Deep Learning for Electricity Consumption Time Series Analytics

Presented by
Adrien Petralia

Supervised by

Prof. Themis PALPANAS

and

Philippe CHARPENTIER

Université Paris Cité

EDF

I. Introduction

II. Contributions

III.Conclusions

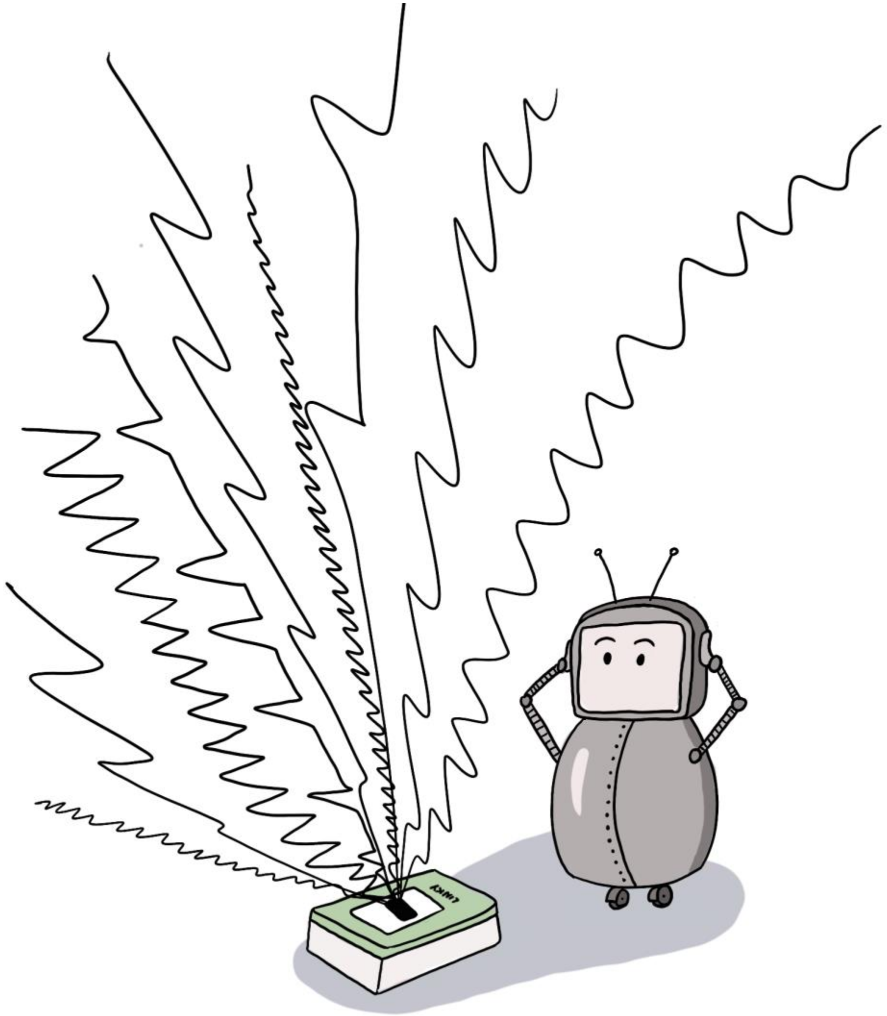

Context: Electricity Production

Renewables 2023 → 2035: From Renewable Surge to the Flexibility Challenge

- COP28 (2023) → “beginning of the end of fossil fuels”
- France 92 % low-carbon mix in 2023
- Renewables share 120 TWh → 270–320 TWh (expected in 2035)

RTE technical reports, 2023.

Renewables 2023 → 2035: From Renewable Surge to the Flexibility Challenge

- COP28 (2023) → “beginning of the end of fossil fuels”
- France 92 % low-carbon mix in 2023
- Renewables share 120 TWh → 270–320 TWh (expected in 2035)

Typical **daily** electricity **grid domestic demand**

Renewable sources peak production

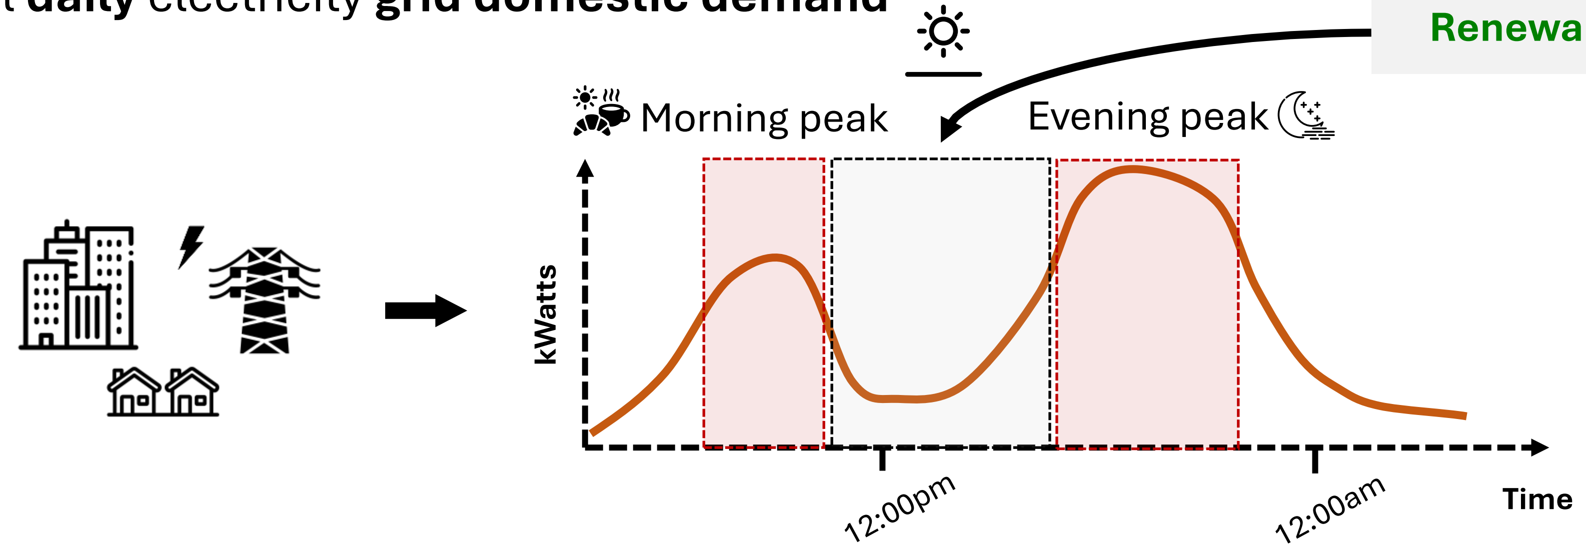

Challenge:
Flexibility

RTE technical reports, 2023.

Context: The Flexibility Challenge

Flexibility = real-time ability to match variable supply & demand

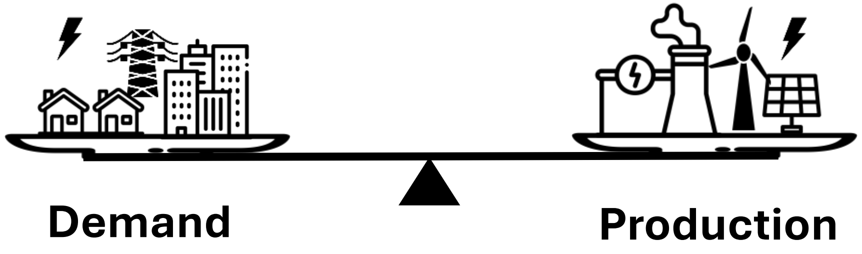

Context: The Flexibility Challenge

Flexibility = real-time ability to match variable supply & demand

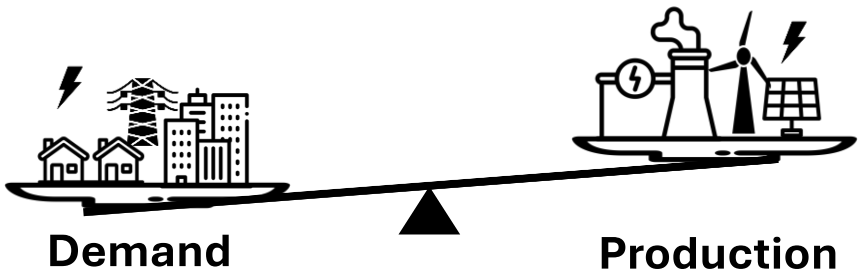

Context: The Flexibility Challenge

Flexibility = real-time ability to match variable supply & demand

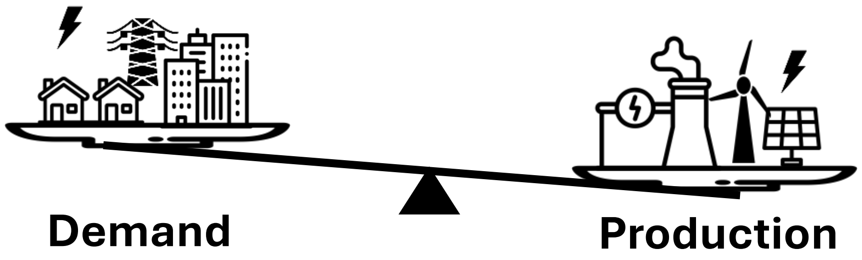

Context: The Flexibility Challenge

Flexibility = real-time ability to match variable supply & demand

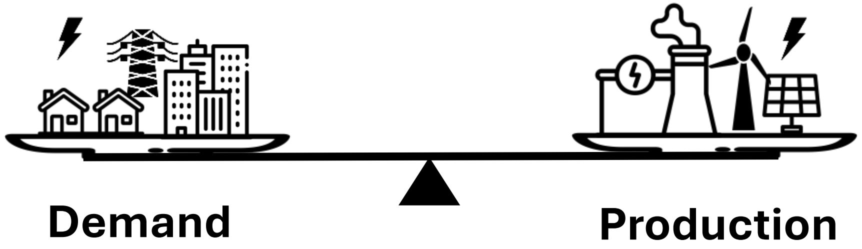

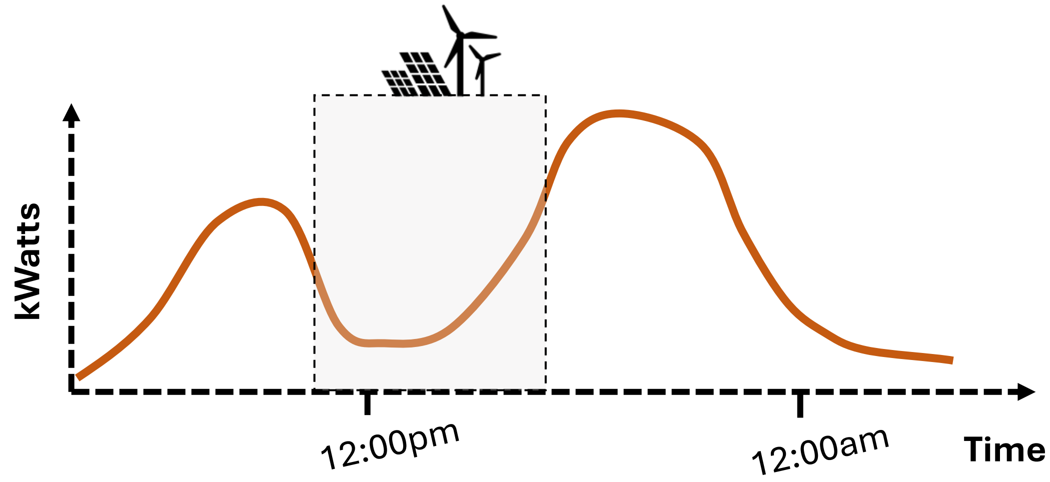

Today daily electricity domestic demand

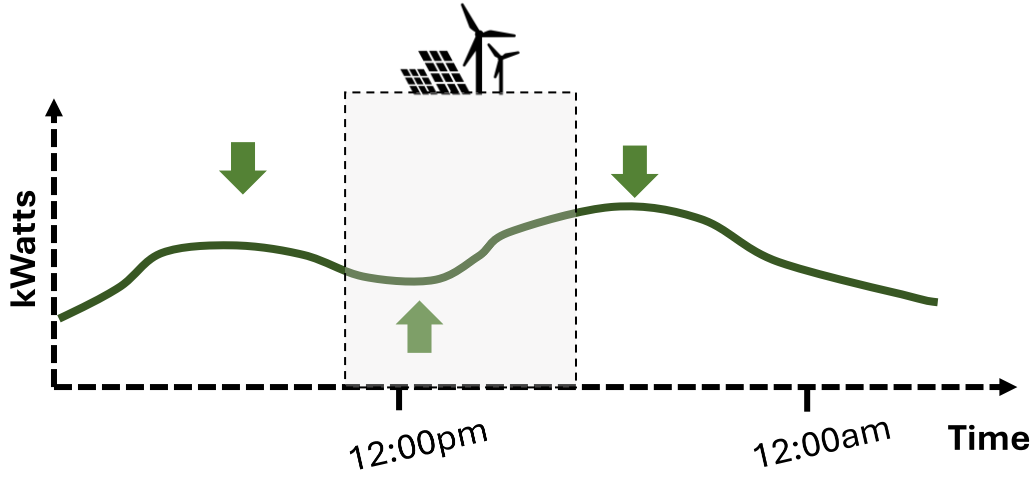

Targeted *smoothed* daily domestic demand

Context: The Flexibility Challenge

Different Possible Levers

Forecasting (solar/wind energy)

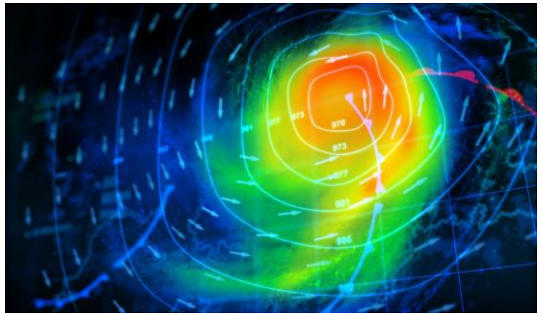

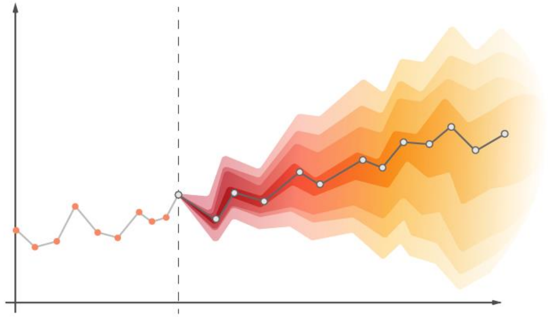

Storage (hydro, batteries)

Thesis focus

Consumer involvement

Some pictures are obtained from the Internet.

Context: Consumers Engagement for Efficient Energy Management

How can suppliers **convince consumers** to **change their behaviors** ?

1. Personalized contracts

Lower off-peak pricing, e.g., for **charging** your **Electric Vehicle** at night

2. Dynamic pricing

Criticial Peak Pricing: Higher rates during peak events
Peak Time Rebeat: Rewards for reducing use during peak times

3. Appliance-level feedback (real-time consumption awareness)

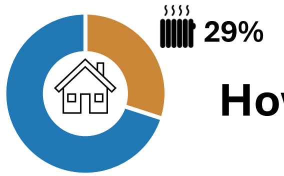

How much does your heater **cost you per month**?

Help **customers reduce** their **bill** *(up to -12 %)* [1, 2]

Solutions based on customers' characteristics/information!

[1] Martinez et al., Advanced Metering Initiatives and Residential Feedback Programs, 2010
[2] Allcott et al., Social norms and energy conservation, Journal of Public Economics, 2011

