $i_t = \{1, \dots, 24\}$ and $p = 24$ for hour encoding, and $i_t = \{1, \dots, 7\}$ and $p = 7$ for days encoding.

3.2 TransApp

TransApp is a deep-learning architecture specifically designed to be used inside the ADF. As shown in Figure 2(a), TransApp results in

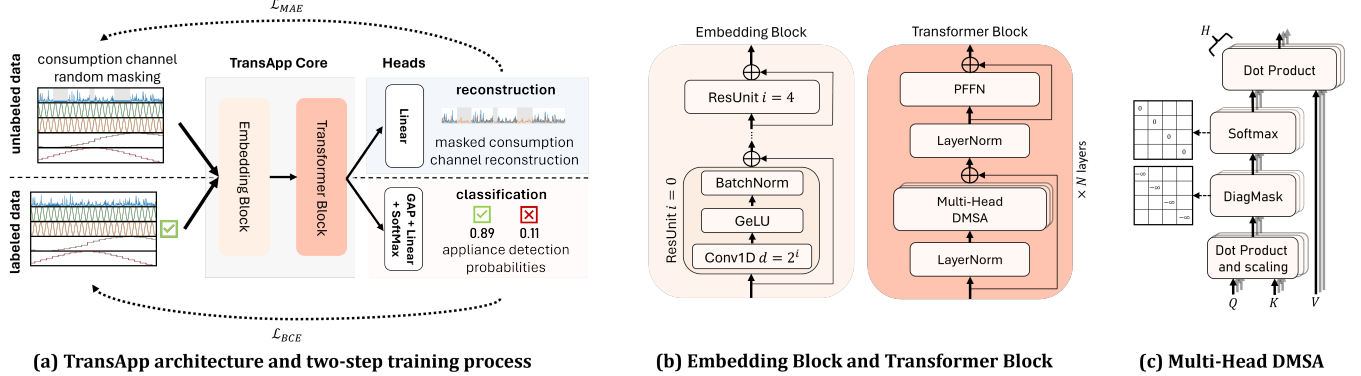


Figure 2: Overview of the TransApp architecture and the two-step training process.

an encoder that can be combined with a specific head to reconstruct a corrupted series during the pretraining or perform classification.

The core of TransApp combines a robust embedding block based on dilated convolutions followed by a Transformer block. The Embedding Block, shown in Figure 2(b), is composed of 4 convolutional Residual Units (ResUnit) that use a dilation parameter $d = 2^i$ (with $i = 1, \dots, 4$) that exponentially increases according to the ResUnit’s depth. This block serves as a features extractor, providing localized patterns to help the model perform better on classification tasks. The Transformer block results in N stacked Transformer layers, depicted in Figure 2(b). It is used to learn long-range dependencies and is a key part of our architecture to extract representation and benefit of our pretraining process. We note that we introduce DMSA instead of the original Self-Attention mechanism [11] in the Transformer layer as a strategy that emphasizes inter-token relations and enhances the model’s ability to capture meaningful dependencies. As depicted in Figure 2(c), DMSA involves applying a mask to the diagonal elements of the attention score matrix that forces the scores to be zero after the softmax operation.

3.2.1 Two-step training process. We introduce a two-step training process for the proposed TransApp architecture (cf. Figure 2(a)).

[Self-supervised Pretraining] Leveraging self-supervised pretraining of a Transformer architecture on auxiliary tasks has previously been employed to enhance model performance on downstream tasks [2, 3]. The proposed pretraining process, inspired by the mask-based pretraining of vision transformers [3], involves utilizing only the input consumption series without any label information. As depicted in Figure 2(a), this step results in a reconstruction objective of a corrupted (masked) time series fed to the model input using a linear layer after the core model architecture. We note that the masking process used in our approach aims to corrupt random segments (i.e., values set to 0) into the consumption series channel of the input sequence.

During the self-supervised process, the model is trained using a Loss function that calculates the Mean Absolute Error between the predicted and true values of the masked elements of an input consumption time series (\mathcal{L}_{MAE}), defined as $\mathcal{L}_{MAE} = \frac{1}{\#M} \sum_i |\hat{x}_i - x_i|_1$, with $\#M$ the number of masked elements in the input series. **[Supervised Training]** As depicted in Figure 2(a), the supervised training results in a simple binary classification process using labeled time series. For this step, we use a classification head after

the core model architecture that comprises a global average pooling, followed by a linear layer and a softmax activation. During this phase, TransApp is trained using the Binary-Cross Entropy Loss (\mathcal{L}_{BCE}). Note that the label of the entire consumption series is assigned to all sliced subsequences during the training process.