Context: Smart Meter Deployement

Millions of Smart Meters deployed in **individual households**

Grid network

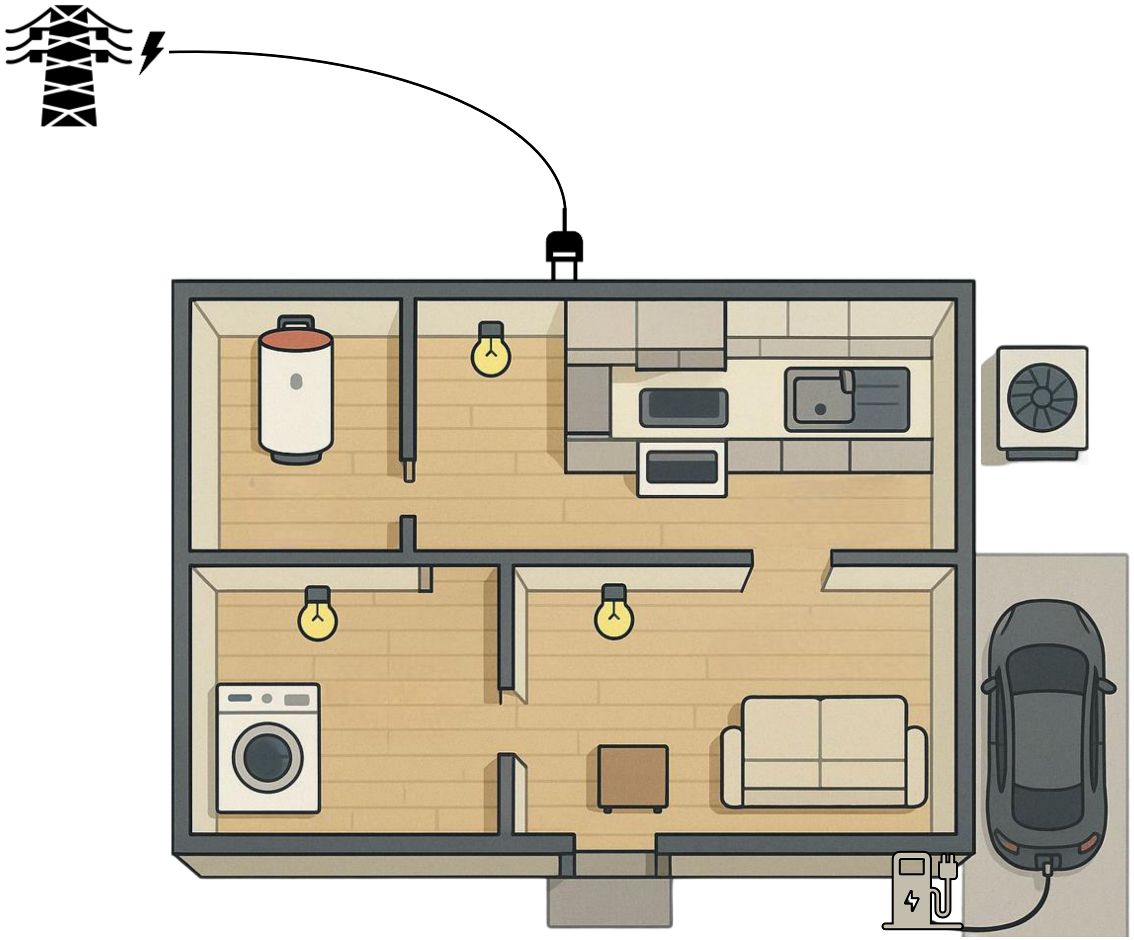

Context: Smart Meter Deployement

Millions of Smart Meters deployed in **individual households**

Bill clients **more accurately**

Better manage the **smart grid**

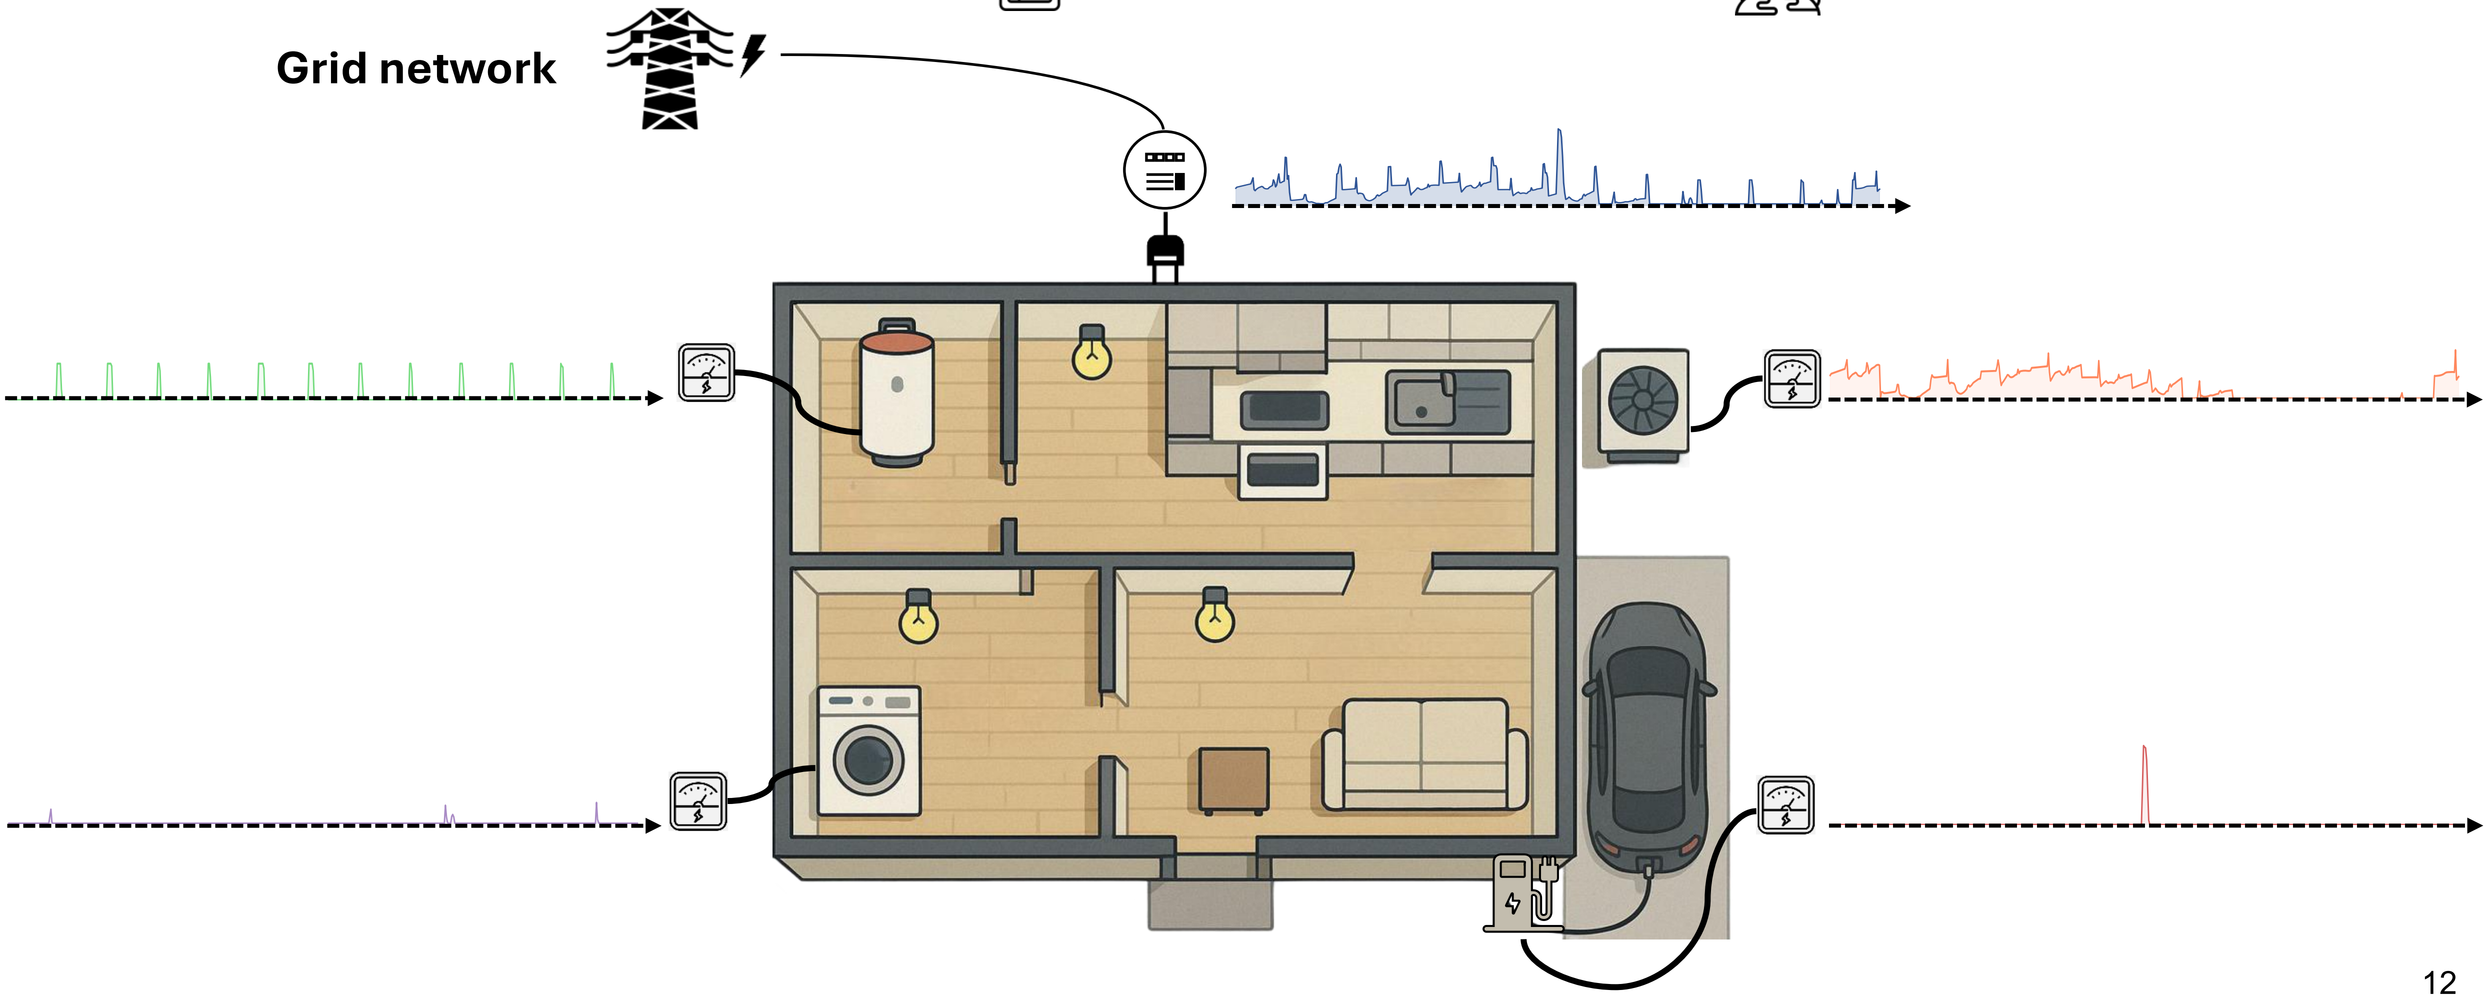

Millions of Smart Meters deployed in **individual households**

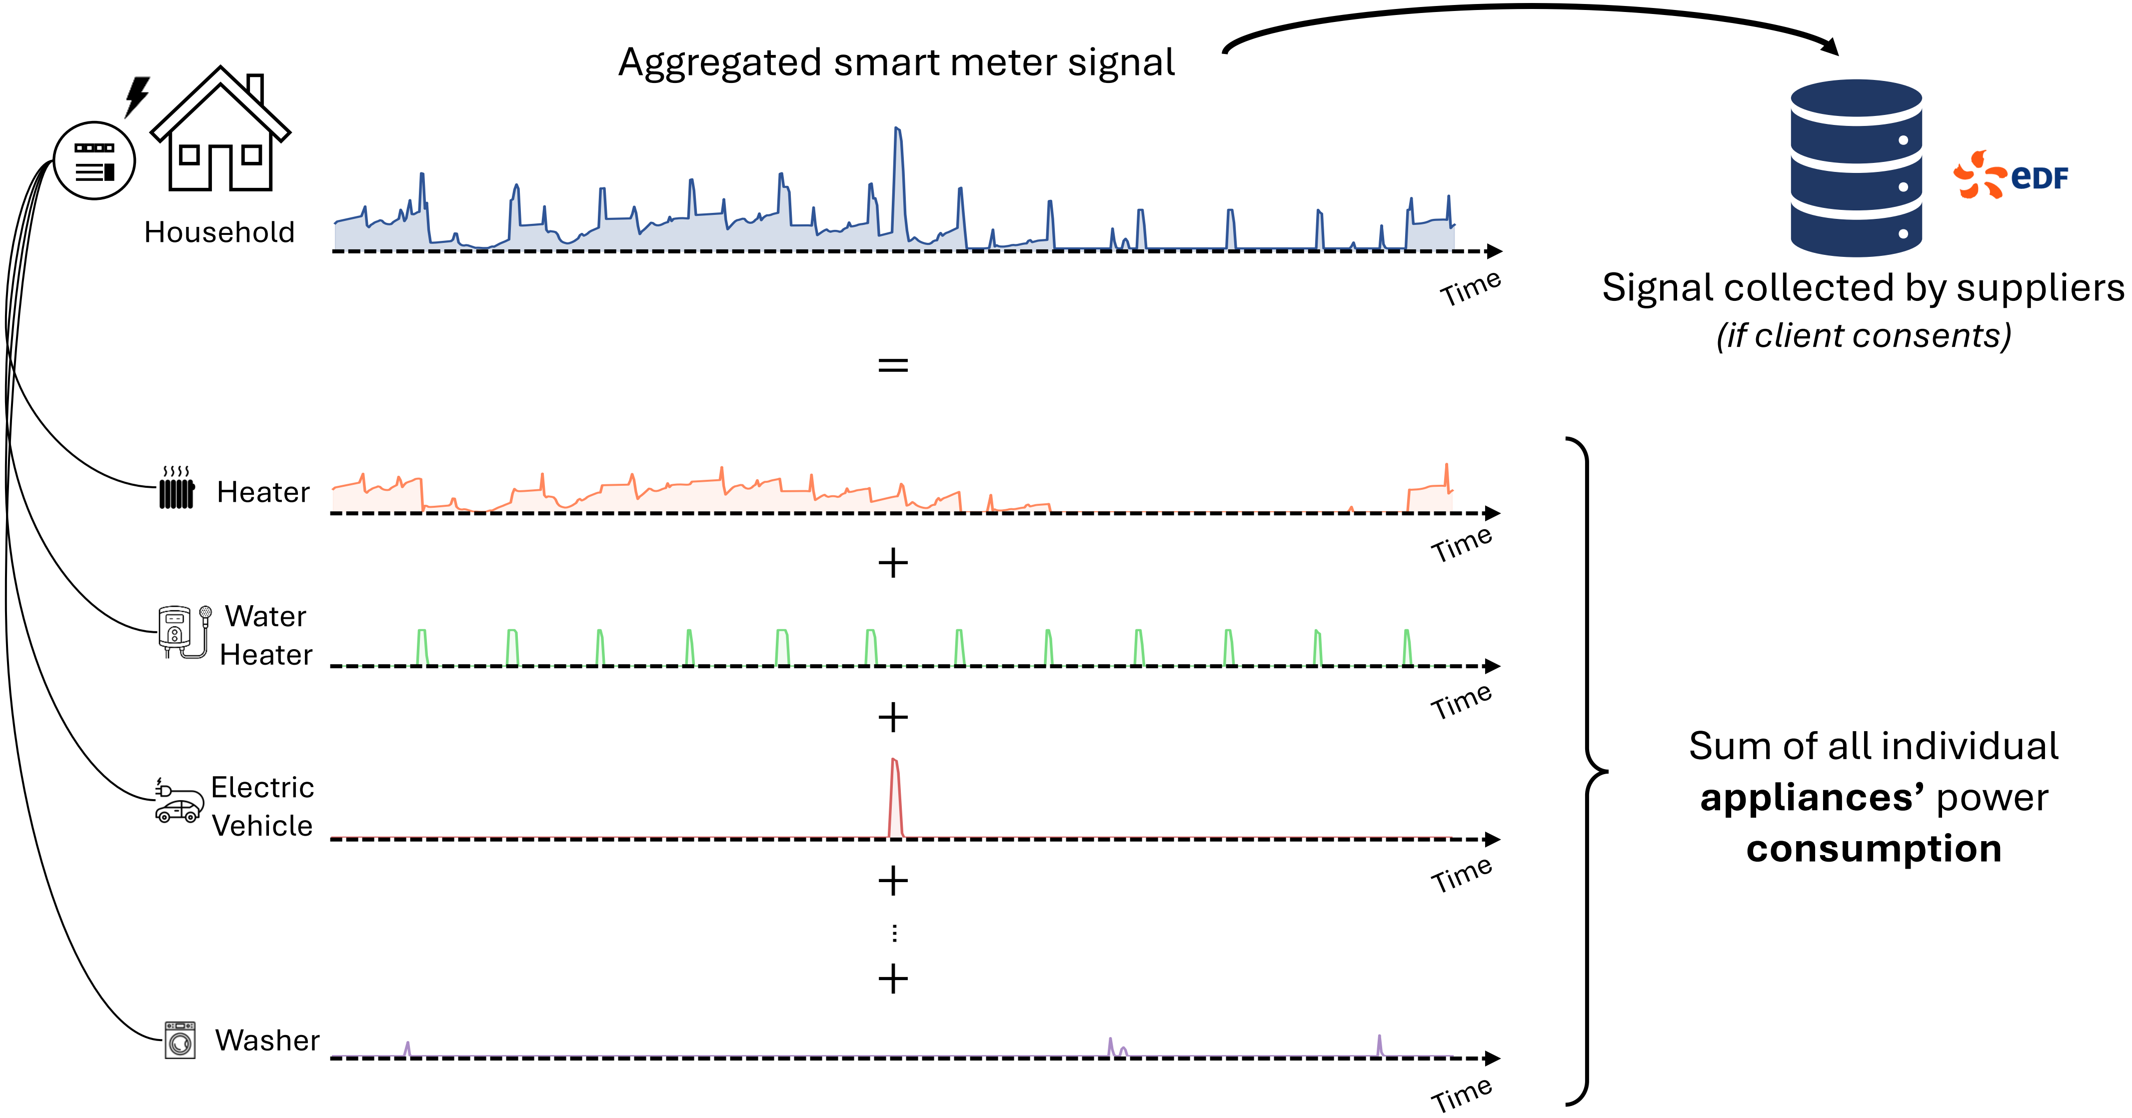

Context: Smart Meter Deployement

These data are stored in large **electricity consumption databases**

Electricity consumption database
(Millions of clients)

Recorded smart meter consumption

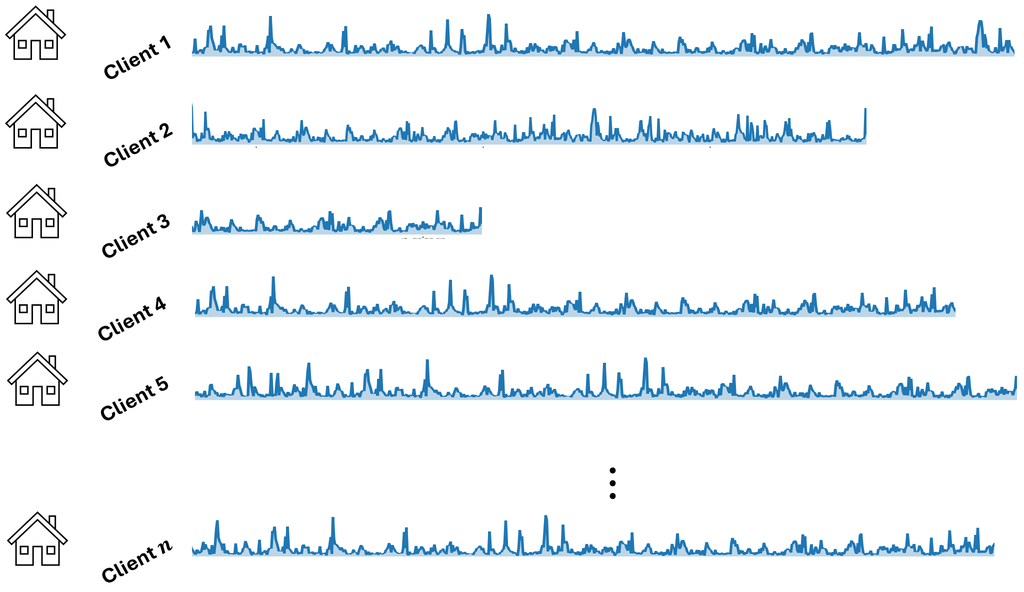

Characteristics	**Used for**
*Which **appliances** are **present** in the **house**?*	**Segment** Customer Base ***Personalized Contracts***
***How** does the client use **them**?*	**Understand** Appliance Uses ***Dynamic pricing***
*What **proportion** of **consumption** does each **appliance** account for?*	**Deliver** Appliance Feedback ***Dynamic Pricing*** *and* ***Real-time consumption awareness***

Intrusive Load Monitoring: sub-metering instrumentation

Expensive *500 to 1500$/house* **Time-consuming** **Not well-received**

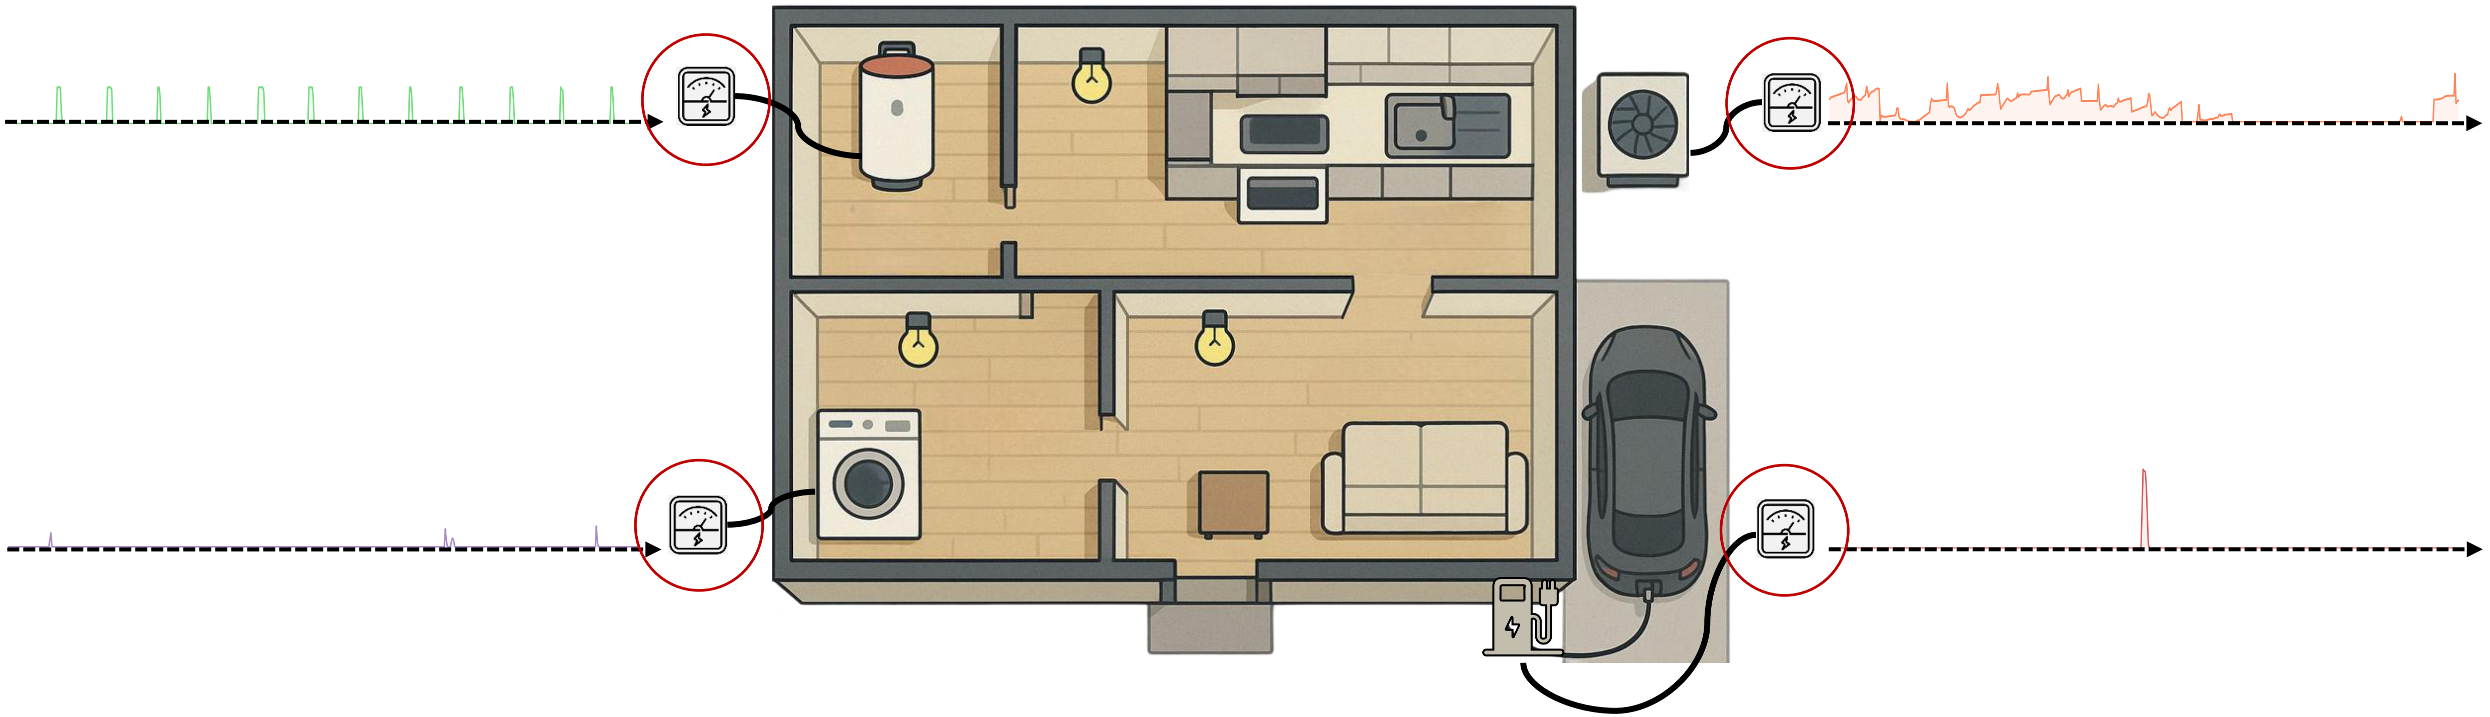

Background: Electricity Consumption Time Series Analytics

Non-Intrusive **L**oad **M**onitoring (NILM): estimates **power consumption**, operational **patterns**, or **on/off state** of individual appliances using **only the total aggregated signal**

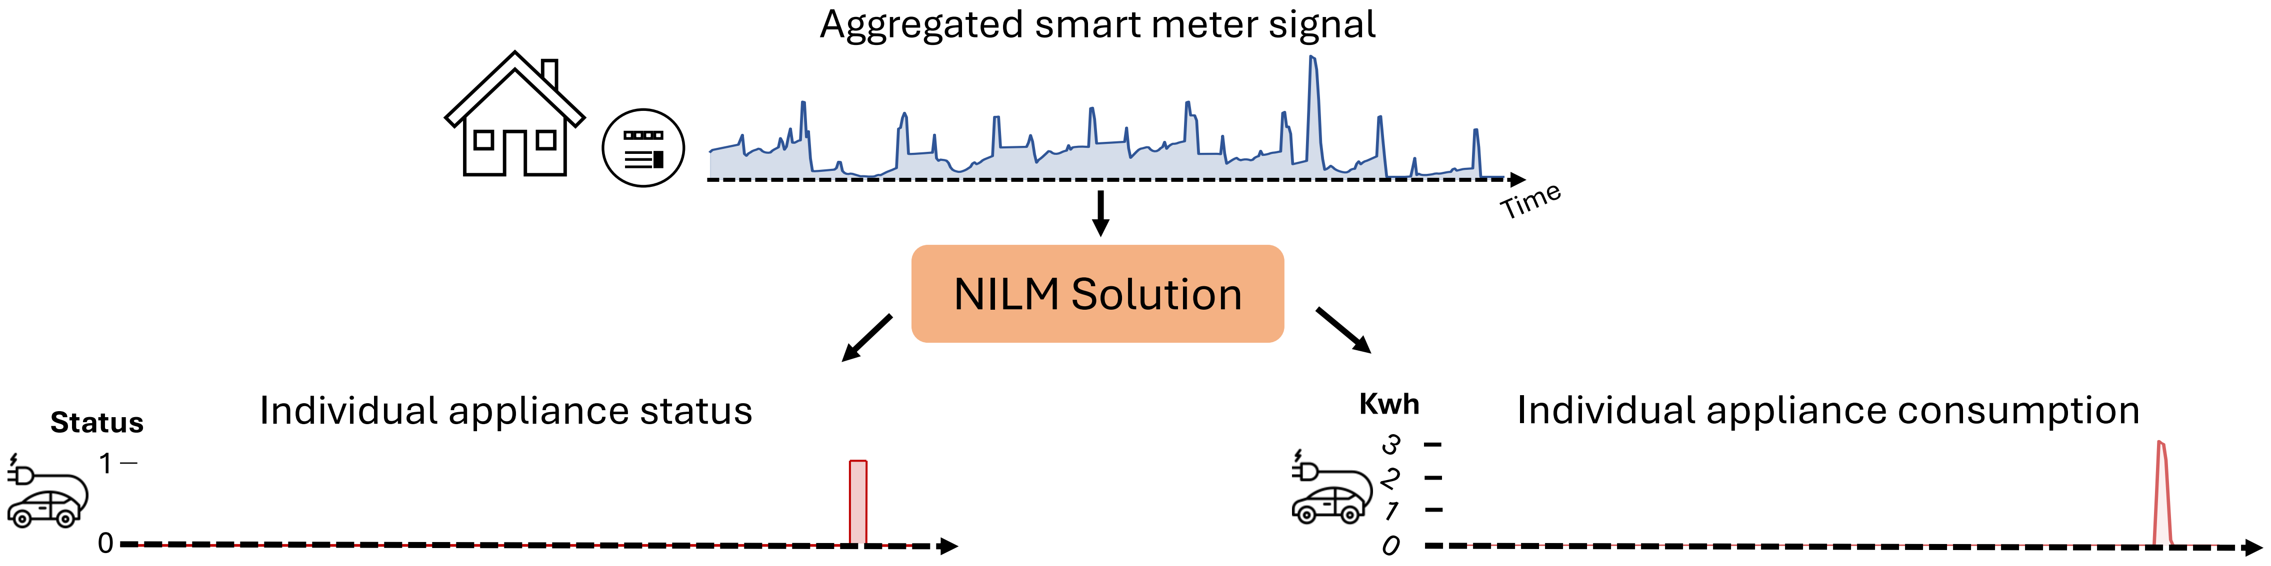

Early research (1992)	ML Area (2010's)	DL Area (2015-now)
Combinatorial Optimization G. W. Hart [1]	Sparse Coding, HMM Andrew Ng [2]	RNN, CNN, Transformer Jack Kelly [3]

[1] G. Hart, Nonintrusive appliance load monitoring, 1992 [2] Andrew Ng, Energy Disaggregation via Discriminative Sparse Coding, NIPS, 2011 [3] Jack Kelly, Neural NILM, ACM BuildSys', 2015

Recorded **Electricity Consumption Signals** are **Time Series** Data

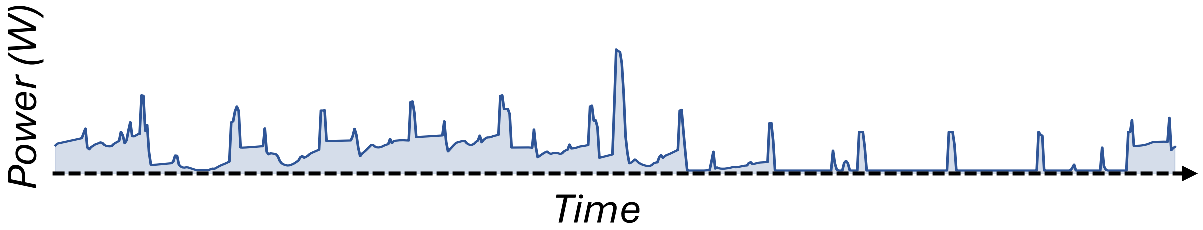

Sequence of ordered consumption index over time

Time Series Analytics

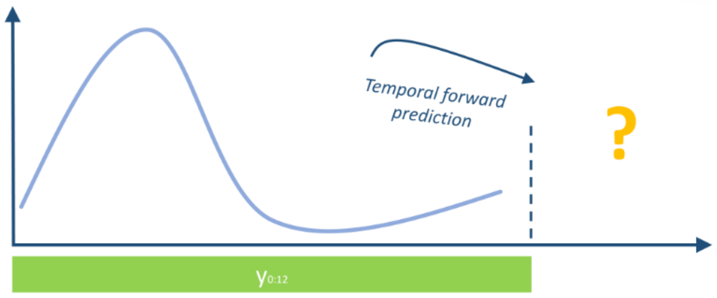

Forecasting

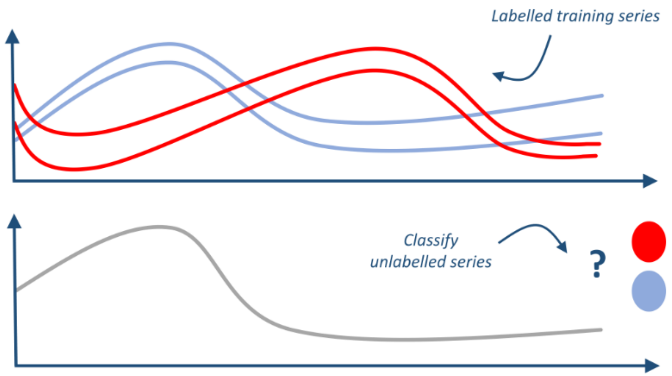

Classification

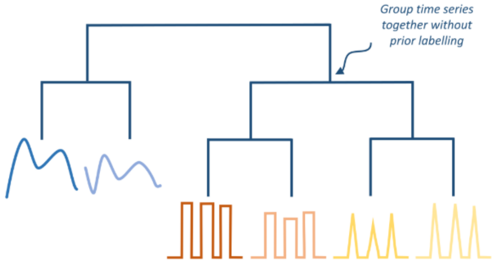

Clustering

...

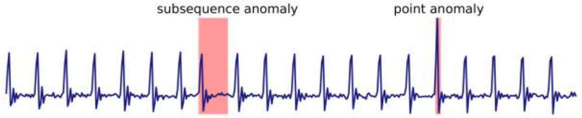

Anomaly Detection

Middlehurst et al., aeon: a Python toolkit for learning from time series, 2024

Challenge: Very Low-Frequency Smart Meter Reading

Common **smart meters** collect data at a **very low-frequency**

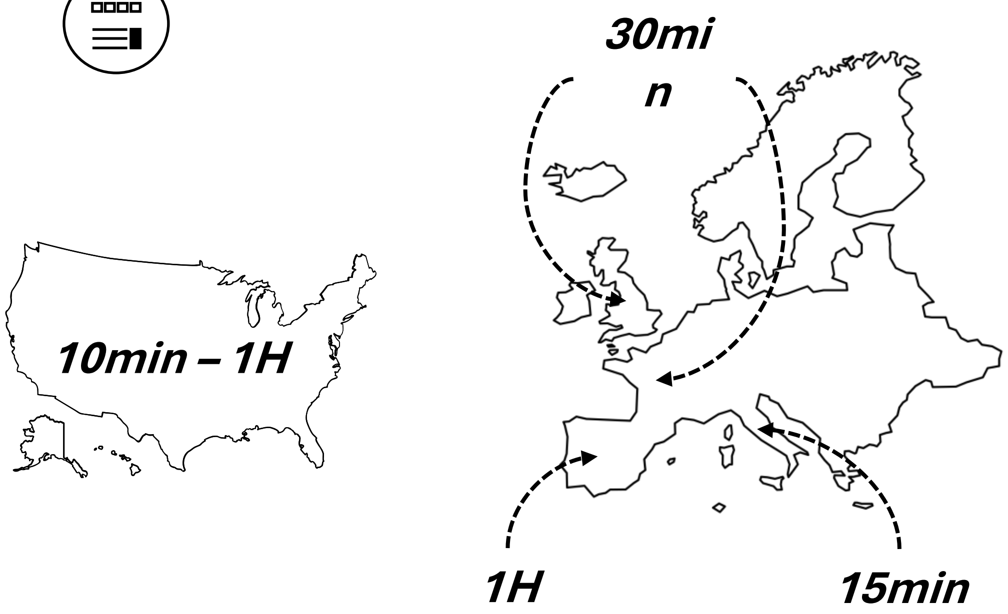

≈ ***10*** to ***60min !***

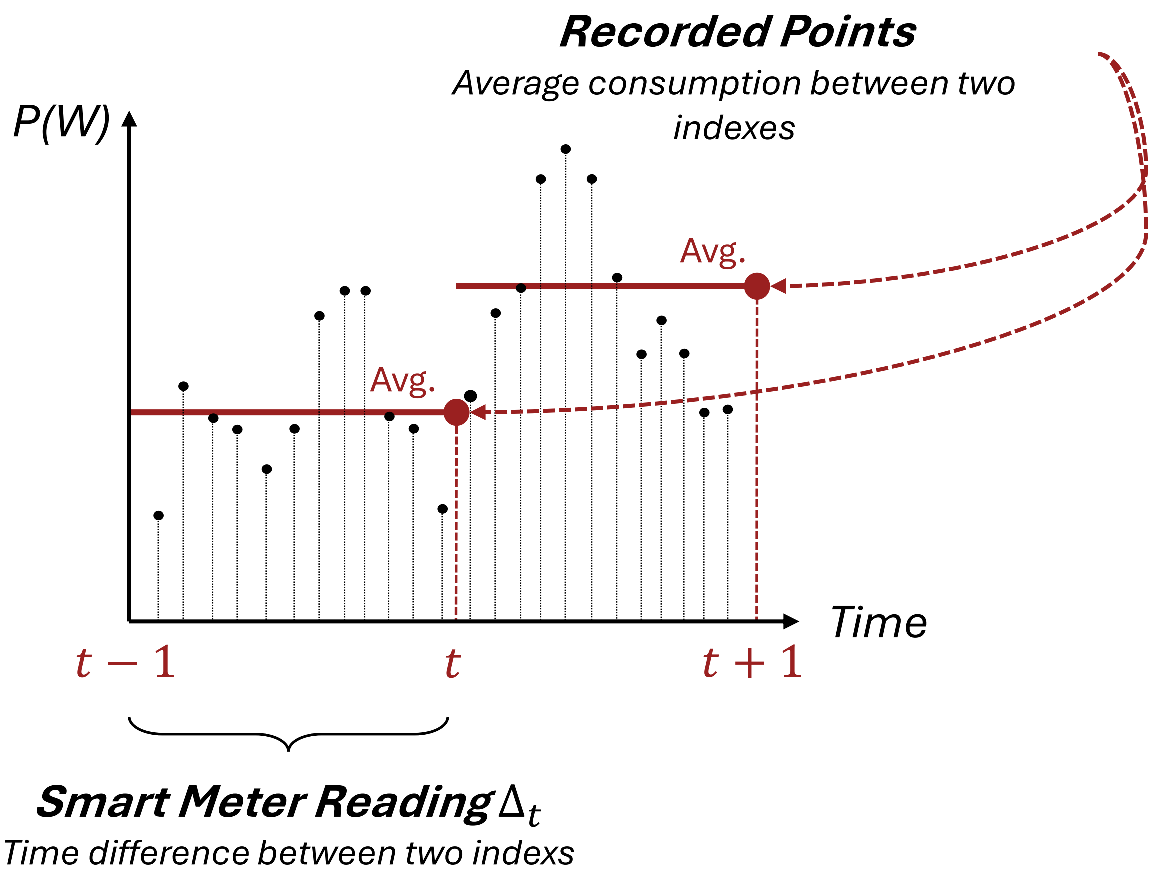

Challenge: Very Low-Frequency Smart Meter Reading

Effect of **Smart Meter Granularity** on Electricity Consumption Curve **Shape**

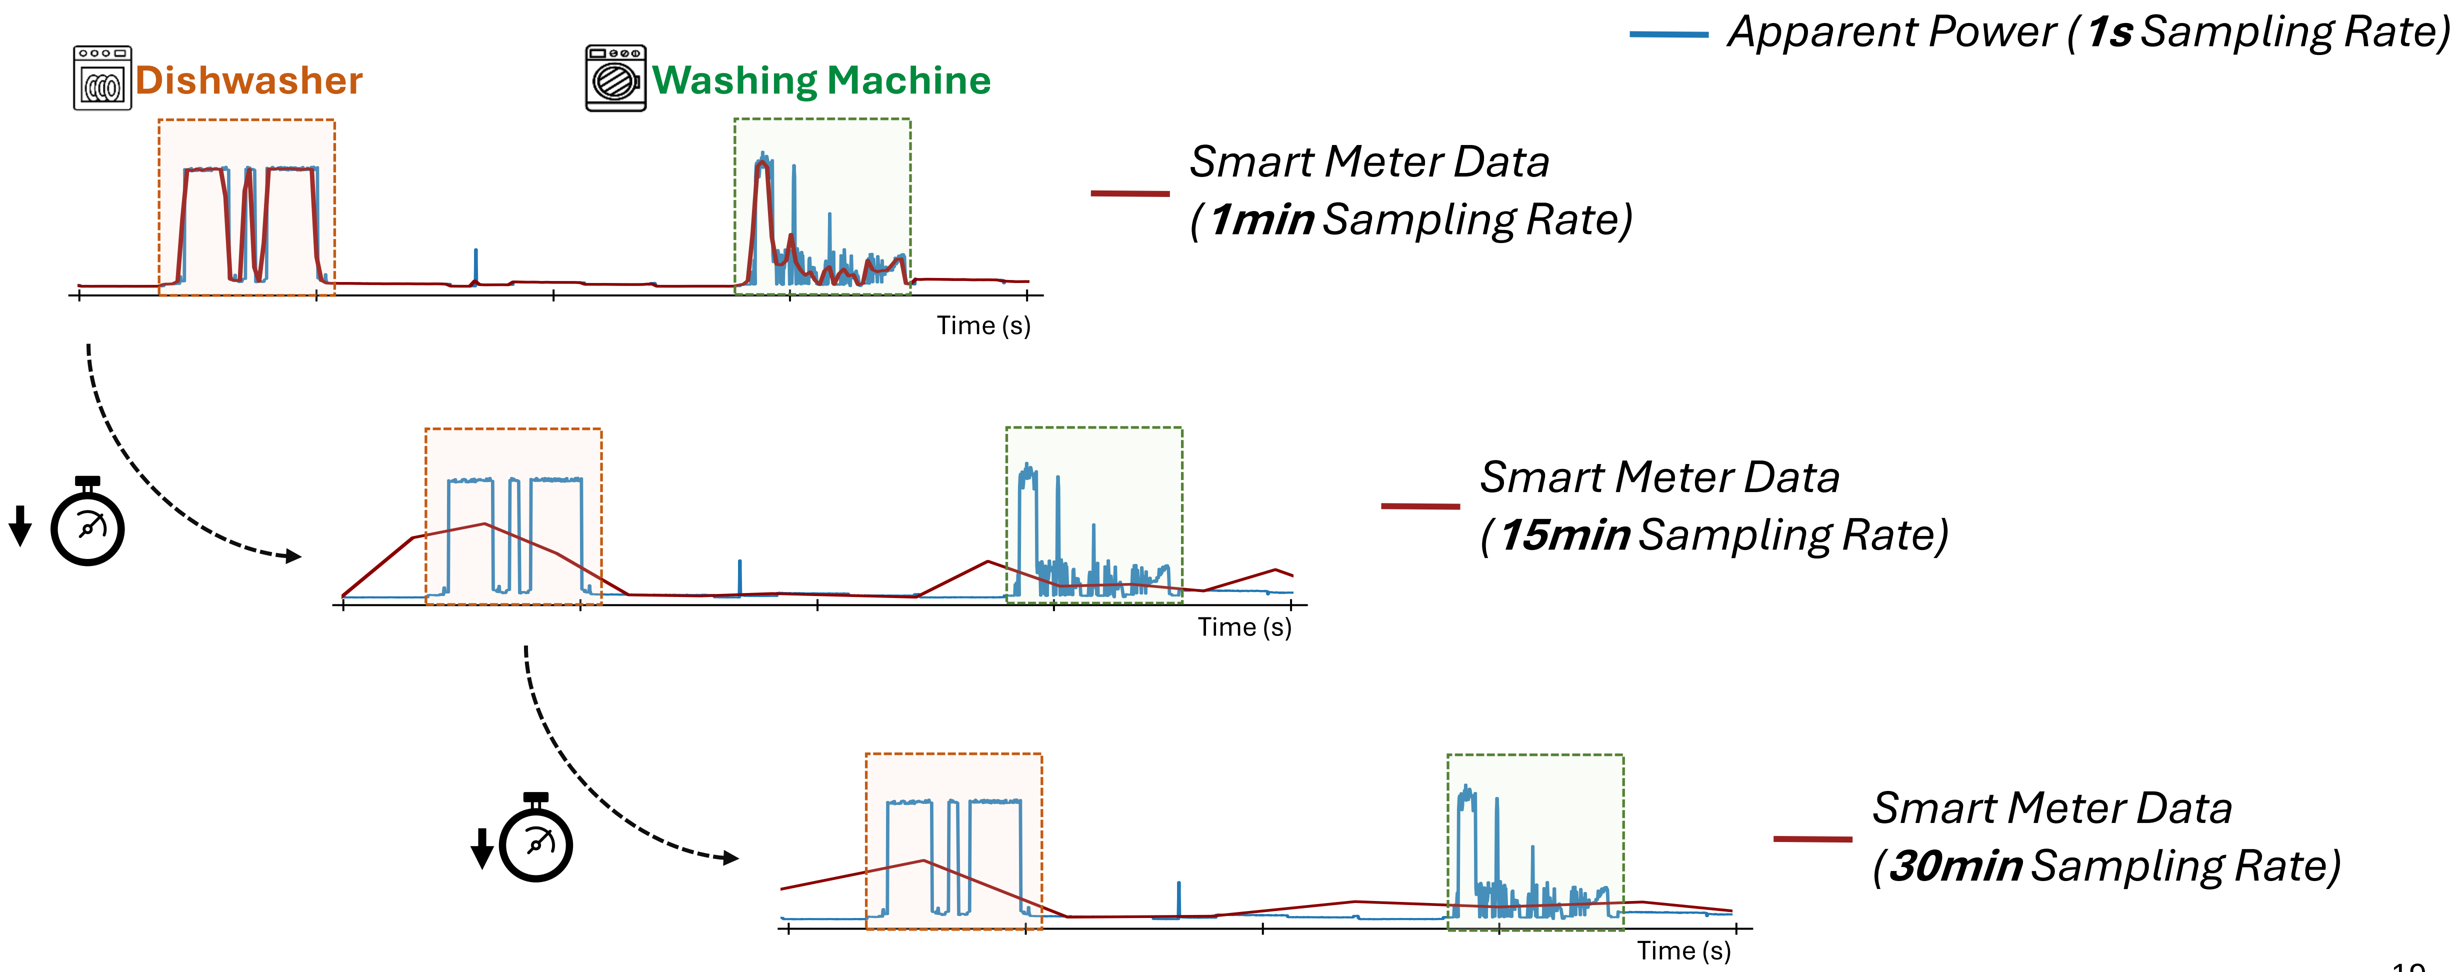

Thesis Challenges

*Can we extract **relevant information** from **electricity consumption time series** collected by **smart meters** at a **very low frequency**?*

Objectives

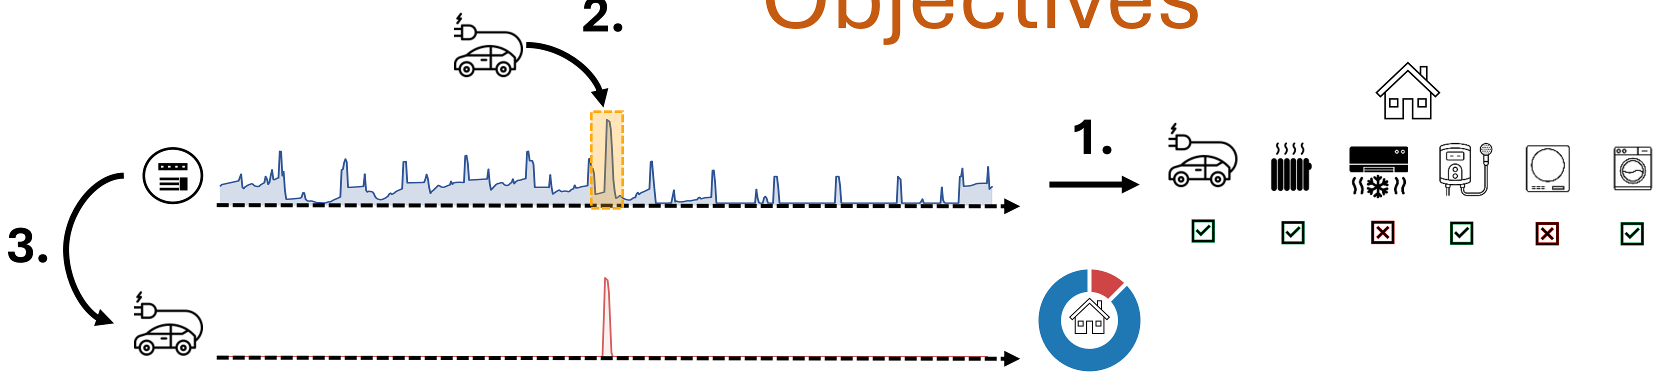

1.

Detect appliances present in the house

Segment
Customer Base

Personalized Contracts

2.

Localize appliance's activation time

Understand
Appliance Uses

Dynamic Pricing

3.

Estimate appliance's consumption

Deliver
Detailed Appliance Feedback

Dynamic Pricing
Real-time Consumption Awareness

I. Introduction

II. Contributions

1. Appliance Detection Presence in Consumer Households

2. Appliance Pattern Localization

3. Energy Disaggregation

III. Conclusions

1. Appliance Detection

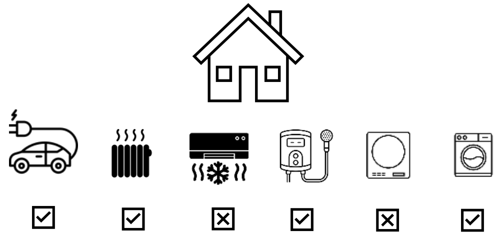

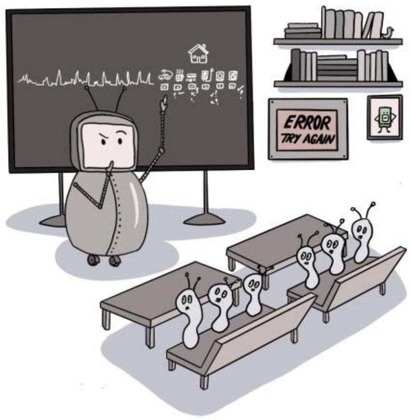

Background: Appliance Recognition

A subfield of **NILM**

*Example of signatures (**1sec** sampled data)*

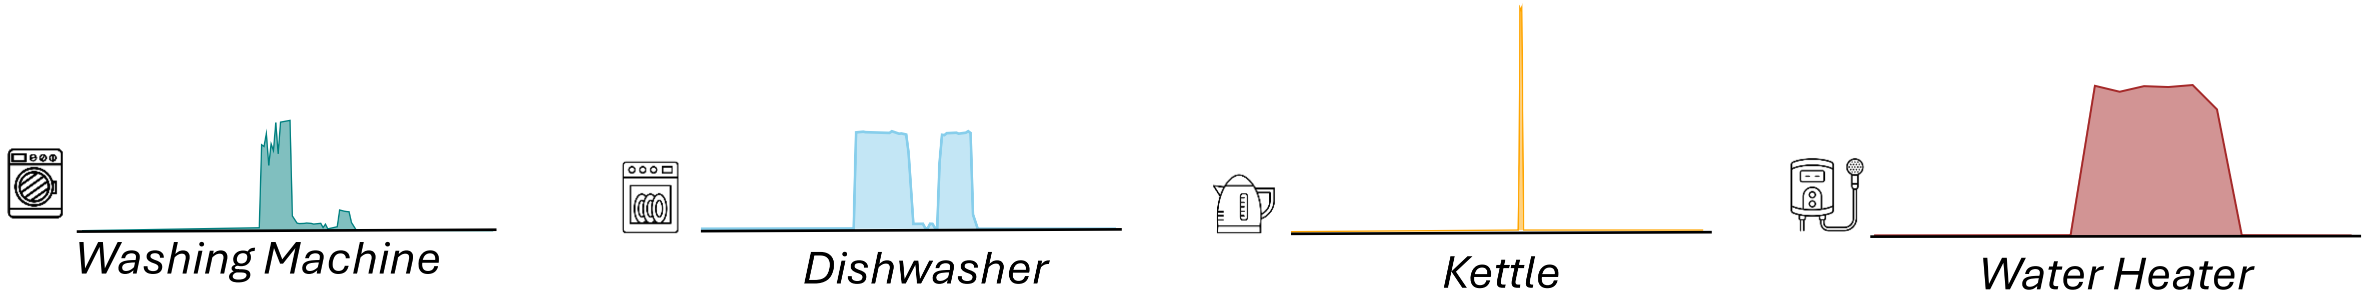

Based on **appliance signature matching** or **event detection** using **high frequency sampled data** (> 1Hz) [1, 2]

Approaches **not applicable** using **30min** sampled data!

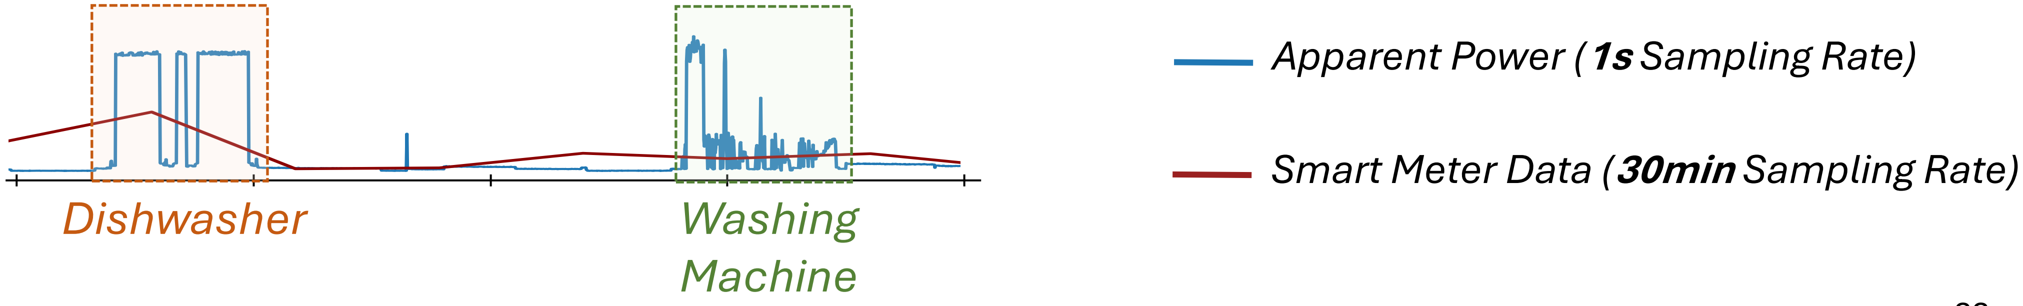

[1] P. Lavrion et al., Electricity Demand Activation Extraction: From Known to Unknown Signatures, Using Similarity Search, ACM e-Energy, 2021
[2] H. Rafiq et al., A review of current methods and challenges of advanced deep learning-based non-intrusive load monitoring (NILM) in residential context, Energy&Building, 2023

Background: EDF Comsumption Surveys Data

EDF conducts **surveys** on **customer samples**

*"Is the appliance **X** present in your household?"*

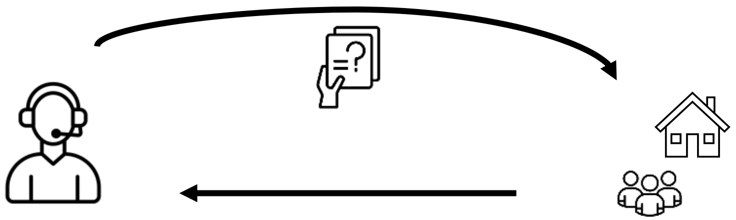

Customers fill out a **questionnaire** in exchange for a small reward

Survey dataset

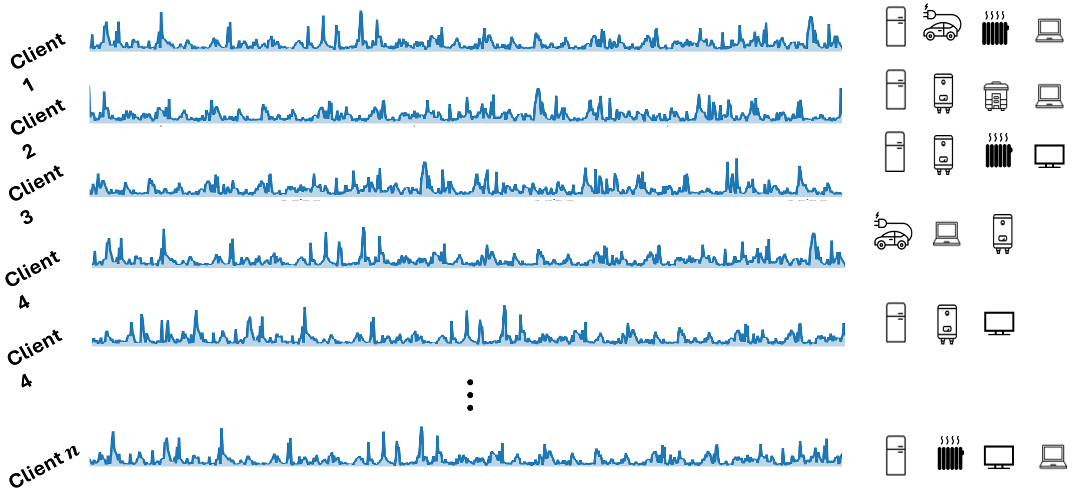

Proposed approach: Appliance Detection as a Time-Series Classification Problem

Survey Dataset

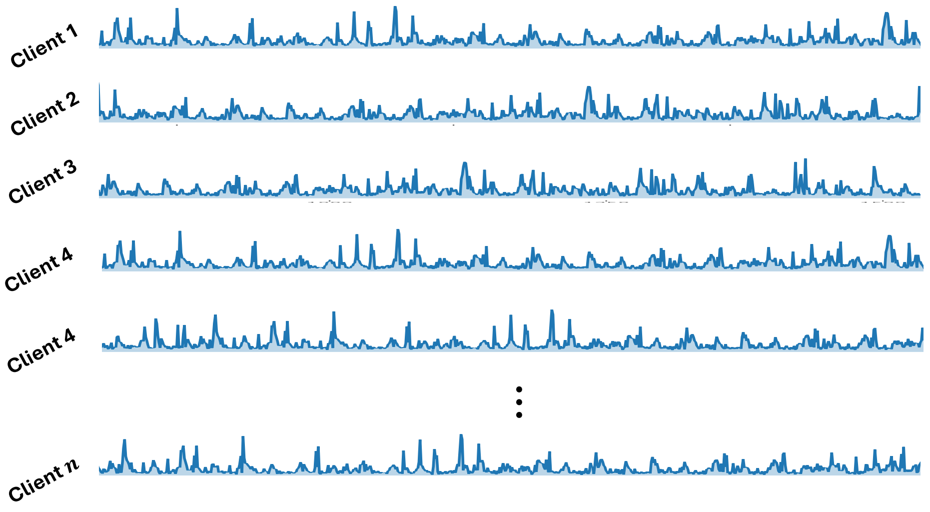

Aggregated Consumption Series

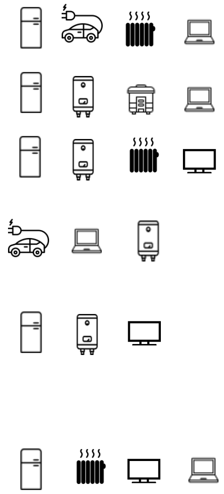

Labels

Weak labels

=

*No guarantee about **when***
(or even if)
*the **appliance is actually in use***

What if we approach this as a
Time Series Classification Problem?

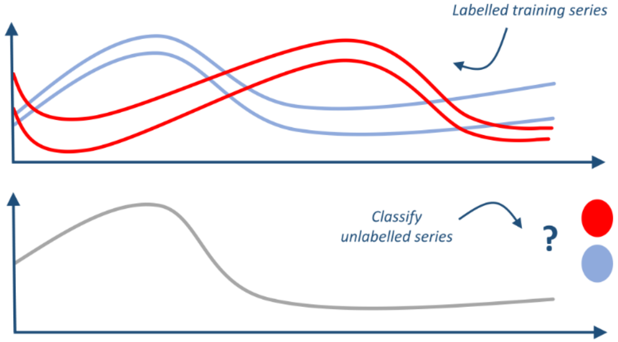

Time Series Classification

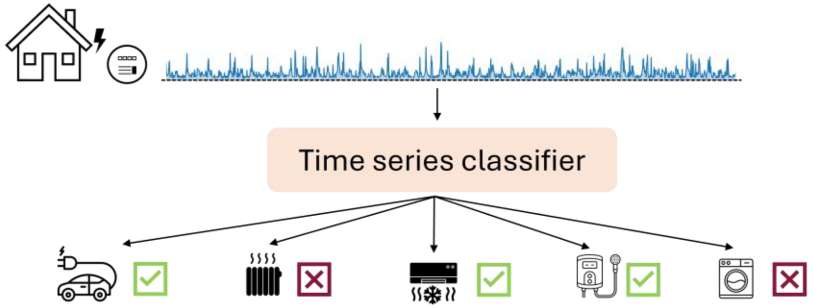

Various **time series classifiers exists** in the literature

- ❑ **Nearest-neighbor based classifiers**
  - KNN with Euclidian distance
  - KNN with Dynamic Time Warping

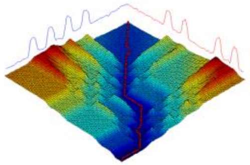

- ❑ **Motifs discovery based classifiers**
  - Shapelets Transform Classifier

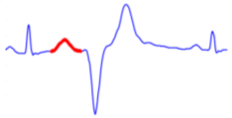

- ❑ **Dictionnary-based classifiers**
  - BOSS
  - BOSS ensemble
  - cBOSS ensemble

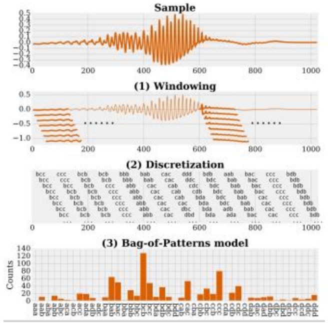

- ❑ **Tree-based classifiers**
  - Time Series Forest
  - RISE
  - DrCIF

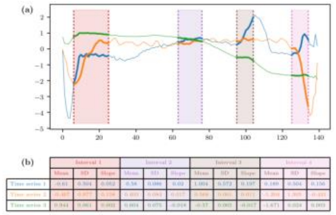

- ❑ **Random Convolutional based classifiers**
  - Rocket
  - MiniRocket
  - Arsenal

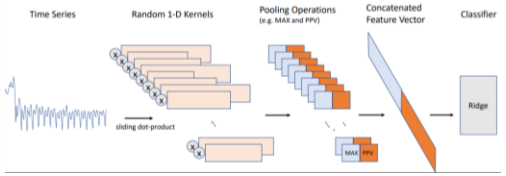

- ❑ **Deep-learning based classifiers**
  - ConvNet
  - ResNet
  - InceptionTime

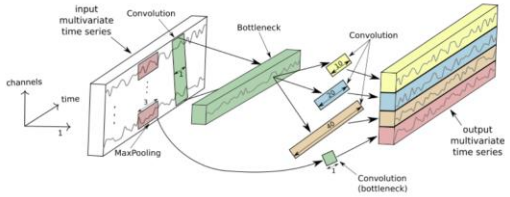

Which type of classifier delivers the best performance for our task?