4 EXPERIMENTAL ANALYSIS

We now present our experimental evaluation of ADF & TransApp. First, we conducted a comprehensive evaluation to compare the state-of-the-art time series classifiers [6], aiming to identify the most accurate and scalable option for appliance detection, considering factors like meter reading and consumption series length [8]. Our study found that deep learning and random convolution-based classifiers are the most effective. Based on these findings, we compared the performance of our solutions (ADF+TransApp) against the best 9 different time series classifiers on 16 appliance detection cases using two labeled electricity consumption datasets. We designated the non-pretrained architecture as TransApp and the pre-trained one as TransAppPT (pretrained only on the labeled dataset). Additionally, we named TransAppPT-1, a larger architecture that was pretrained on a large dataset of 200k consumption series and then fine-tuned on the two labeled datasets. We report detection accuracy (using the F1-Macro Score) and execution time, noting that due to space constraints, only aggregated results are presented here. Figure 3(a) shows detection performance for different baselines (inside/outside ADF) compared to our solutions, indicating that all classifiers benefit from ADF and that our solutions are the most accurate. Figure 3(b) presents the average detection score for TransAppPT-1 with different pretraining data sizes, showing that the score increases proportionally with the data used. Figure 3(c) illustrates the inference time to predict 1K consumption series labels based on input sequence length, indicating that TransApp scales nearly linearly within our framework, while inference time significantly increases for long sequences outside the framework. These preliminary results are auspicious, as they show the superiority of the proposed approaches. Moreover, we demonstrate that pretraining TransApp on large consumption series data can significantly improve its performance on downstream tasks. We note that this solution can be applied to data series in different domains, and improve classification performance in situations with limited labeled data (e.g., in the health domain).

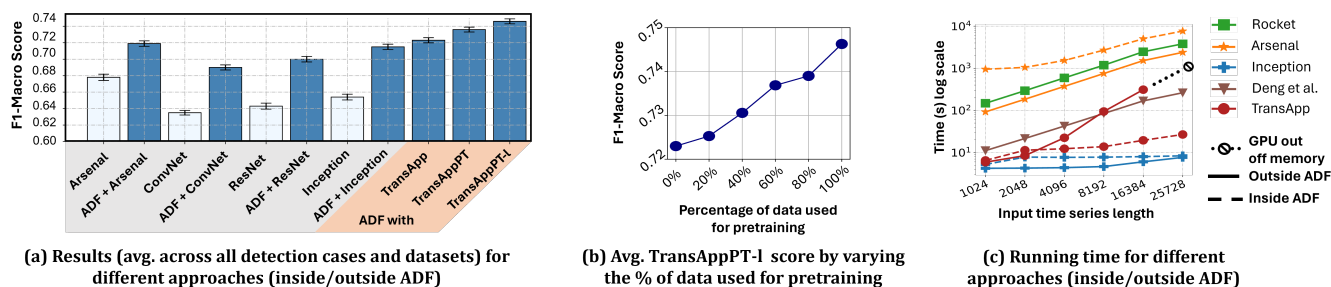


Figure 3: Overview of the experimental evaluation results: (a) Avg. detection score results for the different approaches (inside/outside ADF); (b) Impact of the amount of unlabeled data used for pretraining; (c) Running time for classifiers (inside/outside ADF) to predict 1K instances labels according to the entire input consumption series length.

5 ONGOING AND FUTURE WORK

In this work, we proposed new solutions to identify *if* the presence of a specific appliance; we are now focusing on proposing solutions to detect *when* the appliances have been used, as well as their individual associated power consumption. We briefly describe our ongoing and future work during this Ph.D.’s final year.

[DeviceScope] We are developing an interactive system named DeviceScope [7] to allow non-expert users to better understand electricity consumption data by detecting and localizing individual appliance usage patterns. The system’s core is based on a combination of trained time series classifiers and explainable classification approaches [12]. More specifically, if a time series classifier detects the presence of an appliance over a time period, a classification-based explainability method is then applied to identify the part of the time series that contributed to the label prediction, enabling the highlighting of the appliance usage pattern. To the best of our knowledge, DeviceScope is the first system to enable appliance localization using scarce labels for training.