Middlehurst et al., aeon: a Python toolkit for learning from time series, 2024

Key Takeaway: Appliance Detection as a Time-Series Classification Problem

Overall results for all **classifiers**, using **30min sampled data**

Convolutional-based methods *and, specifically,* ***deep-learning*** *approaches* ***perform the best*** *on this task****!***

Average detection score (F1-Score Macro)

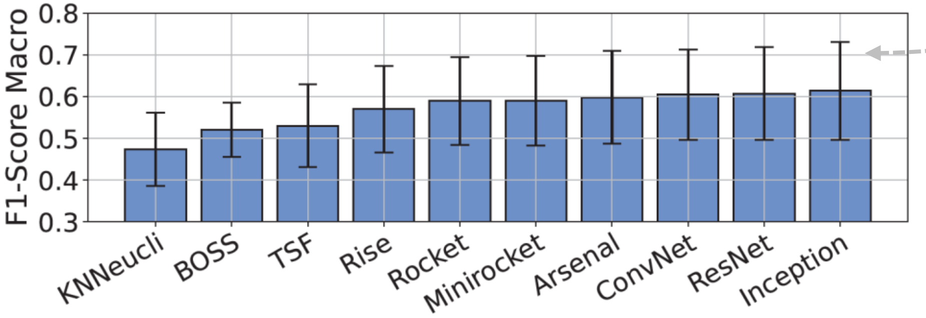

Variability across appliance detection scenarios

However, using **30min** sampled data **does not** allow for the detection of **all appliances**

Detectable, fair

- Heater
- Water Heater

Detectable, but **insufficient**

- Electric Vehicle
- Air Conditionner
- Heat Pump
- Dishwasher
- Tumble Dryer
- TV/Computer

Undetectable

- Microwave
- Kettle
- Oven
- Washer

A. Petralia et al., Appliance Detection Using Very Low-Frequency Smart Meter Time Series, ACM e-Energy 2023

Key Takeaway: Appliance Detection as a Time-Series Classification Problem

Synthesis

Detecting appliances can be **cast** as a **Time Series Classification Problem** using **30min** sampled data

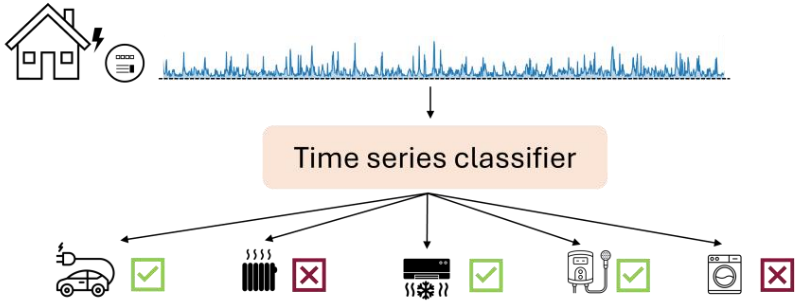

Deep Learning *solutions are the most* ***efficient*** *and* ***accurate***

Limitations

Reported **accuracy** is still **relatively low** for **real-world applications…**

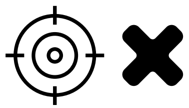

Doesn't take into account the **variable length aspect** of the series

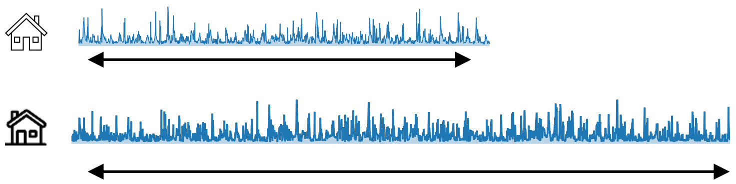

A. Petralia et al., Appliance Detection Using Very Low-Frequency Smart Meter Time Series, ACM e-Energy 2023

Lever: Large Amount of Unlabeled Electricity Consumption Data

Suppliers collect increasingly larger amounts of **electricity consumption data**

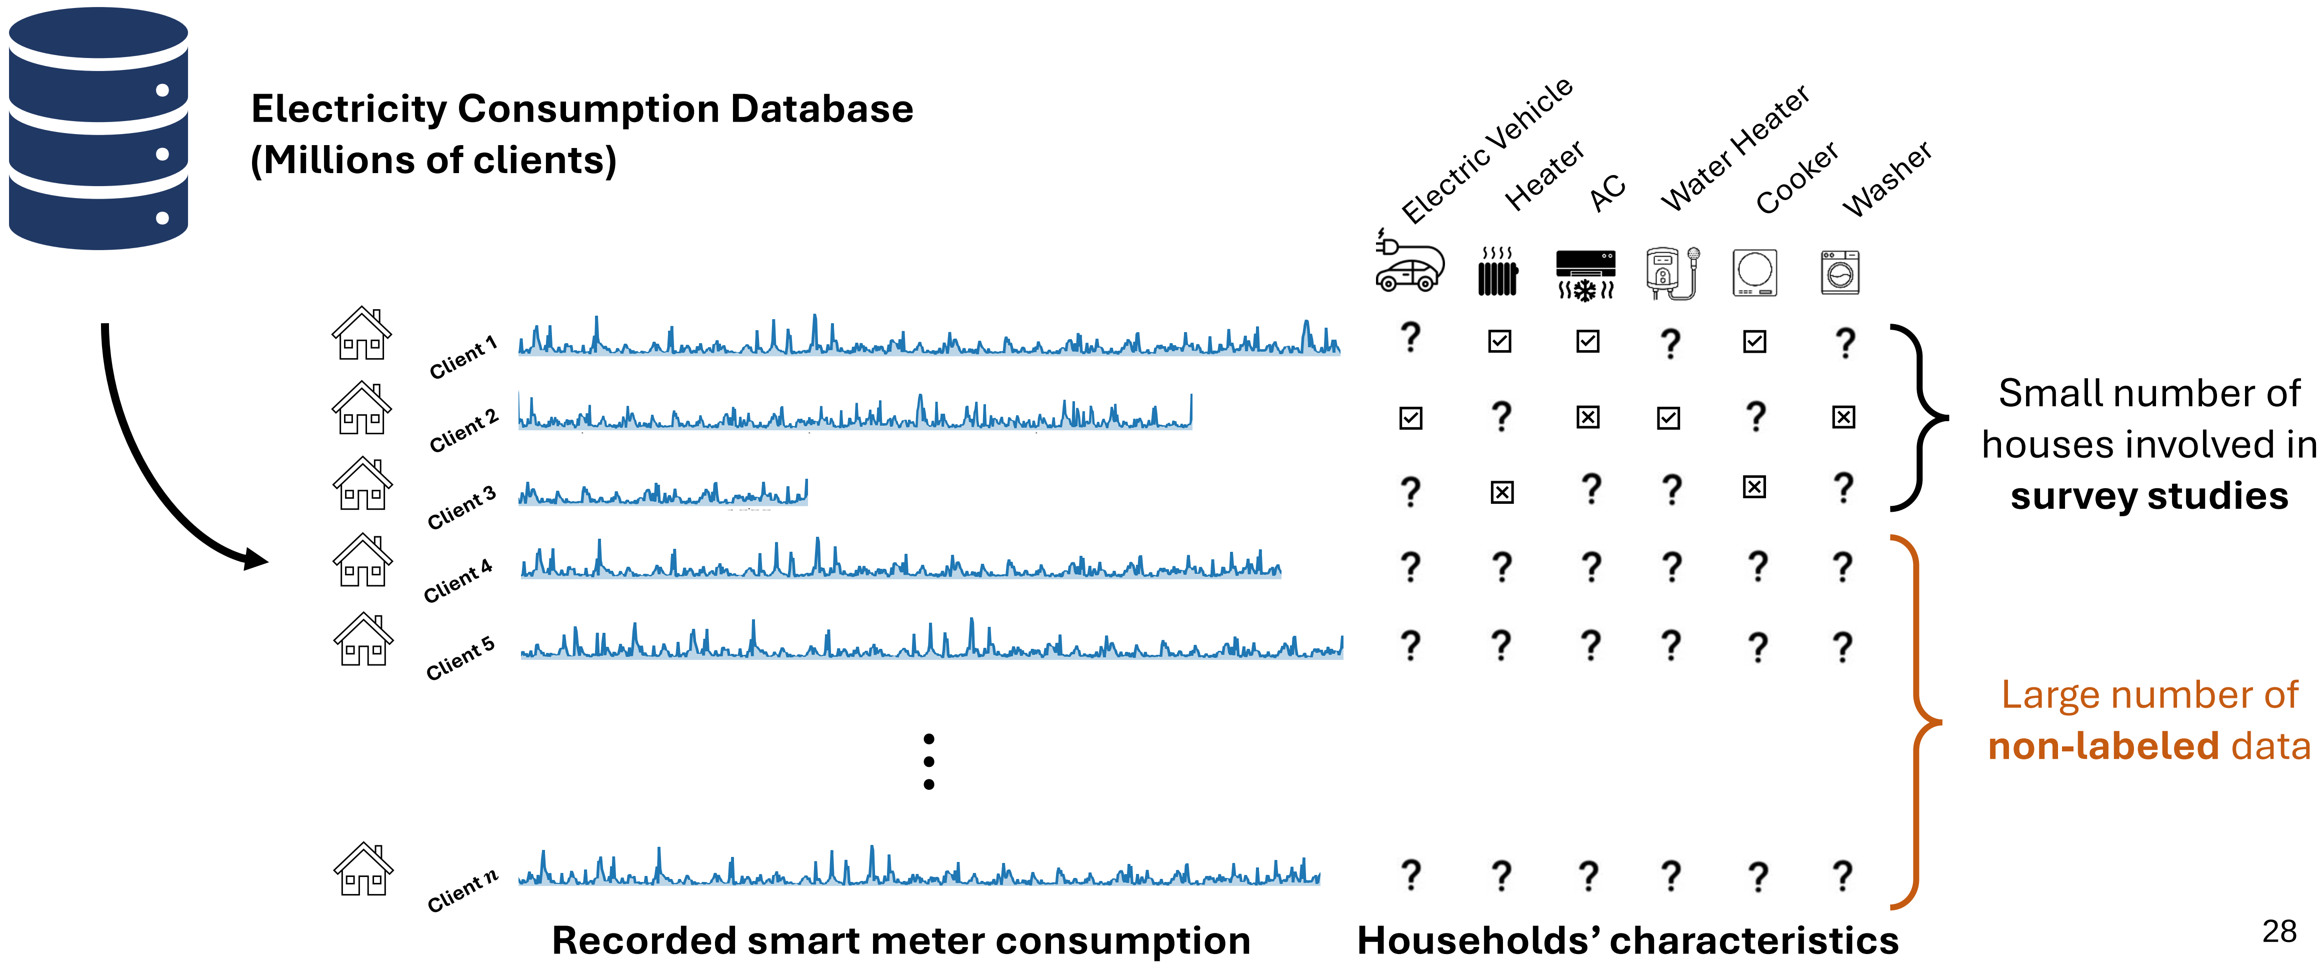

Problem

*How to **enhance** the **accuracy** of **Appliance Detection Presence** in households using **very** low-frequency smart meter data?*

Challenges

1. **Nature** of **electricity consumption data**
   Very low frequency reading used by Smart Meters
   Long and **variable length** consumption series

2. **Data size**
   Few labeled data for training a solution
   Large amount of non-labeled data

Problem

*How to **enhance** the **accuracy** of **Appliance Detection Presence** in households using **very** low-frequency smart meter data?*

Challenges

1. **Nature** of **electricity consumption data**
   Very low frequency reading used by Smart Meters
   Long and **variable length** consumption series

2. **Data size**
   Few labeled data for training a solution
   Large amount of non-labeled data

Solutions

✓ The **A**ppliance **D**etection **F**ramework (**ADF**)

✓ **TransApp**: a deep-learning time series classifier

Proposed Approach: ADF

The **A**ppliance **D**etection **F**ramework

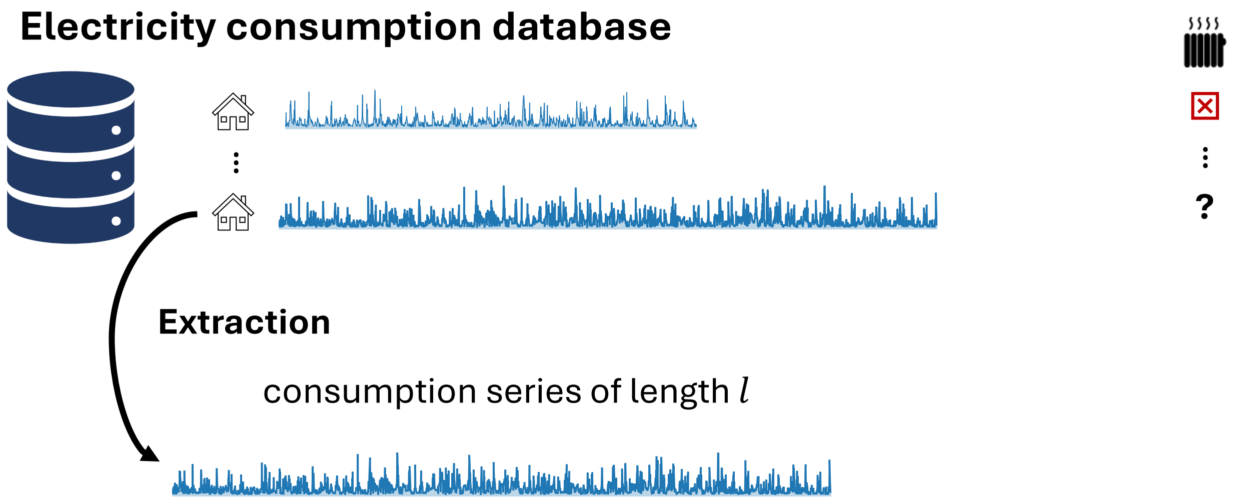

A. Petralia et al., ADF & TransApp: A Transformer-Based Framework for Appliance Detection Using Smart Meter Consumption Series, PVLDB, 2024.

Proposed Approach: ADF

The **A**ppliance **D**etection **F**ramework

1. **Slice** series into subsequences and **concatenate** them with timestamp-encoded information

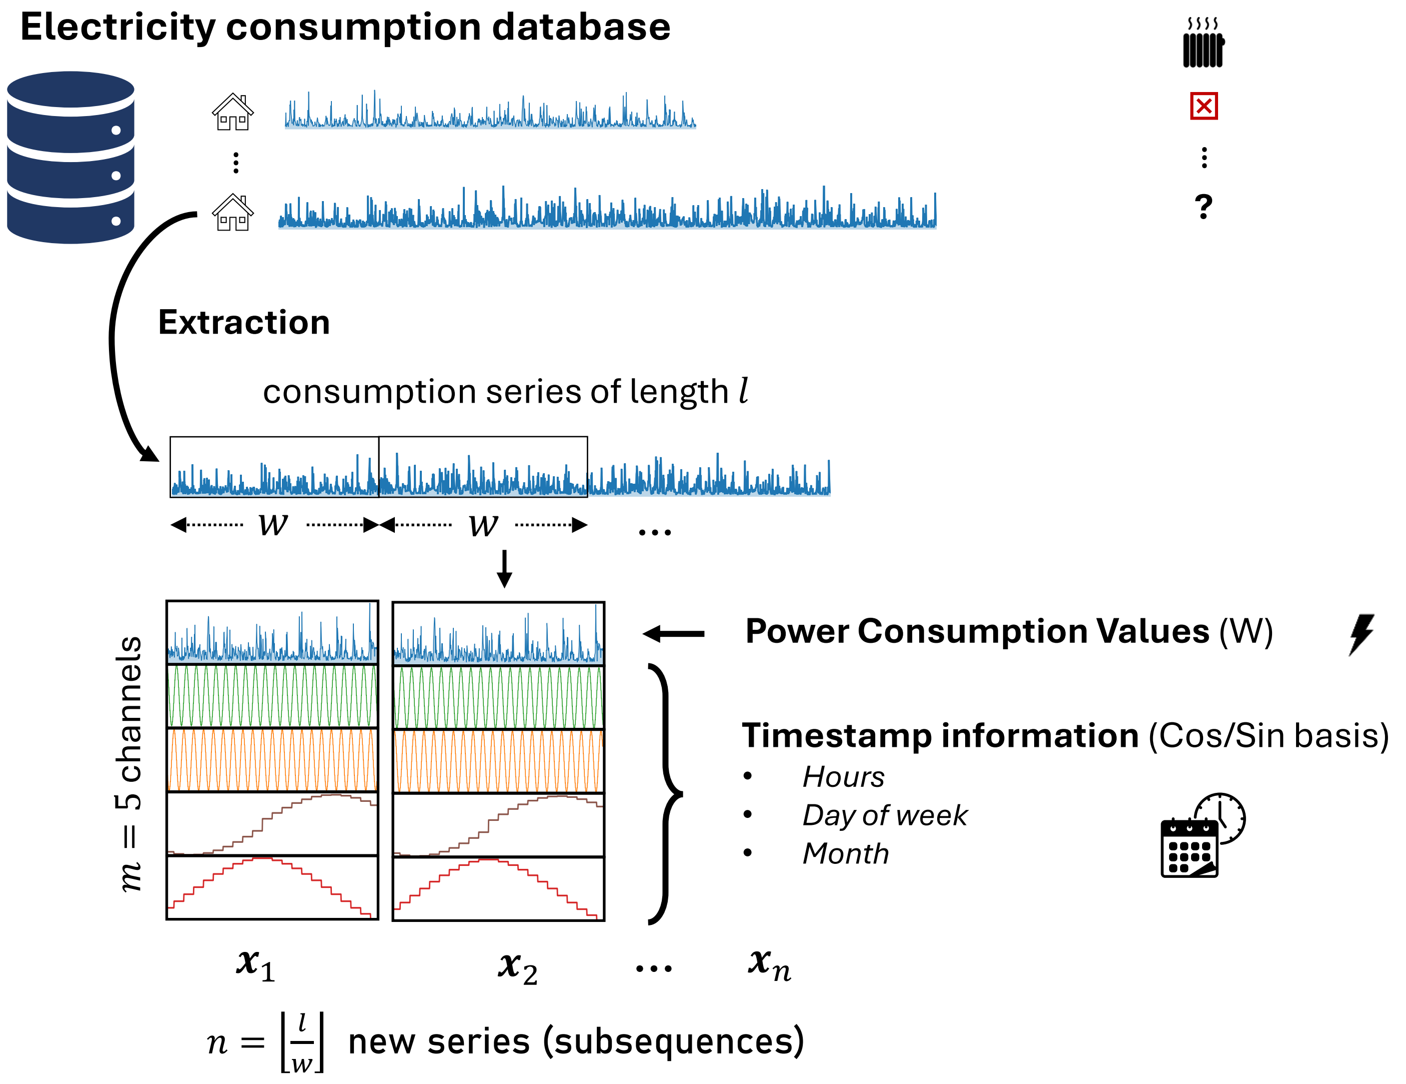

A. Petralia et al., ADF & TransApp: A Transformer-Based Framework for Appliance Detection Using Smart Meter Consumption Series, PVLDB, 2024.

Proposed Approach: ADF

The Appliance Detection Framework

1. **Slice** series into subsequences and **concatenate** them with timestamp-encoded information

2. TransApp predicts probabilities for **each subsequences**

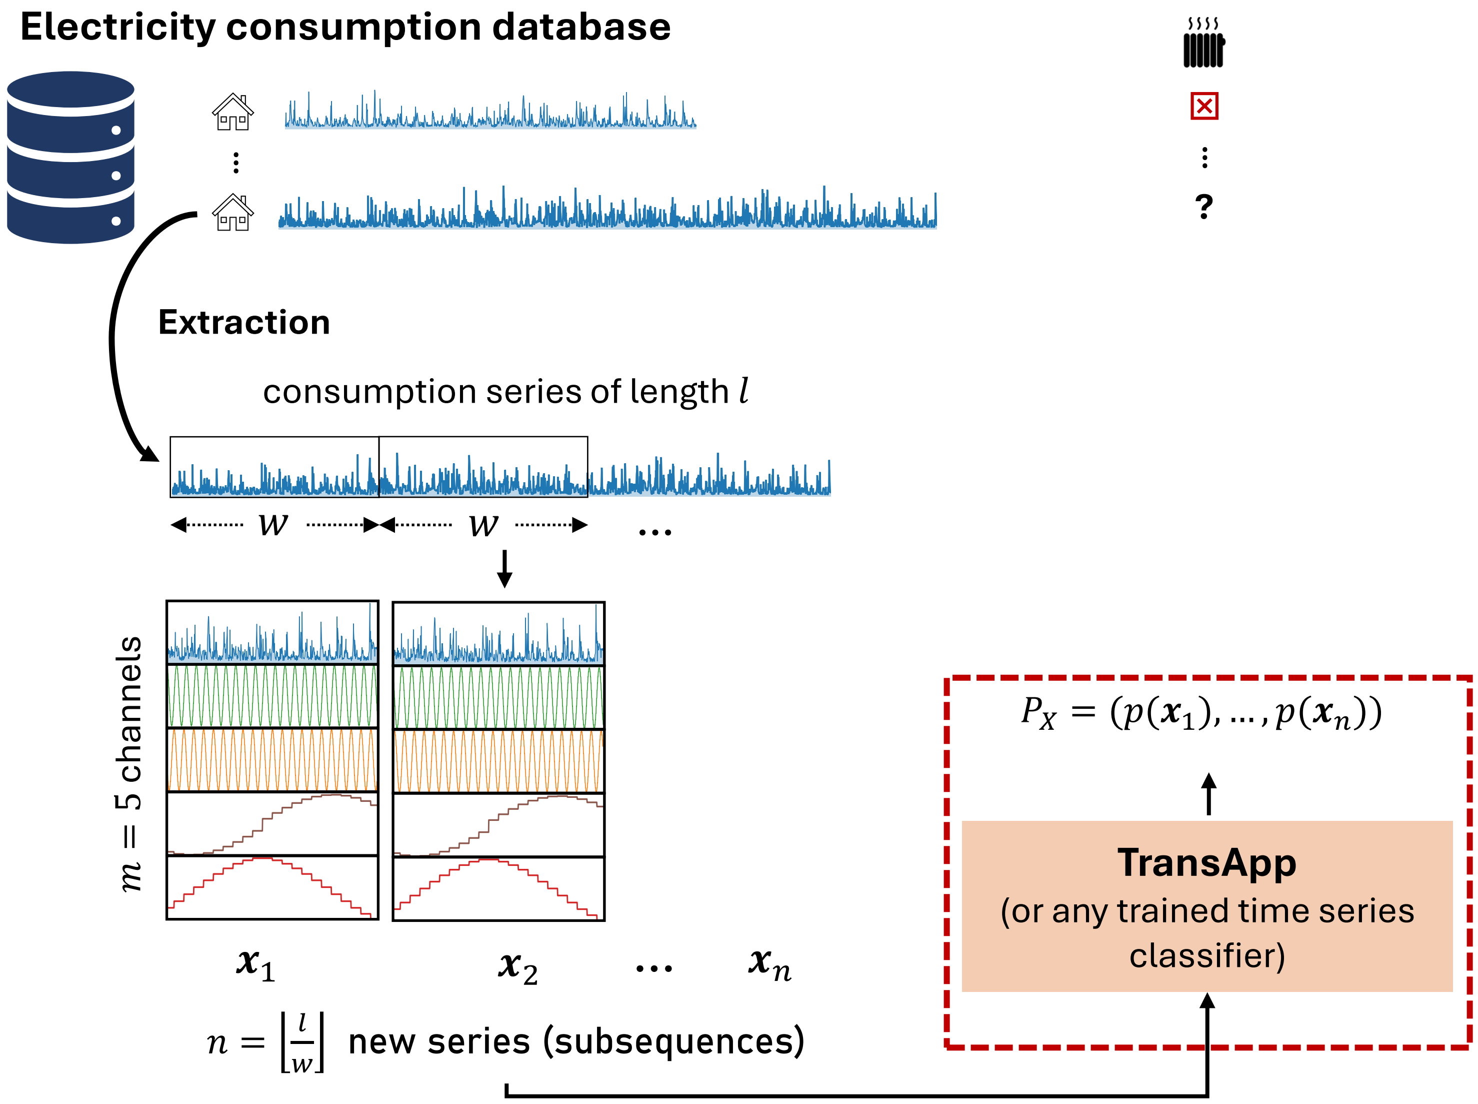

A. Petralia et al., ADF & TransApp: A Transformer-Based Framework for Appliance Detection Using Smart Meter Consumption Series, PVLDB, 2024.

Proposed Approach: ADF

The **A**ppliance **D**etection **F**ramework

1. **Slice** series into subsequences and **concatenate** them with timestamp-encoded information

2. TransApp predicts probabilities for **each subsequences**

3. **Merge predicted probabilities** by extracting best quantile

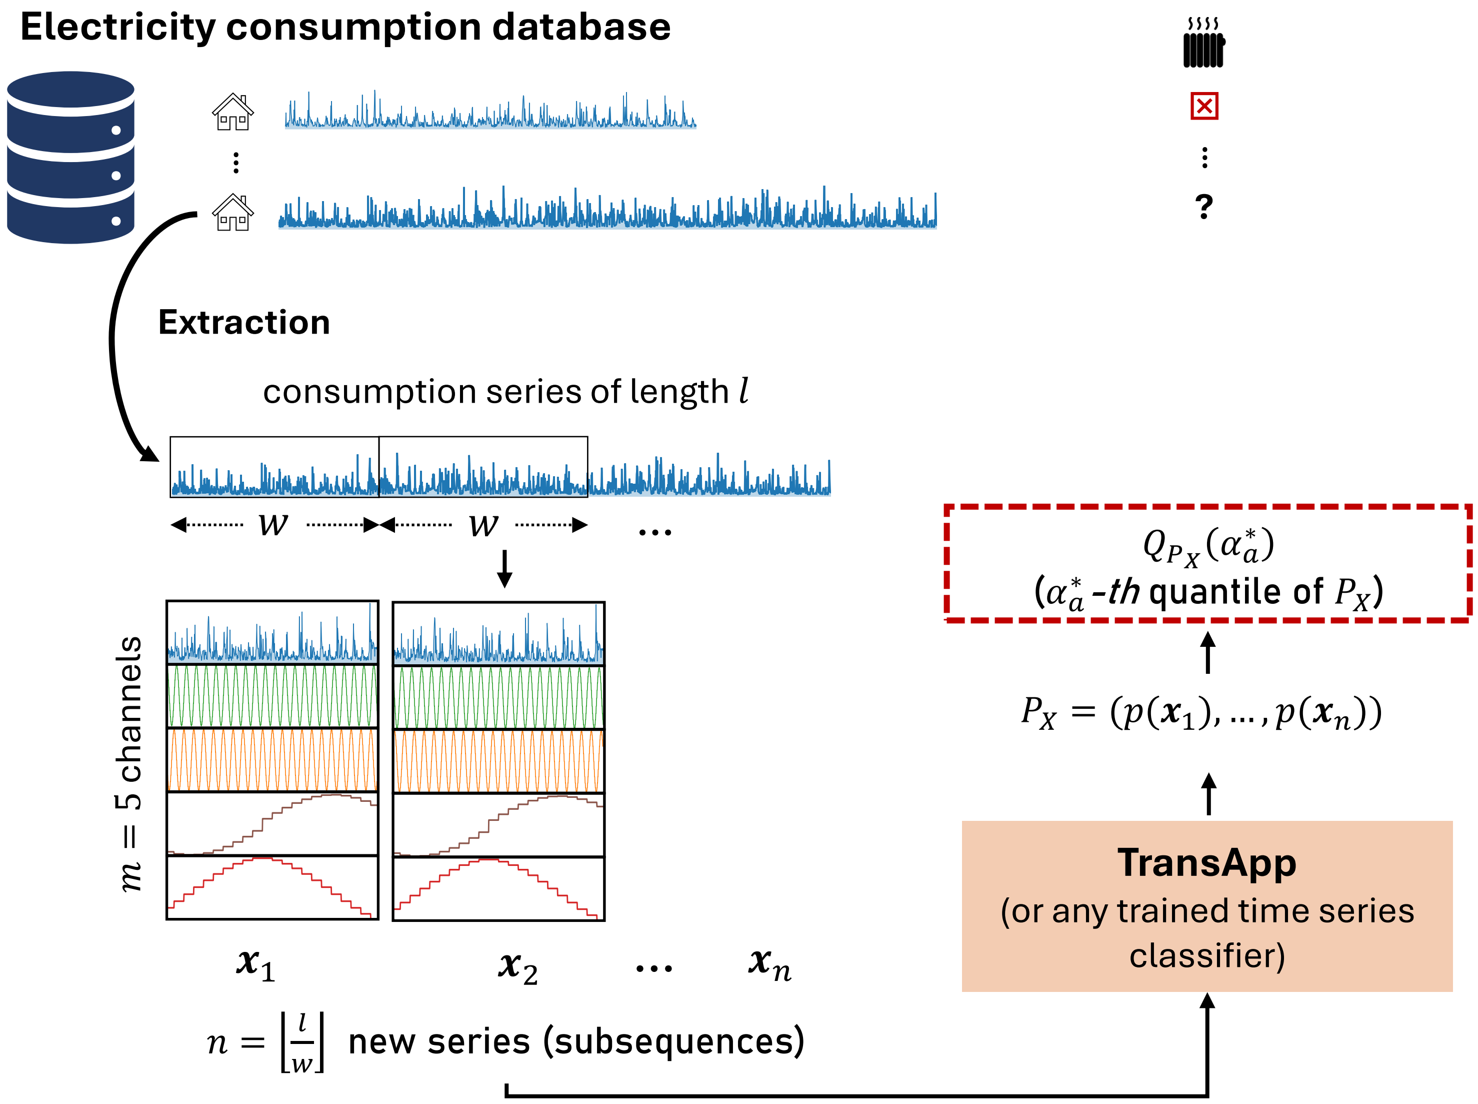

A. Petralia et al., ADF & TransApp: A Transformer-Based Framework for Appliance Detection Using Smart Meter Consumption Series, PVLDB, 2024.

Proposed Approach: ADF

The **A**ppliance **D**etection **F**ramework

1. **Slice** series into subsequences and **concatenate** them with timestamp-encoded information

2. TransApp predicts probabilities for **each subsequences**

3. **Merge predicted probabilities** by extracting best quantile

4. Determine the final **label prediction**

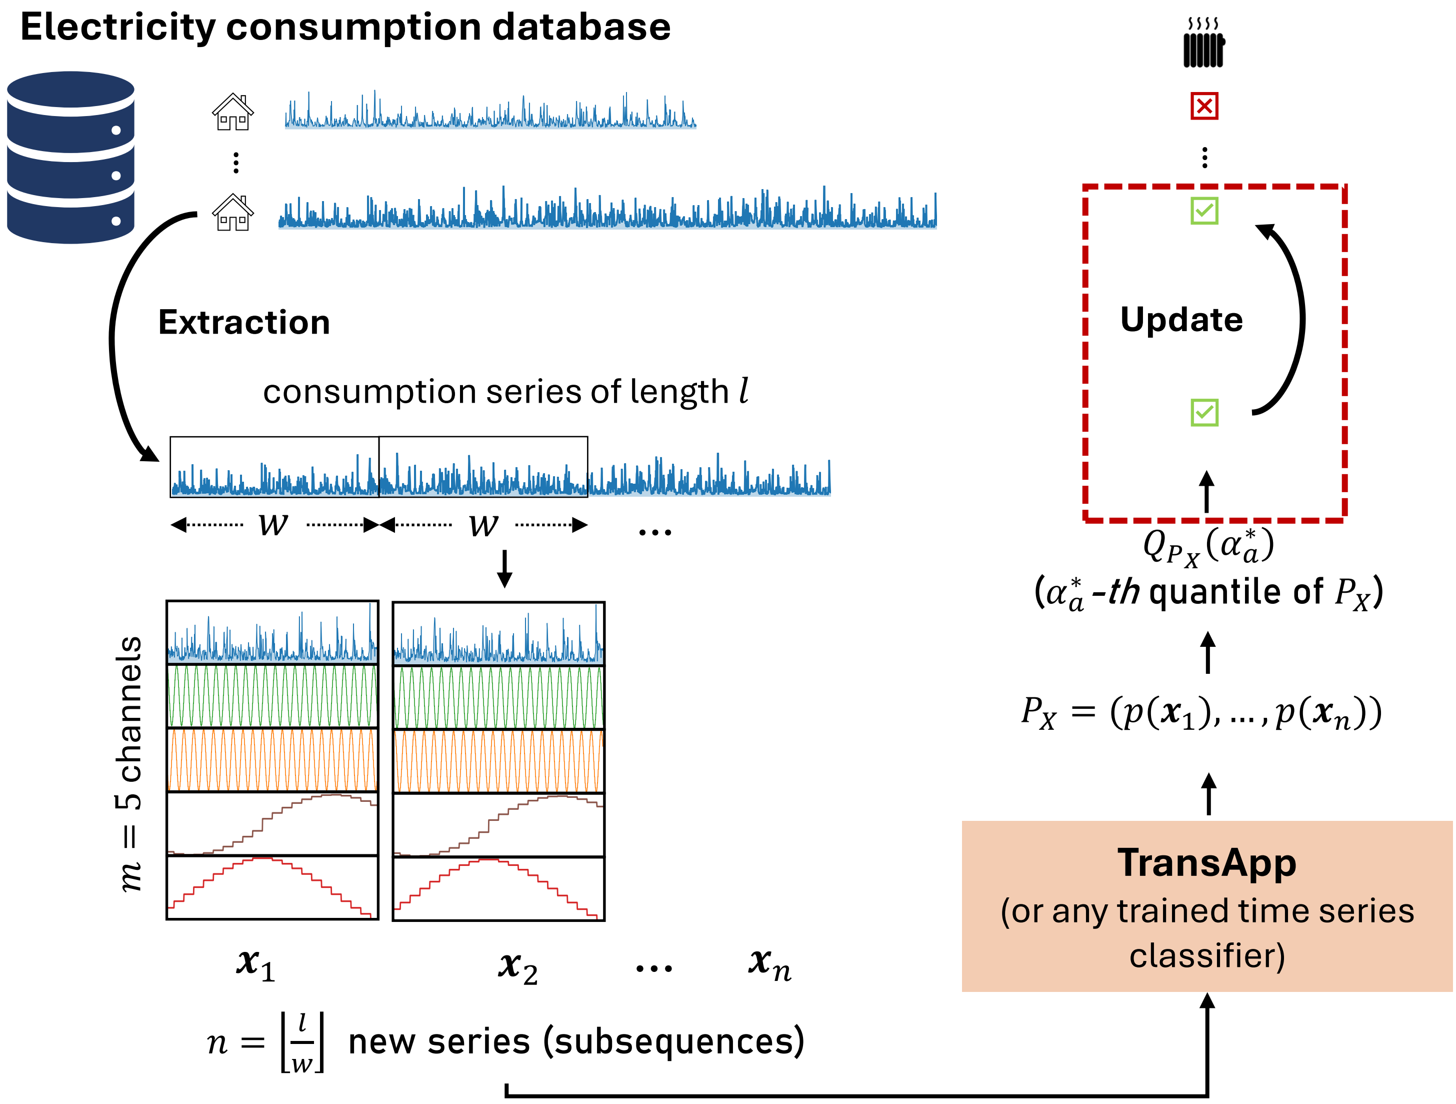

A. Petralia et al., ADF & TransApp: A Transformer-Based Framework for Appliance Detection Using Smart Meter Consumption Series, PVLDB, 2024.

TransApp: A simple deep-learning architecture

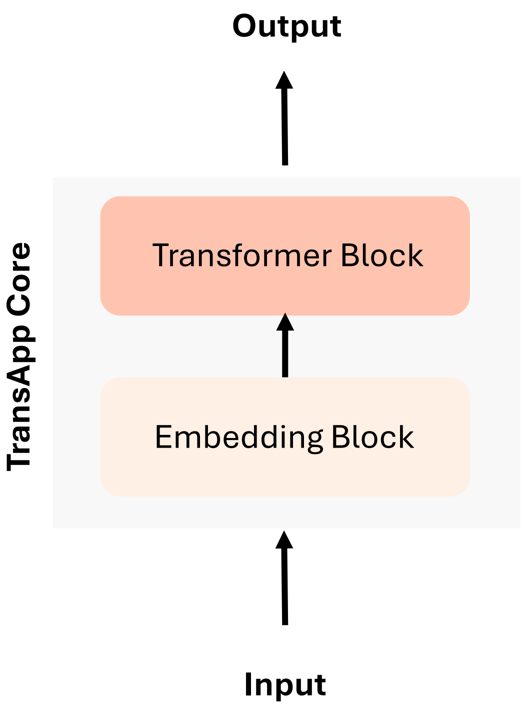

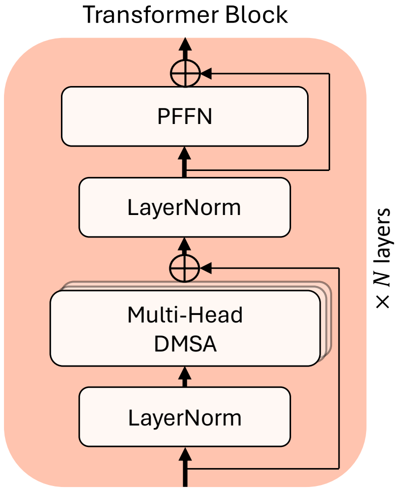

A **Transformer block** to learn useful representations and capture long-range dependencies

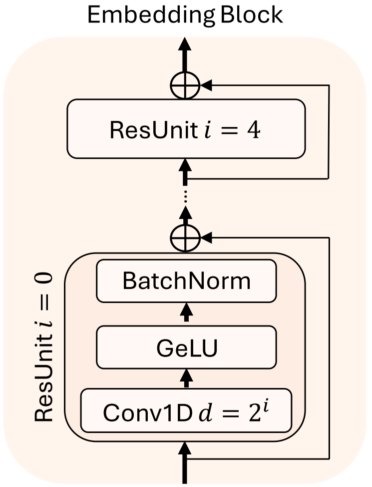

A strong **Convolutional Block** to extract localized features

A. Petralia et al., ADF & TransApp: A Transformer-Based Framework for Appliance Detection Using Smart Meter Consumption Series, PVLDB, 2024.

Proposed Approach: TransApp

TransApp: Two-steps training process

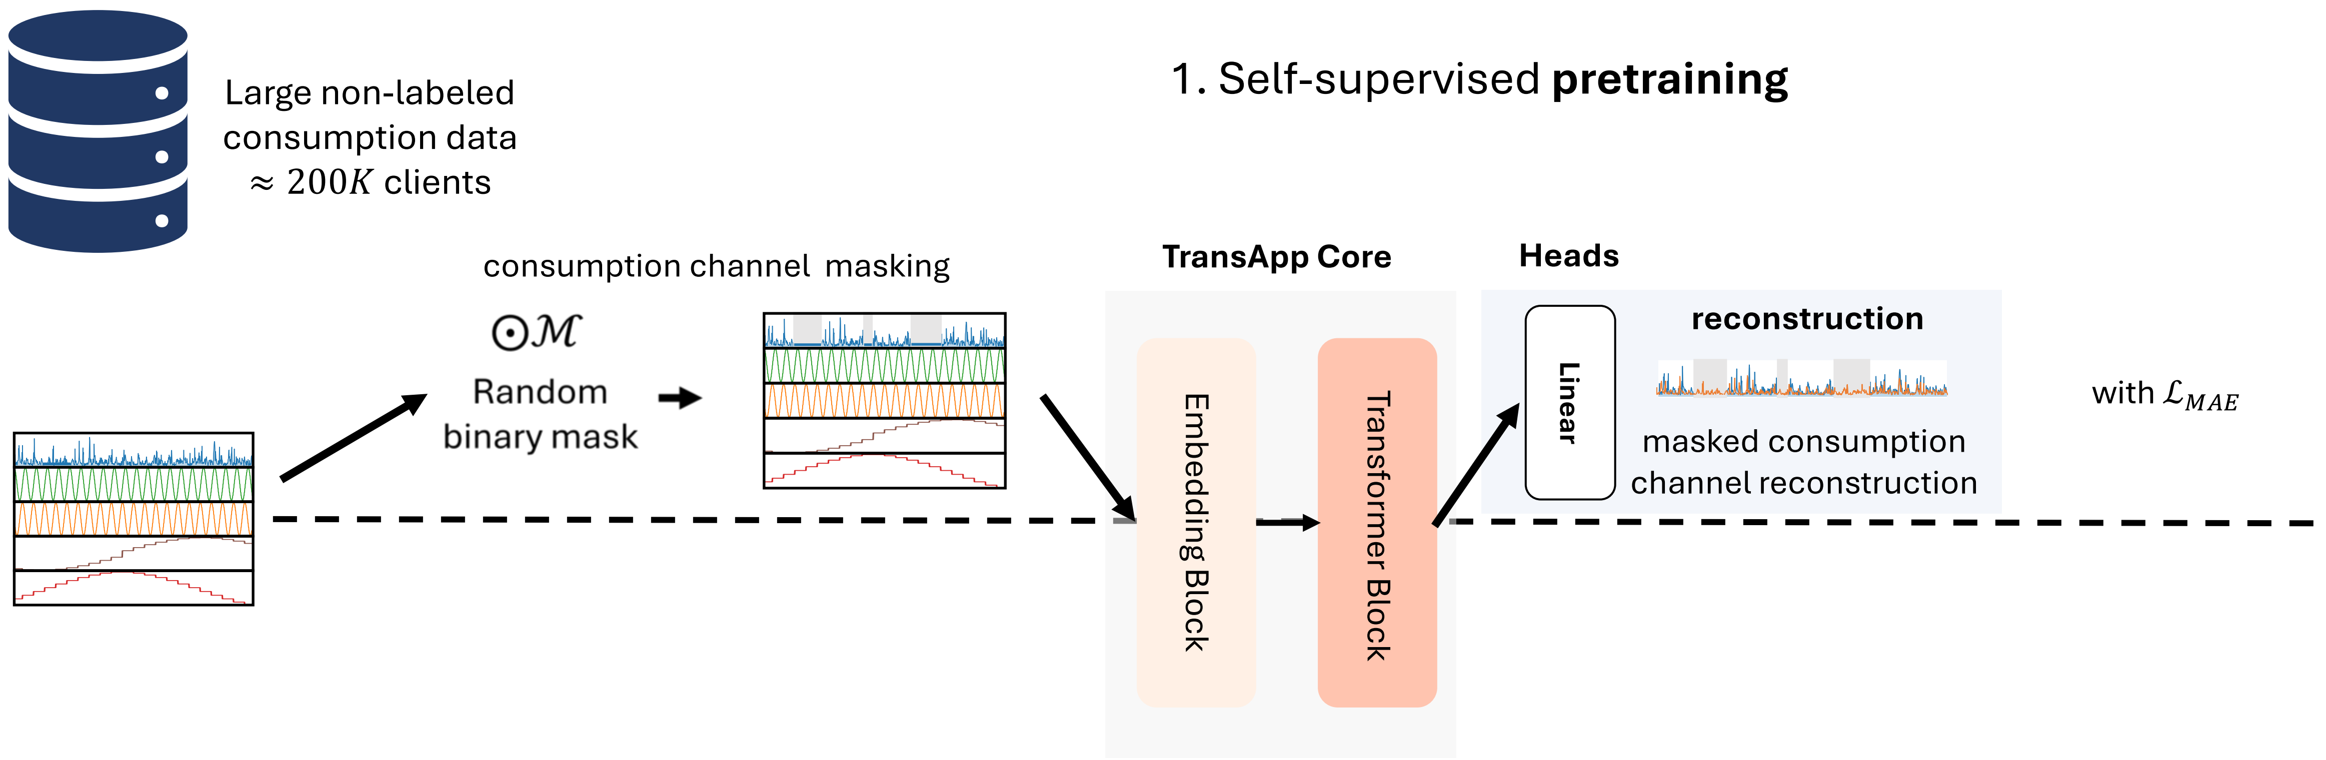

A. Petralia et al., ADF & TransApp: A Transformer-Based Framework for Appliance Detection Using Smart Meter Consumption Series, PVLDB, 2024.

TransApp: Two-steps training process

Large non-labeled consumption data
$\approx 200K$ clients

1. Self-supervised **pretraining**

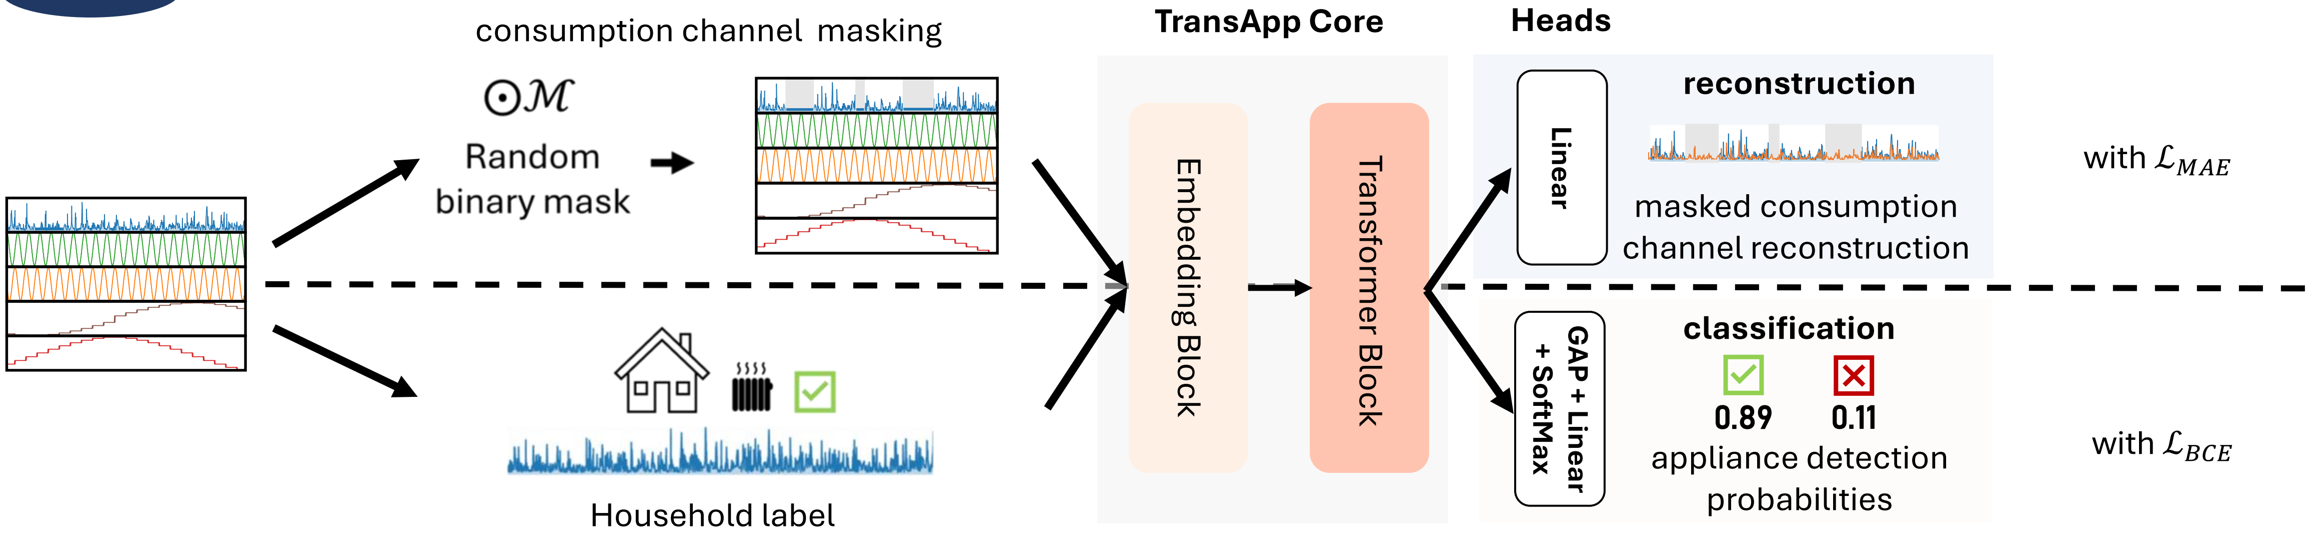

Small labeled dataset
$\approx 100 - 1000$ instances

2. Supervised **finetuning**

A. Petralia et al., ADF & TransApp: A Transformer-Based Framework for Appliance Detection Using Smart Meter Consumption Series, PVLDB, 2024.

Detection Accuracy Results

Average results over 7 different detection cases (EDF dataset)

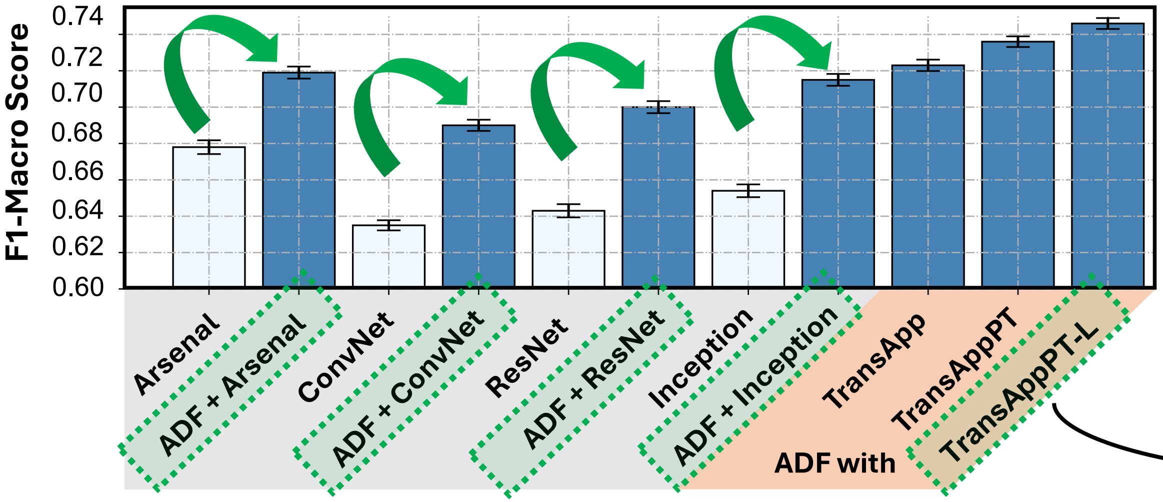

Pretrained on a **large unlabeled dataset** of **200K** households

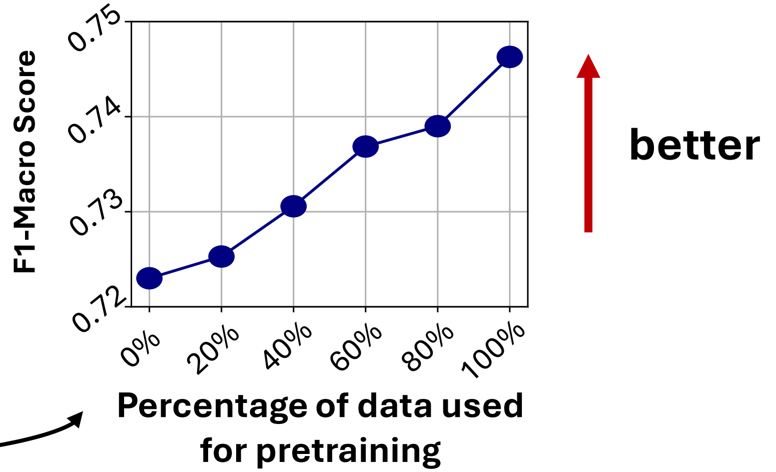

≈ ***+10% performance gain*** by using **SotA Time Series Classifiers** within the **ADF**

Best solution: TransAppPT-L, ***+8.5% increase*** compared to the **2nd -best solution**

Our solution **accurately detects** different appliances in real-world scenarios

Electric Vehicle **Heater** **AC** **Heat Pump** **Water Heater**

A. Petralia et al., ADF & TransApp: A Transformer-Based Framework for Appliance Detection Using Smart Meter Consumption Series, PVLDB, 2024.

ADF makes **TransApp scalable** to **long consumption series**

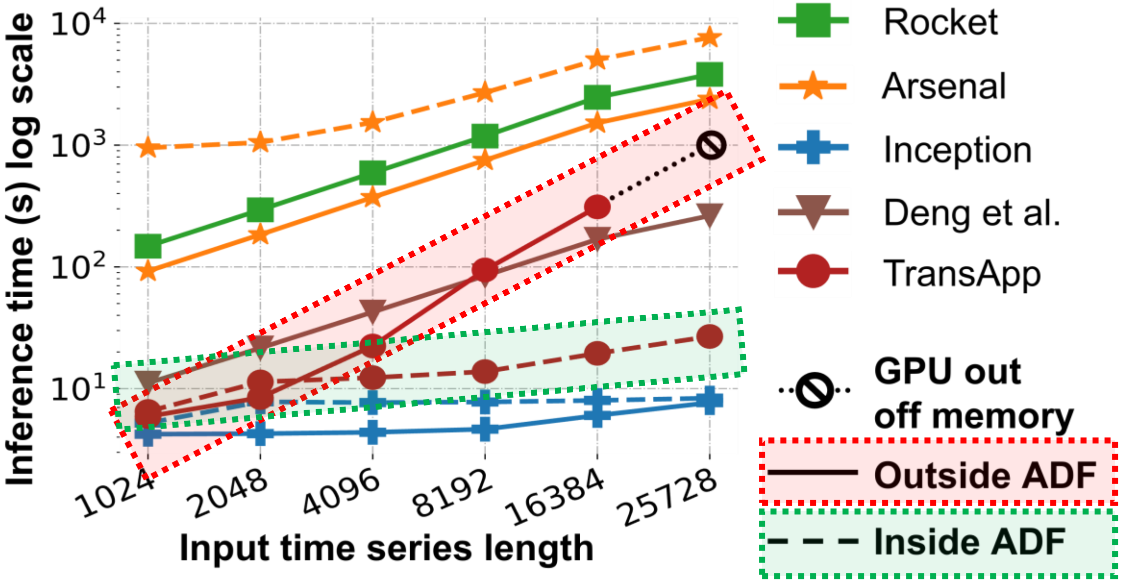

EDF database
4M clients recorded ≈1years

To run through the **entire EDF's client consumption database**

ADF & TransApp		**ADF & Arsenal** (2nd most accurate solution)
≈ 1days	**<<<**	**≈ 42days**

More than 40x faster

A. Petralia et al., ADF & TransApp: A Transformer-Based Framework for Appliance Detection Using Smart Meter Consumption Series, PVLDB, 2024.