[TransApp for energy disaggregation] Energy disaggregation aims to retrieve individual power consumption using only the total aggregated main power recorded by a smart meter. Achieving this task is valuable as this information can be used to provide a detailed breakdown of power consumption for each appliance on your electricity bill. This task is challenging as (1) the individual appliance ground-truth data are scarce for training accurate solutions, and (2) the non-stationary nature of consumption series can cause shifts in data distribution and significantly affect model performance [4]. Therefore, our future work involves enhancing the TransApp architecture to perform energy disaggregation, while accounting for the non-stationary aspect of electricity consumption data and proposing new pretraining methods to benefit from the large amount of unlabeled data currently available to suppliers.

6 CONCLUSION

This work introduces new approaches to address the appliance detection task using real-world smart meter consumption series, characterized by low-frequency readings and scarce labels. We introduced ADF, a framework designed to enhance the performance of time series classifiers on this task, and TransApp, a novel deep-learning time series classifier that is first trained in a self-supervised manner to leverage large amounts of unlabeled data. We showed that ADF significantly improves time series classifier performance

for appliance detection tasks, and TransApp combined with ADF is the most accurate solution. Based on this work, we introduce DeviceScope, an interactive application for localizing individual appliance usage patterns, using only weak labels. These preliminary results are very promising; they set the ground for further advancements in this area, including the use of weakly supervised-based approaches for appliance pattern localization and large-scale pretraining for energy disaggregation.

ACKNOWLEDGMENTS

Work partially supported by EDF R&D and the ANRT French program (Grant 2022/0179).

REFERENCES

- [1] Stanislav Chren, Bruno Rossi, and Tomáš Pitner. 2016. Smart grids deployments within EU projects: The role of smart meters. In *SCSP*.
- [2] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *North American Chapter of the Association for Computational Linguistics*.
- [3] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. 2021. Masked Autoencoders Are Scalable Vision Learners. arXiv:2111.06377 [cs.CV]
- [4] Taesung Kim, Jinhee Kim, Yunwon Tae, Cheonbok Park, Jang-Ho Choi, and Jaegul Choo. 2021. Reversible Instance Normalization for Accurate Time-Series Forecasting against Distribution Shift. In *ICLR*.
- [5] Pauline Laviron, Xueqi Dai, Bérénice Huquet, and Themis Palpanas. 2021. Electricity Demand Activation Extraction: From Known to Unknown Signatures, Using Similarity Search. In *e-Energy ’21*.
- [6] Matthew Middlehurst, Patrick Schäfer, and Anthony Bagnall. 2024. Bake off redux: a review and experimental evaluation of recent time series classification algorithms. *Data Mining and Knowledge Discovery* (19 April 2024).
- [7] Adrien Petralia, Paul Boniol, Philippe Charpentier, and Themis Palpanas. [n.d.]. *DeviceScope*. <https://devicescope.streamlit.app/>
- [8] Adrien Petralia, Philippe Charpentier, Paul Boniol, and Themis Palpanas. 2023. Appliance Detection Using Very Low-Frequency Smart Meter Time Series. In *ACM International Conference on Future Energy Systems (e-Energy)*.
- [9] Adrien Petralia, Philippe Charpentier, and Themis Palpanas. 2023. ADF & TransApp: A Transformer-Based Framework for Appliance Detection Using Smart Meter Consumption Series. *PVLDB* 17, 3.
- [10] Hasan Rafiq, Prajowal Manandhar, Edwin Rodriguez-Ubinas, Omer Ahmed Qureshi, and Themis Palpanas. 2024. A review of current methods and challenges of advanced deep learning-based non-intrusive load monitoring (NILM) in residential context. *Energy and Buildings*.
- [11] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems*, Vol. 30. Curran Associates, Inc.
- [12] Zhiguang Wang, Weizhong Yan, and Tim Oates. 2016. Time series classification from scratch with deep neural networks: A strong baseline. *IJCNN*.
- [13] Bochao Zhao, Minxiang Ye, Lina Stankovic, and Vladimir Stankovic. 2020. Non-intrusive load disaggregation solutions for very low-rate smart meter data. *Applied Energy* 268 (2020).