Solution

*How to **enhance** the **accuracy** of **Appliance Detection Presence** in households using **very** low-frequency smart meter data?*

Challenges

1. **Nature** of **electricity consumption data**
   Very low frequency reading used by Smart Meters
   Long and **variable length** consumption series

2. **Data size**
   Few labeled data for training a solution
   Large amount of non-labeled data

Solutions

✓ The **A**ppliance **D**etection **F**ramework (**ADF**)

✓ **TransApp**: a deep-learning time series classifier

Solution

*How to **enhance** the **accuracy** of **Appliance Detection Presence** in households using **very** low-frequency smart meter data?*

Challenges

1. **Nature** of **electricity consumption data**
   - **Very low frequency** reading used by Smart Meters
   - **Long** and **variable length** consumption series

2. **Data size**
   - **Few** labeled data for training a solution
   - **Large amount** of non-labeled data

Solutions

✓ The **A**ppliance **D**etection **F**ramework (**ADF**)
- ➢ **Improve** classifiers detection accuracy
- ➢ Make classifiers less **sensitive** to the **entire series length**

✓ **TransApp**: a deep-learning time series classifier
- ➢ **Pretrained** on large amount of non-labeled data to improve its accuracy
- ➢ **Scalable** to large database of long series

I. Introduction

II. Contributions

1. Appliance Detection Presence in Consumers Household

2. Appliance Pattern Localization

3. Energy Disaggregation

III. Conclusions

1. Appliance Detection - *Time Series Classification*

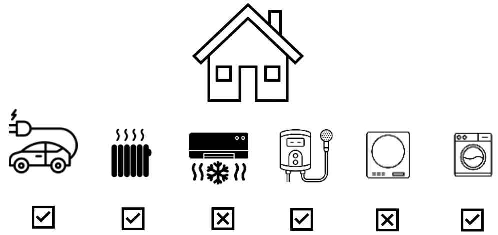

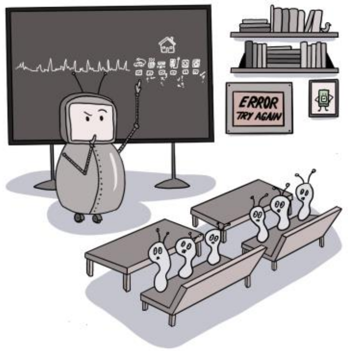

2. Appliance Pattern Localization – *Pattern Identification*

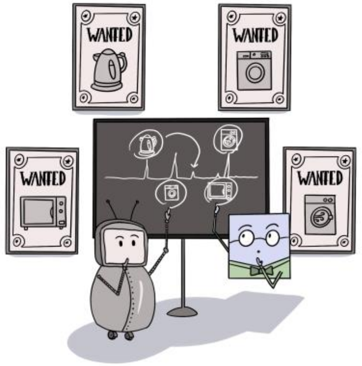

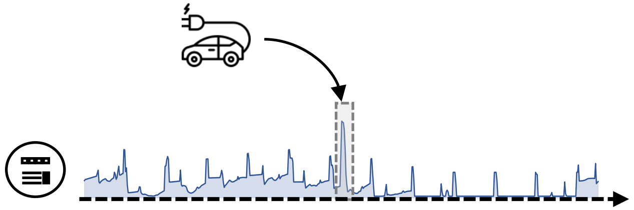

Appliance Pattern Localization

Recent **State-of-the-Art solutions** are based on deep-learning and a **Strongly Supervised Paradigm**

"Strong" Labeled data

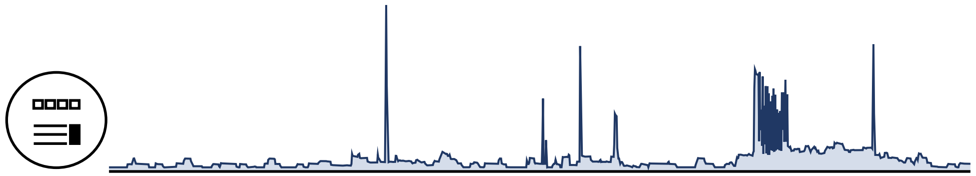

Aggregated Consumption Series

\+

Ground True Appliance **Status**

(one label for each timestep)

Ground true appliance signals are **extremely expensive** to **collect**

500 to 1500$/house

Large amount of « weak » labeled data

(one label per household)

Max 50$/house

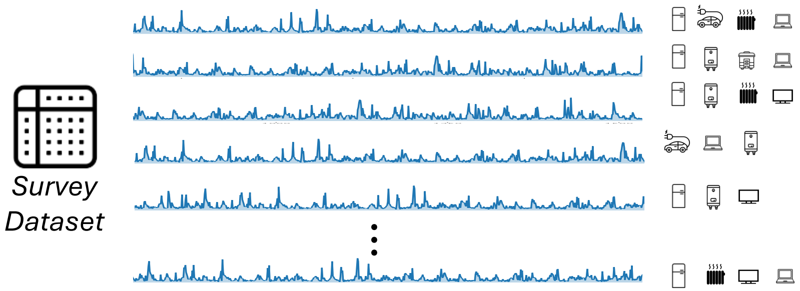

Can we tackle the ***Appliance Pattern Localization*** *problem using* ***minimal supervision****?*

Challenges

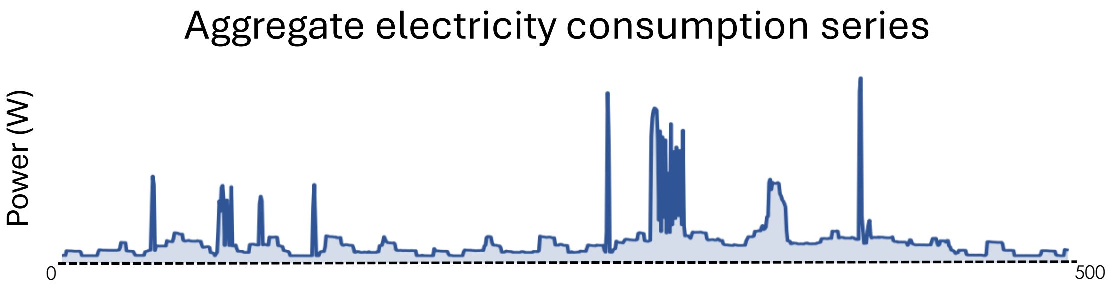

*Can we tackle the **Appliance Pattern Localization** problem using **minimal supervision**?*

Challenges

Aggregate electricity consumption series

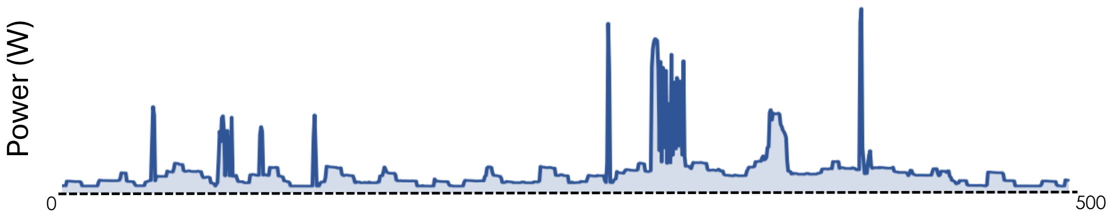

One label per time series and per appliance

(Problem 1)
Appliance Detection

Washer:	Yes
Microwave:	No
Dishwasher:	No
Kettle:	Yes

Problem

*Can we tackle the **Appliance Pattern Localization** problem using **minimal supervision**?*

Challenges

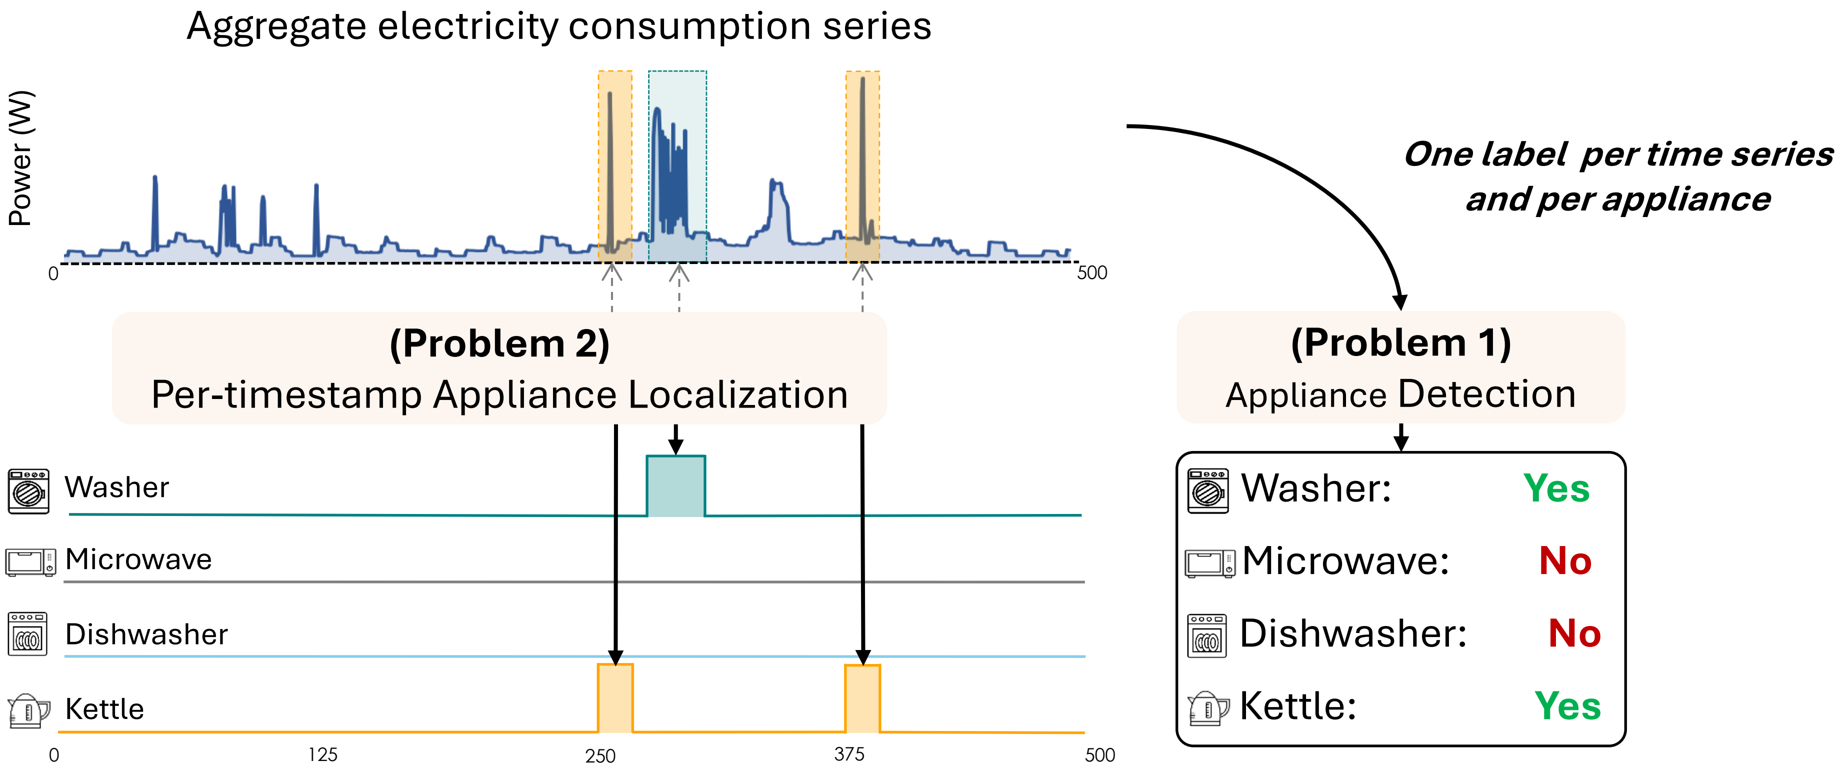

*Can we tackle the **Appliance Pattern Localization** problem using **minimal supervision**?*

Challenges

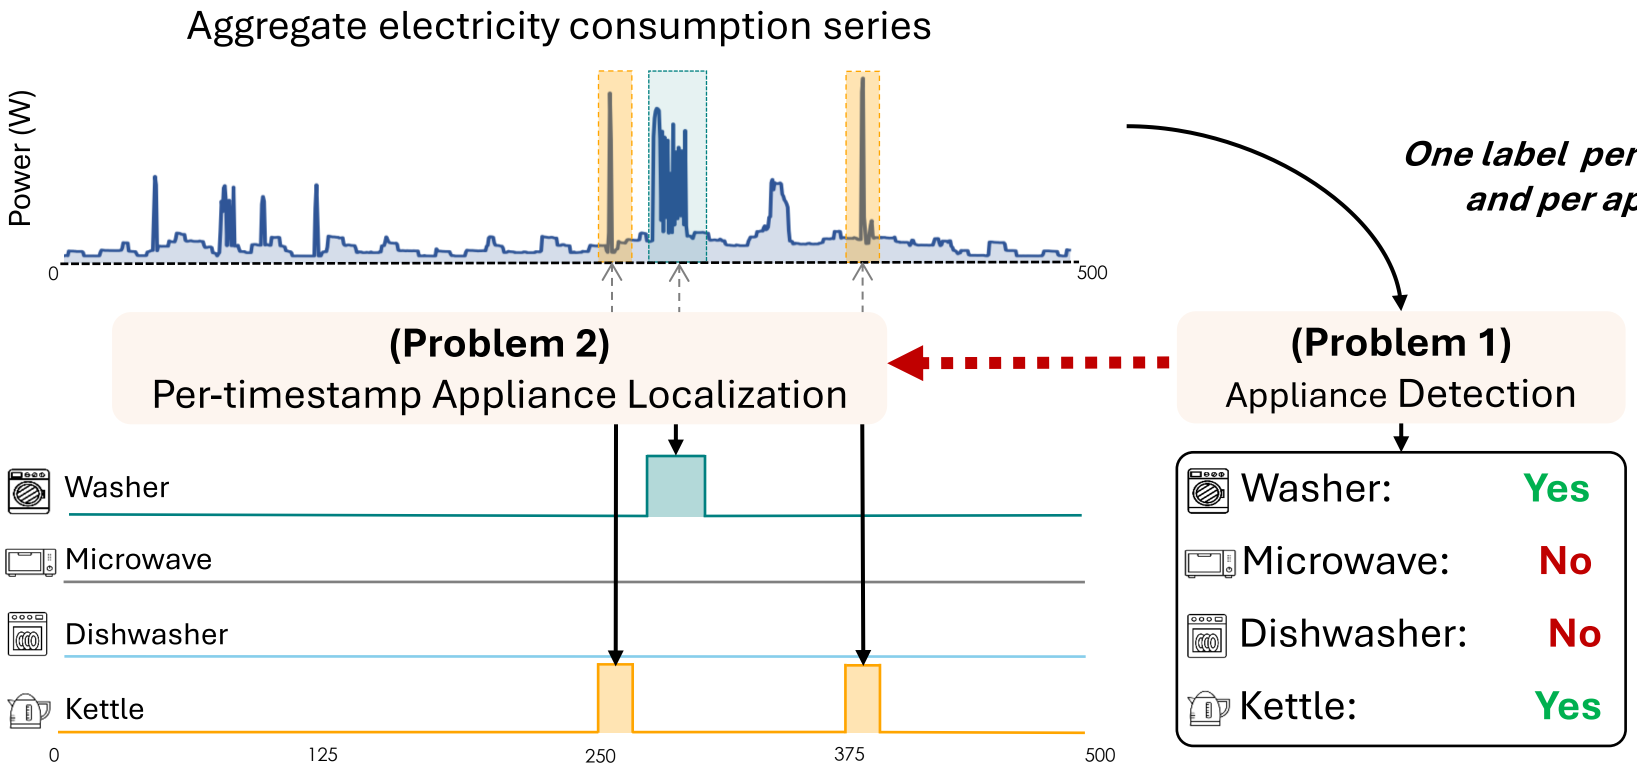

Solving Problem 2 from Problem 1

Problem

*Can we tackle the **Appliance Pattern Localization** problem using **minimal supervision**?*

Challenge

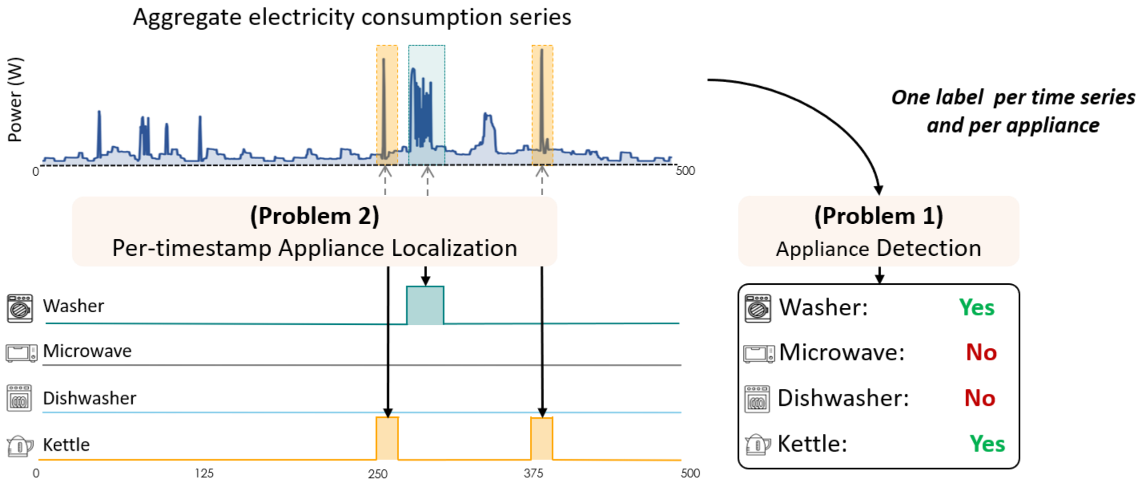

Solution

✓ **CamAL** (**C**lass **A**ctivation **M**ap based **A**ppliance **L**ocalization)

Explainable AI - Class **A**ctivation **M**ap (CAM)

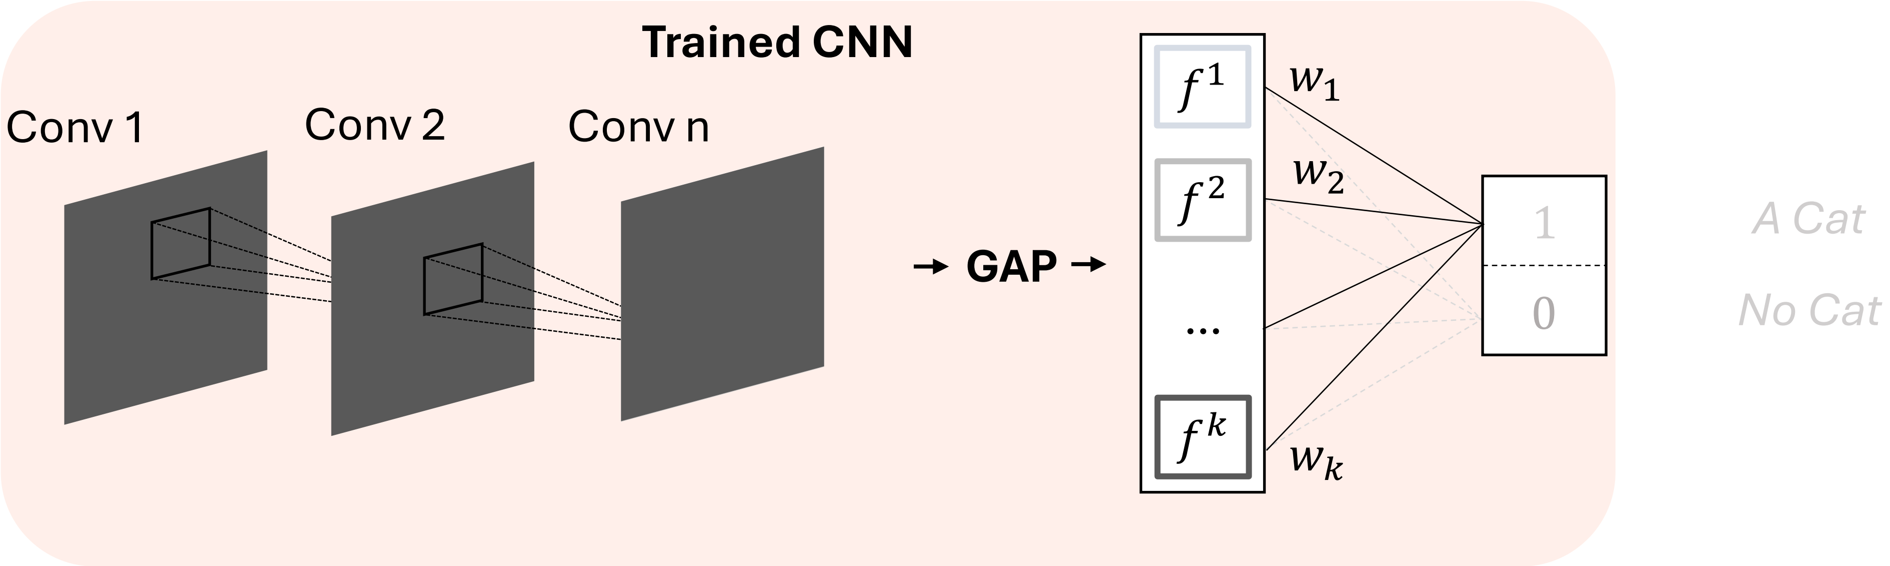

Explainable AI - Class **A**ctivation **M**ap (CAM)

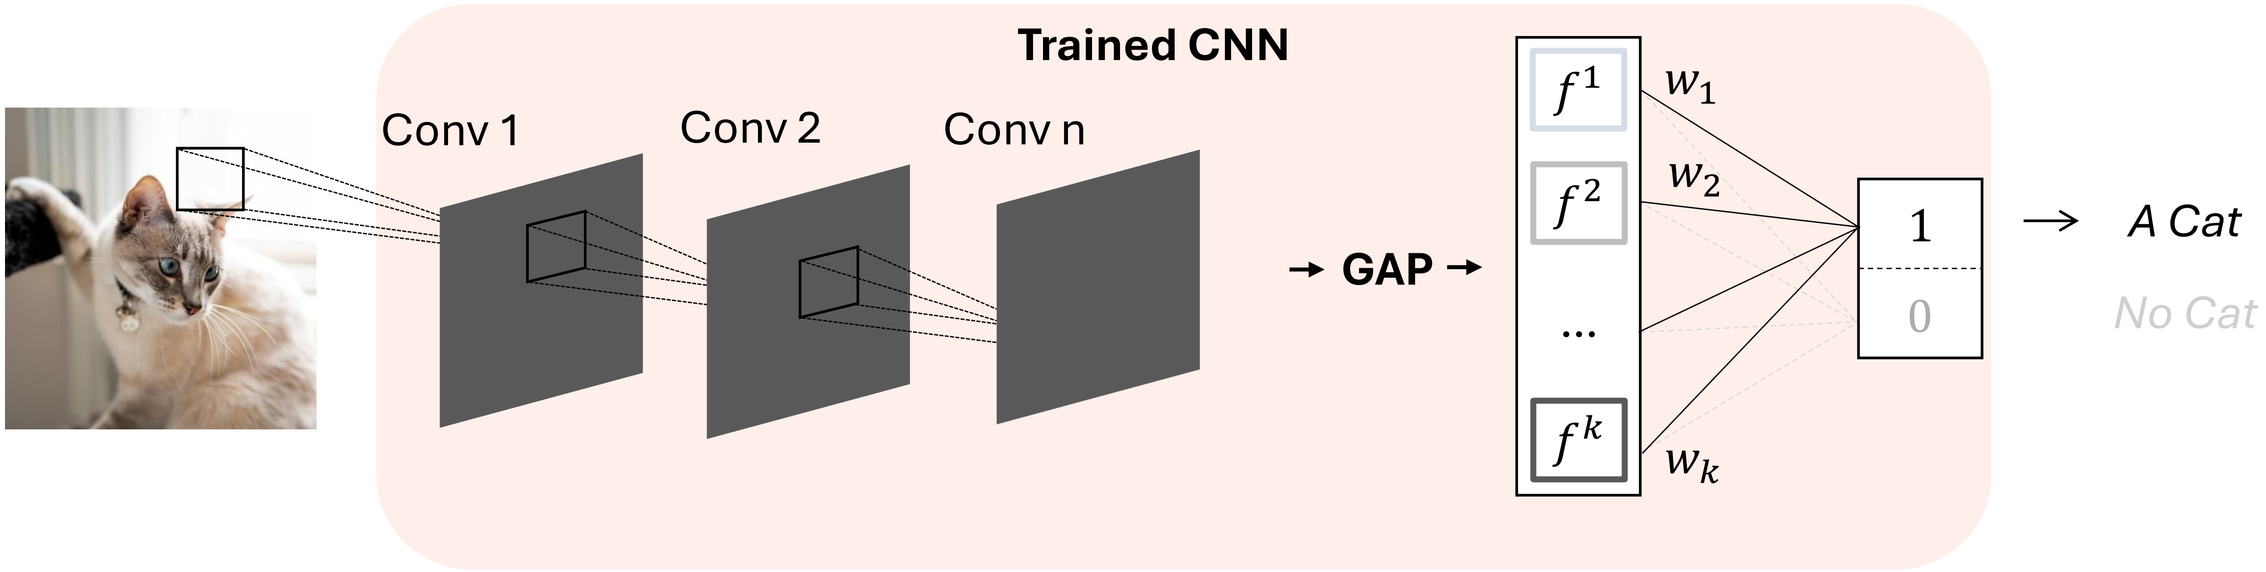

Explainable AI - Class **A**ctivation **M**ap (CAM)

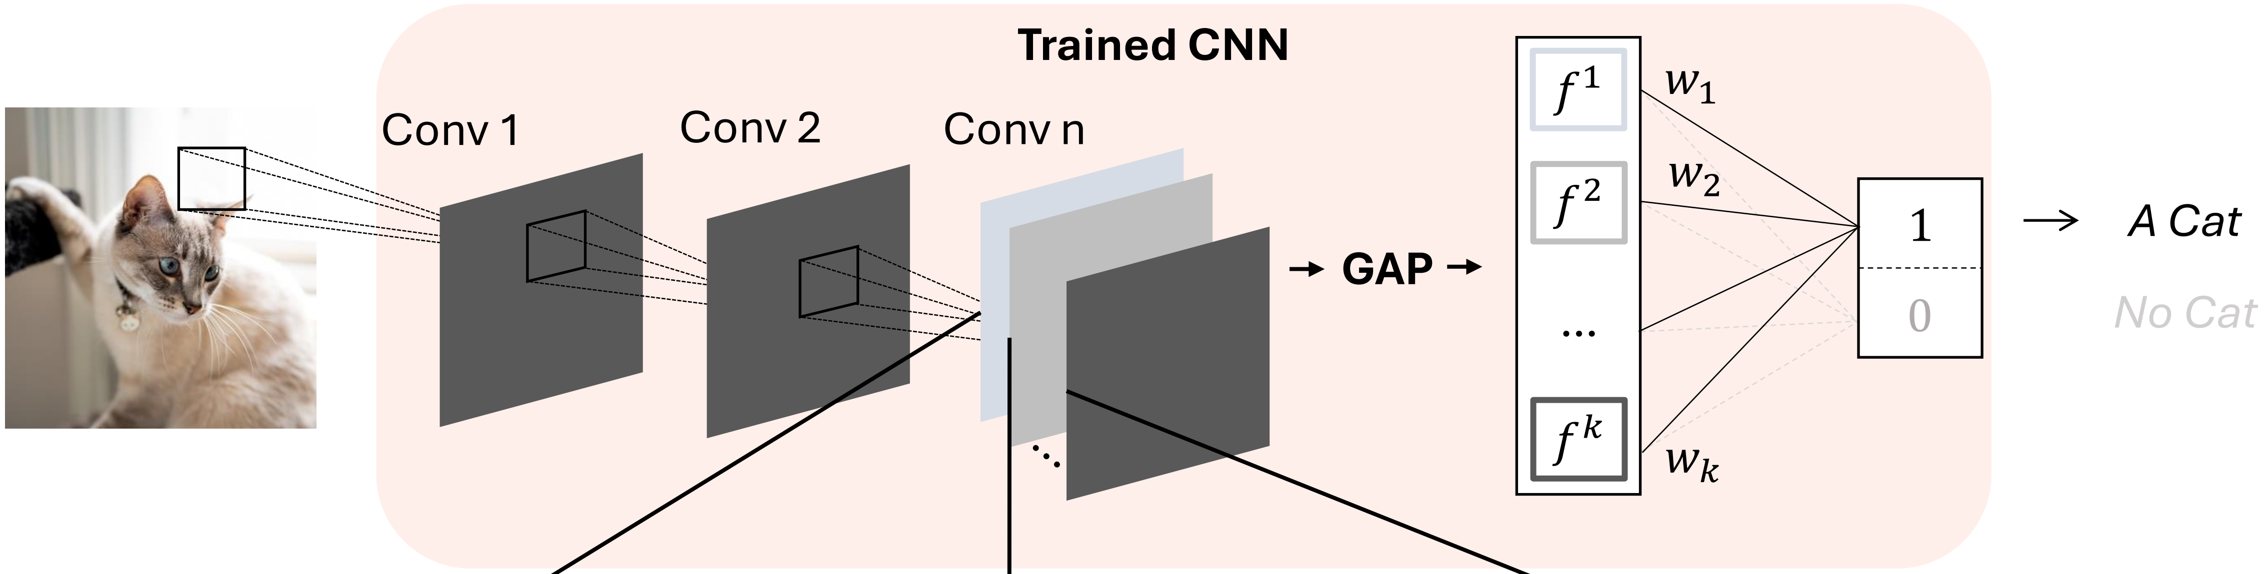

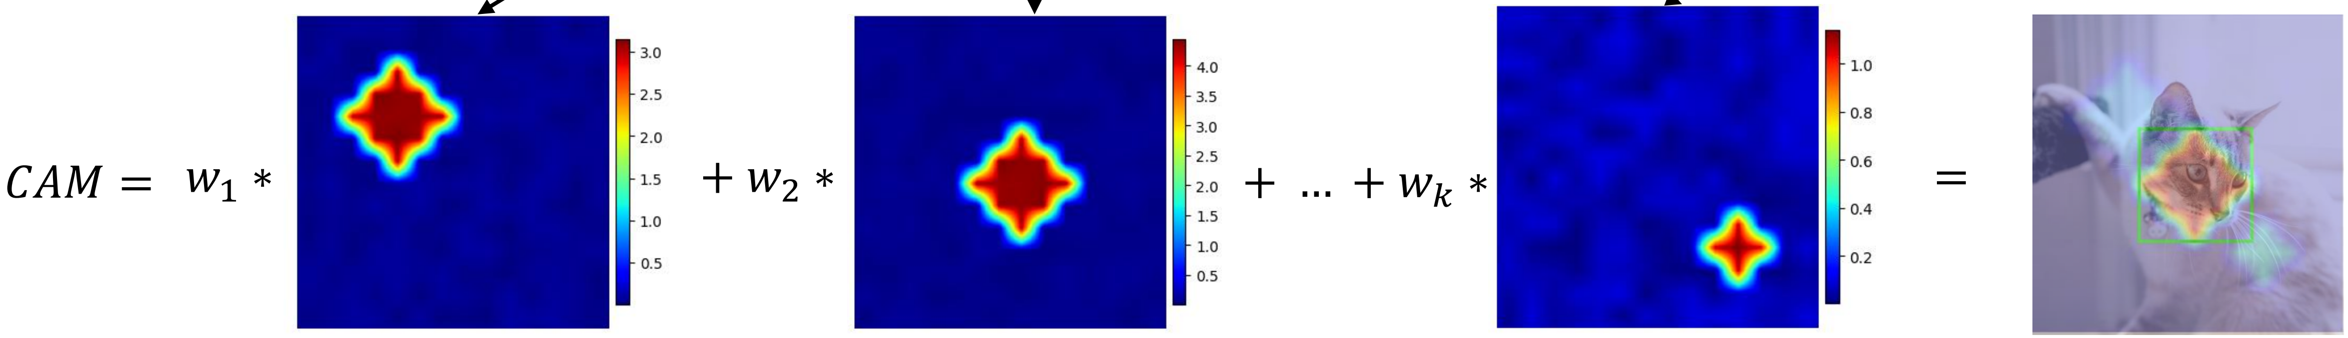

Proposed Approach: CAM for Appliance Pattern Localization?

CNNs (ResNet, Inception) perform well on the Appliance Detection task

Is CAM a « Free Lunch » for **Appliance-Pattern Localization** ?

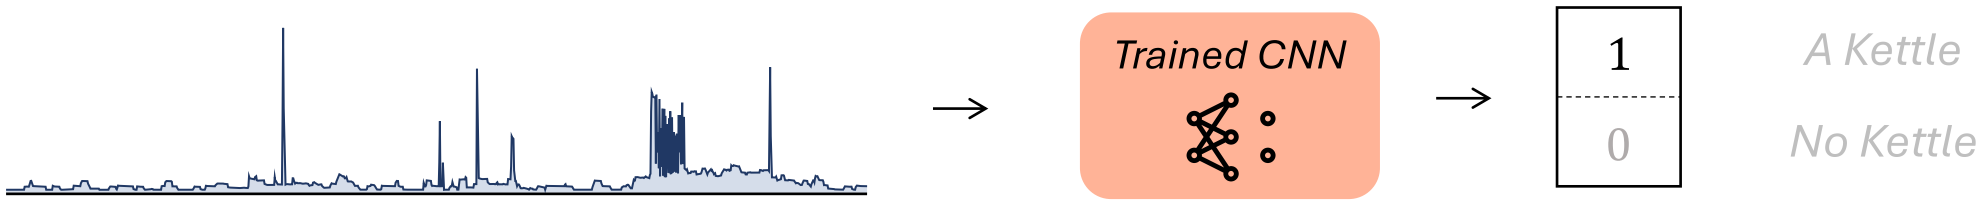

Proposed Approach: CAM for Appliance Pattern Localization?

CNNs (ResNet, Inception) perform well on the Appliance Detection task

Is CAM a « Free Lunch » for **Appliance-Pattern Localization** ?

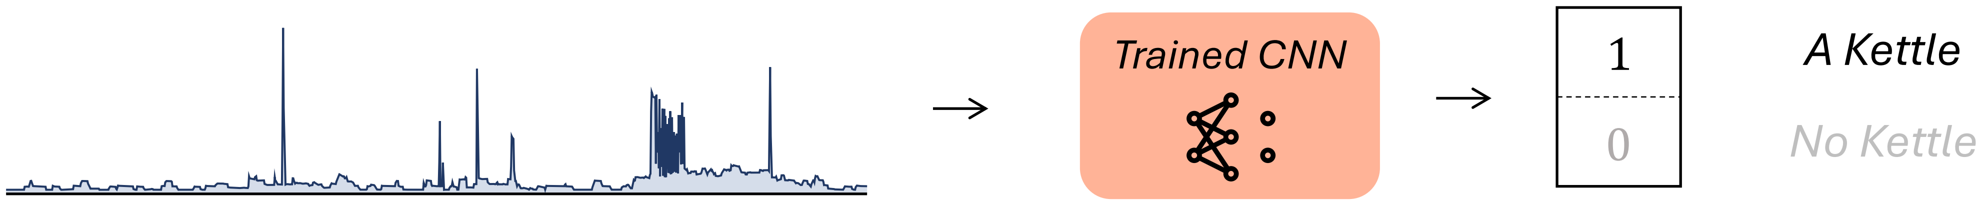

Proposed Approach: CAM for Appliance Pattern Localization?

CNNs (ResNet, Inception) perform well on the Appliance Detection task

Is CAM a « Free Lunch » for **Appliance-Pattern Localization** ?

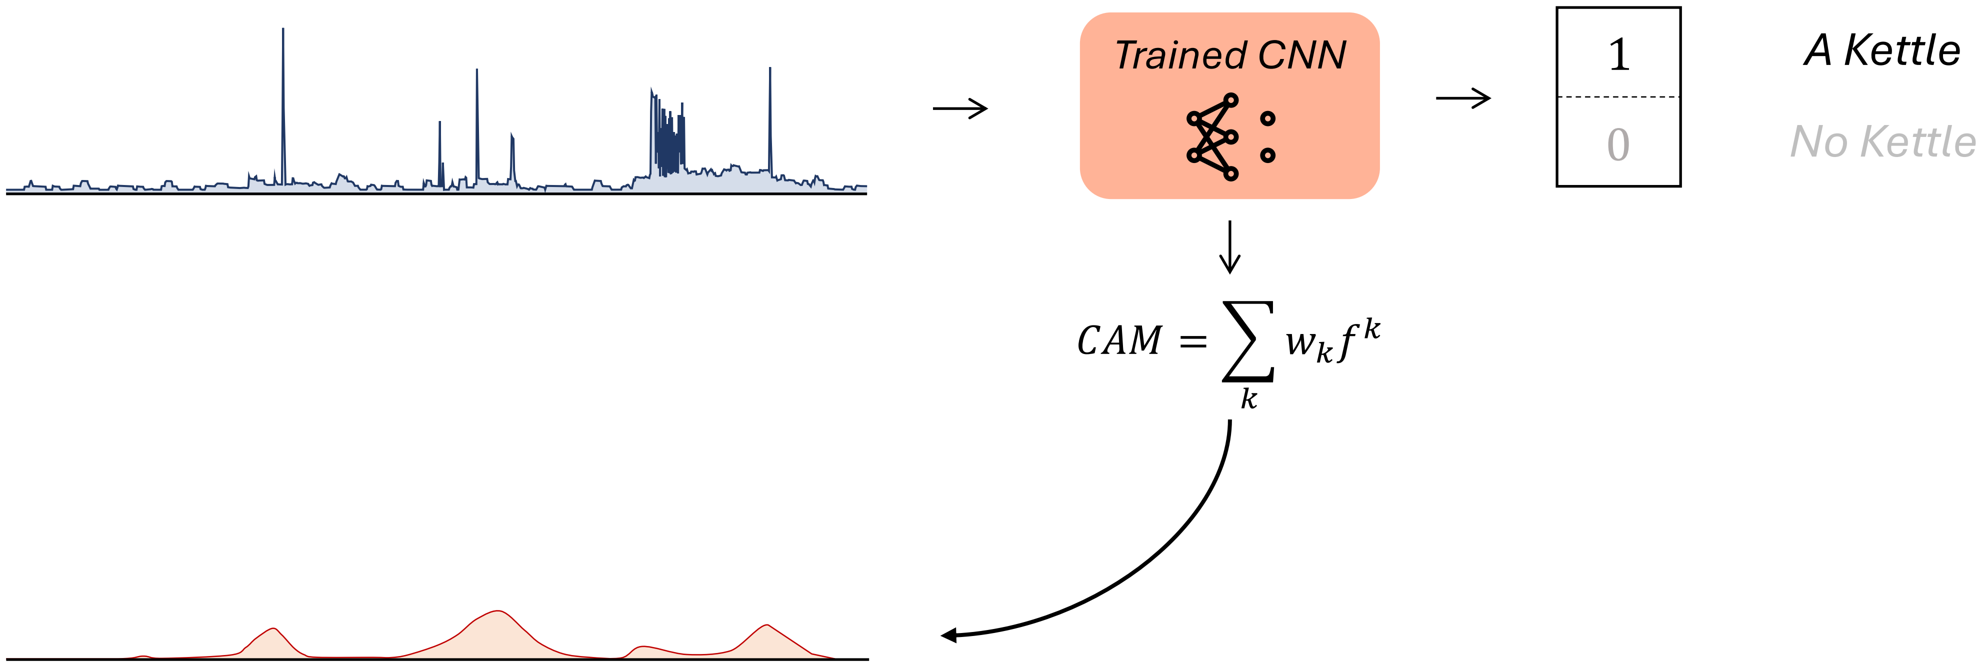

Class Activation Map

Proposed Approach: CAM for Appliance Pattern Localization?

CNNs (ResNet, Inception) perform well on the Appliance Detection task

Is CAM a « Free Lunch » for **Appliance-Pattern Localization** ?

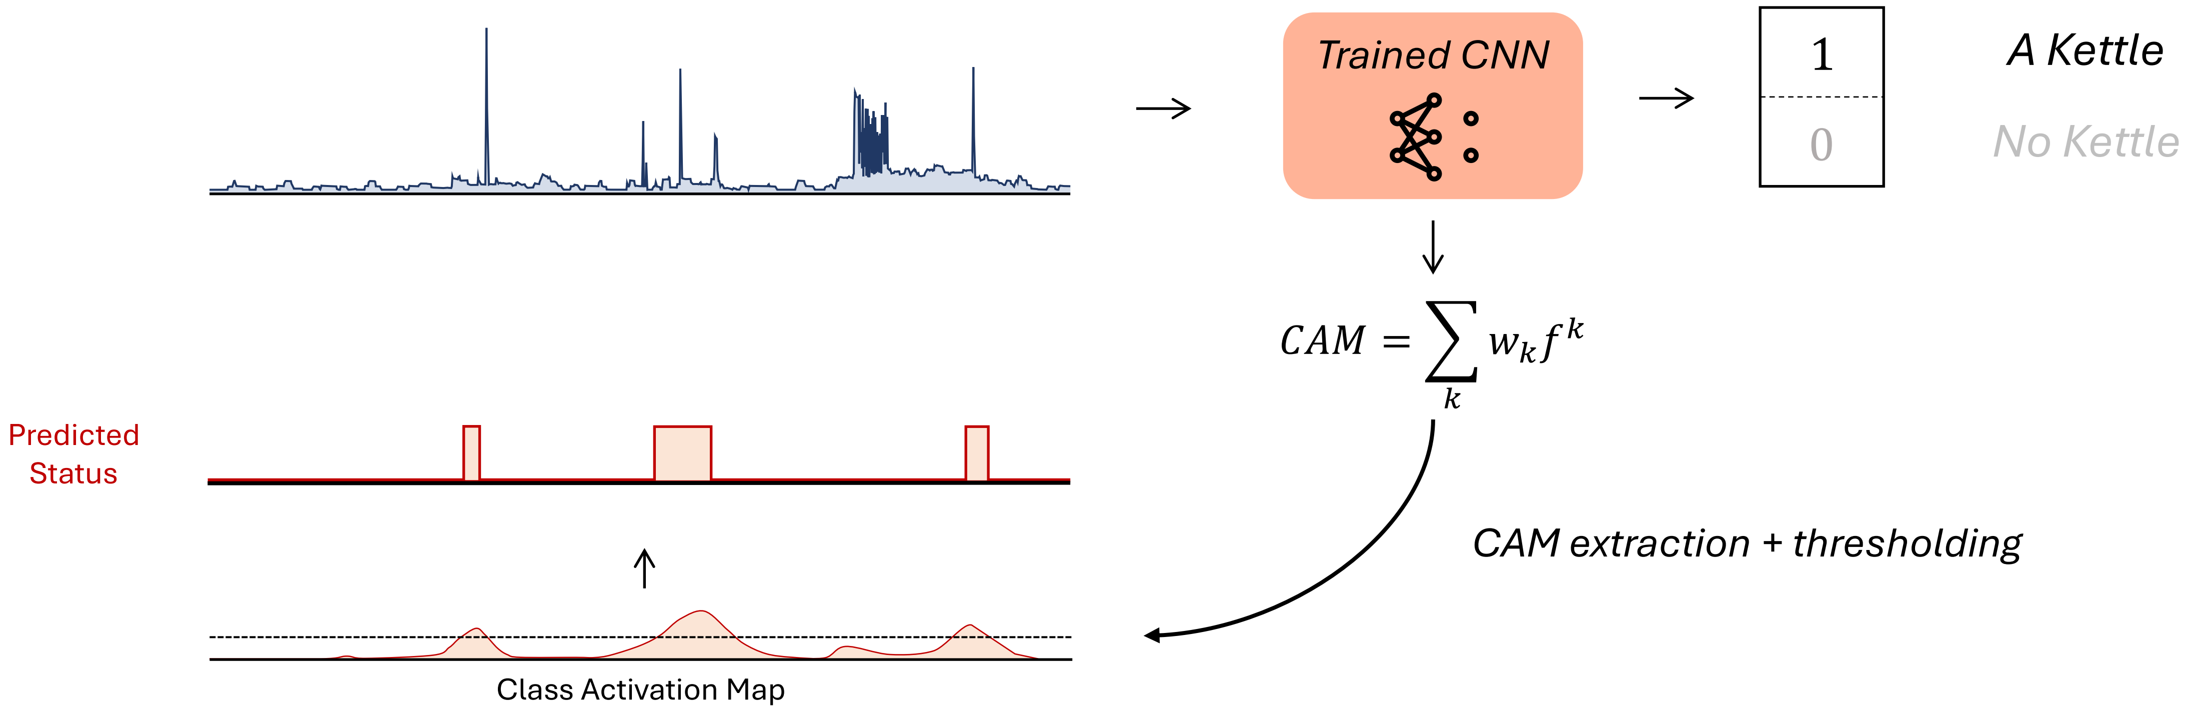

Proposed Approach: CAM for Appliance Pattern Localization?

CNNs (ResNet, Inception) perform well on the Appliance Detection task

Is CAM a « Free Lunch » for **Appliance-Pattern Localization** ?

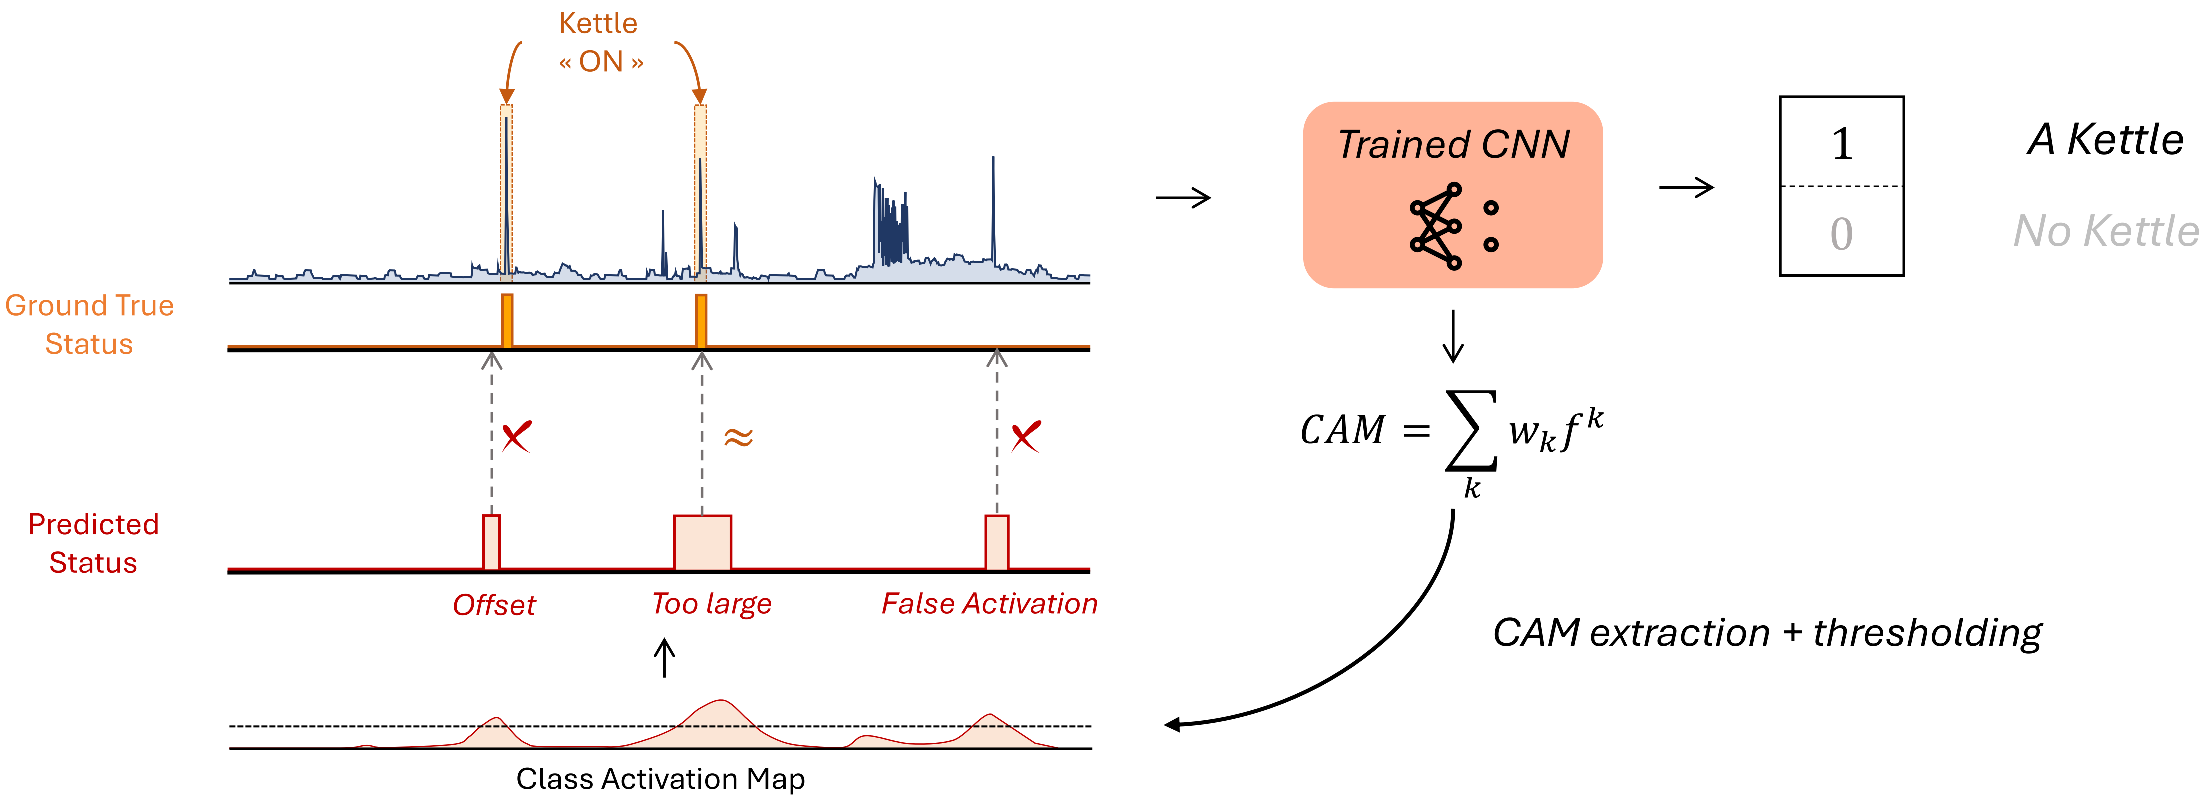

Proposed Approach: CAM for Appliance Pattern Localization?

CNNs (ResNet, Inception) perform well on the Appliance Detection task

Is CAM a « Free Lunch » for **Appliance-Pattern Localization** ? → ***Not that simple…***

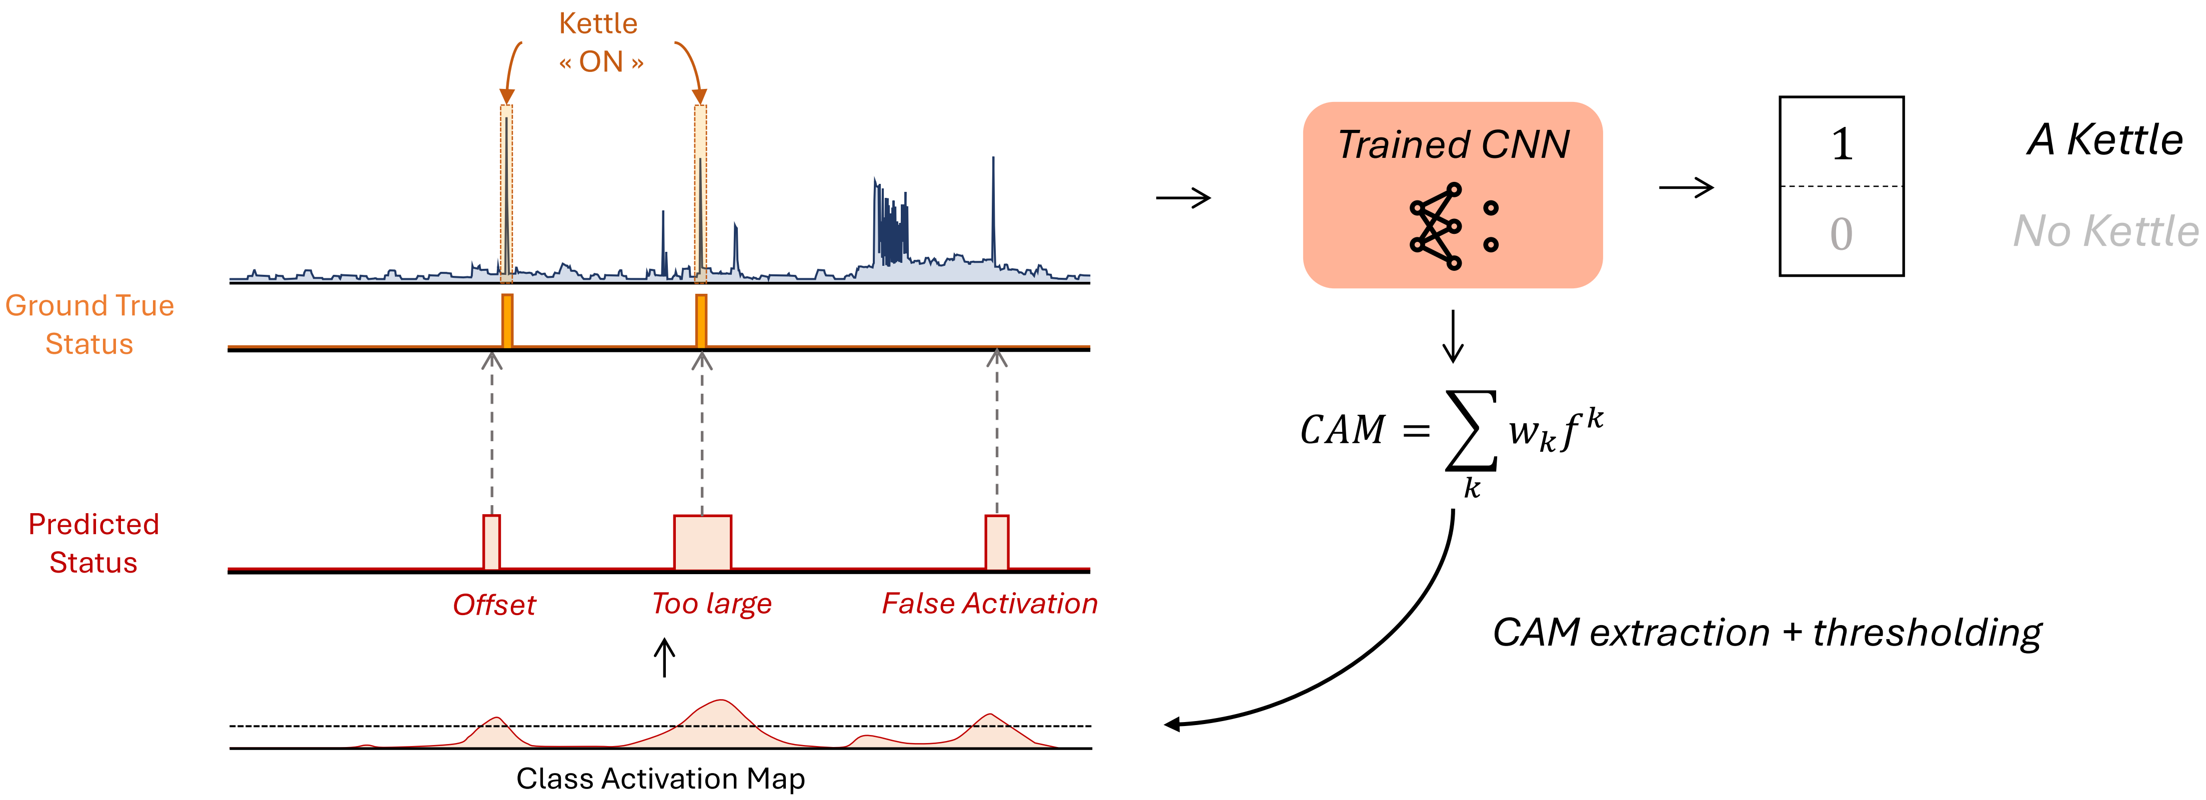

$$CAM = \sum_k w_k f^k$$

Improving **CAM** for **A**ppliance **L**ocalization

A ResNet ensemble with varying kernel sizes

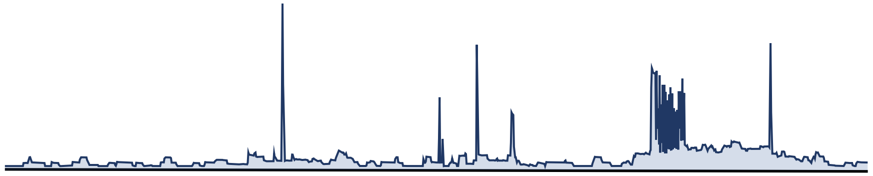

→

Trained ResNet Ensemble

ResNet 1 | *ResNet 2* | ··· | *ResNet n*

A. Petralia et al., *Few Labels are all you need: A Weakly Supervised Framework for Appliance Localization in Smart-Meter Series*, ICDE, 2025.

Improving **CAM** for **A**ppliance **L**ocalization

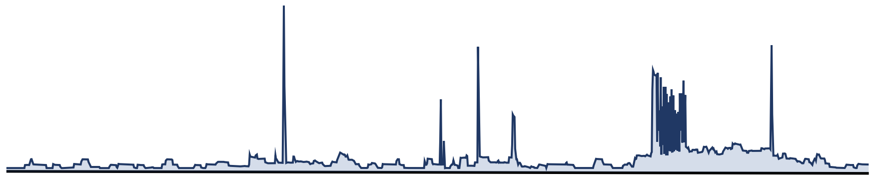

→

A **ResNet ensemble** with **varying kernel sizes**

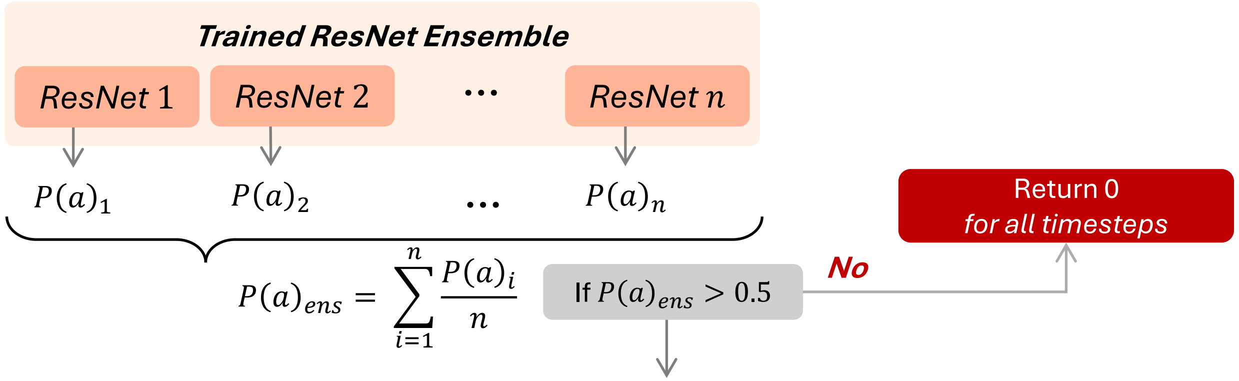

A. Petralia et al., *Few Labels are all you need: A Weakly Supervised Framework for Appliance Localization in Smart-Meter Series*, ICDE, 2025.

Improving CAM for Appliance Localization

A **ResNet ensemble** with **varying kernel sizes**

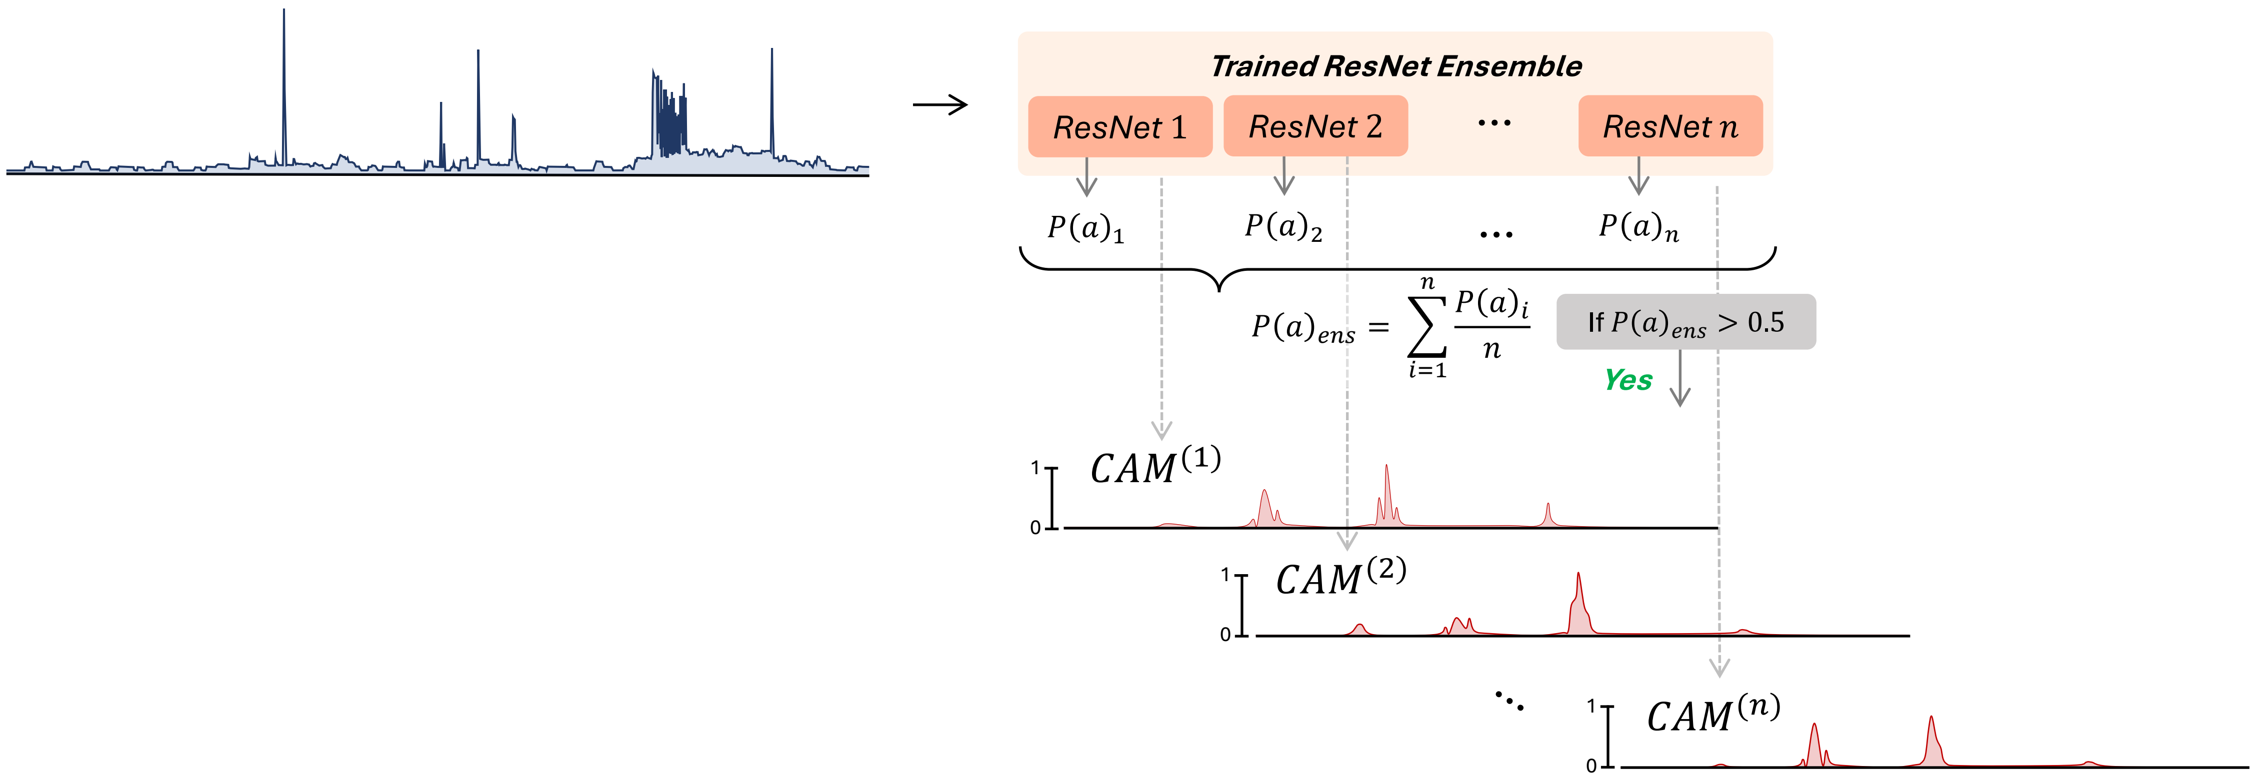

A. Petralia et al., *Few Labels are all you need: A Weakly Supervised Framework for Appliance Localization in Smart-Meter Series*, ICDE, 2025.

Improving **CAM** for **A**ppliance **L**ocalization

A **ResNet ensemble** with **varying kernel sizes**

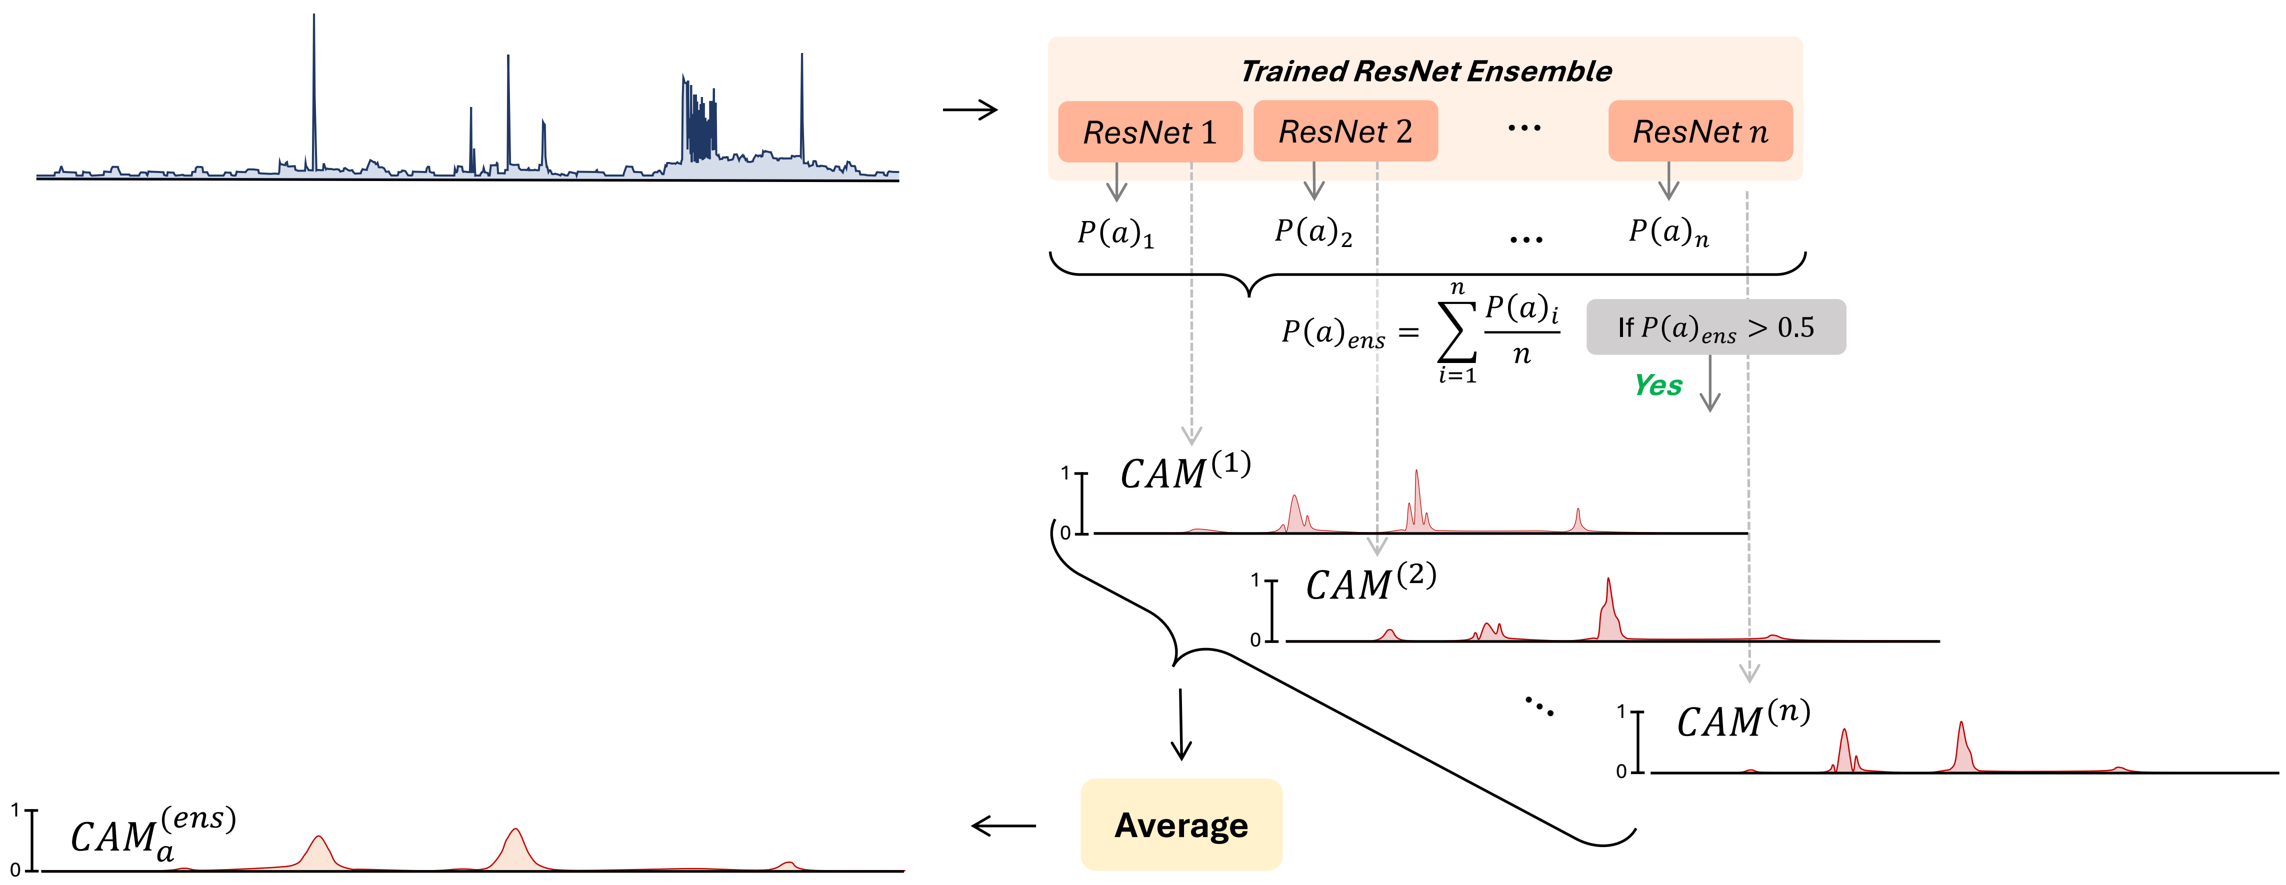

A. Petralia et al., *Few Labels are all you need: A Weakly Supervised Framework for Appliance Localization in Smart-Meter Series*, ICDE, 2025.

Proposed Approach: CamAL

Improving **CAM** for **A**ppliance **L**ocalization

A ResNet ensemble with varying kernel sizes

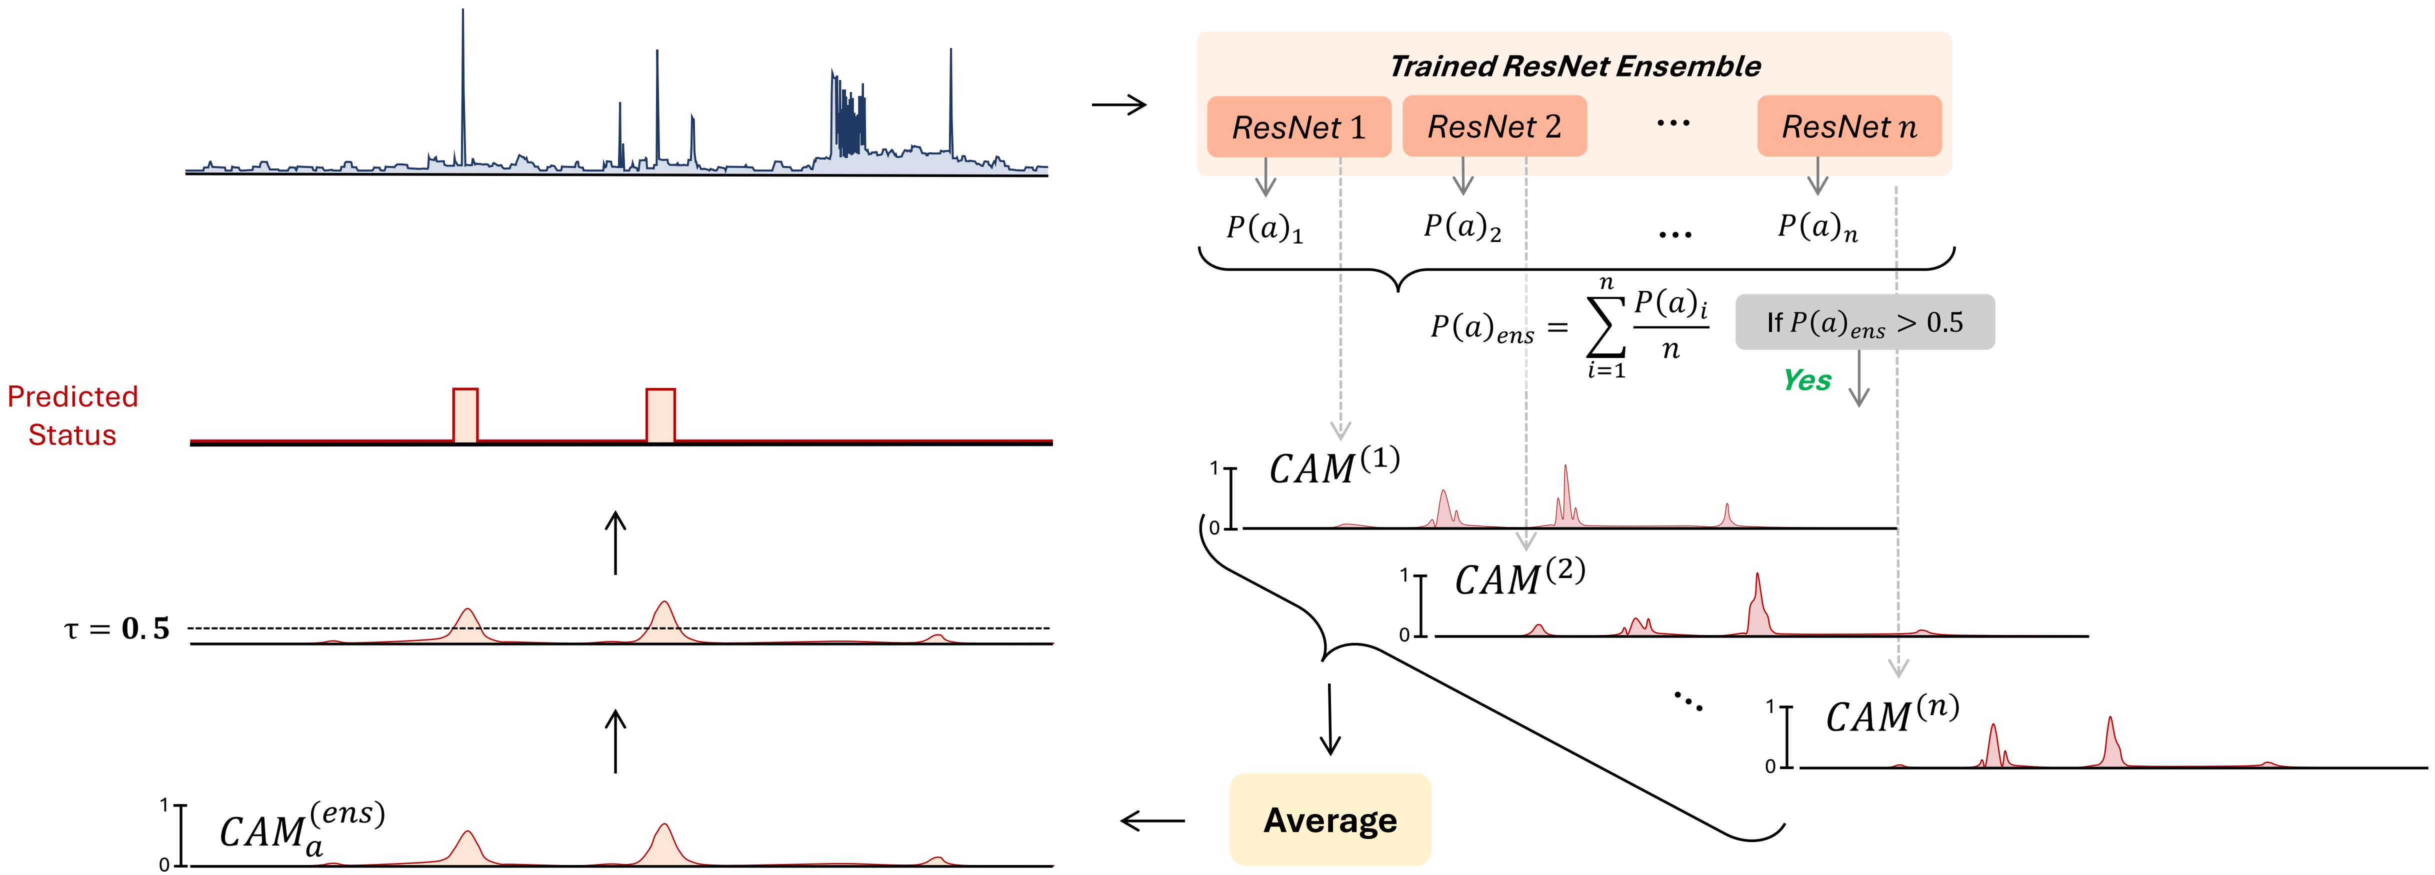

A. Petralia et al., *Few Labels are all you need: A Weakly Supervised Framework for Appliance Localization in Smart-Meter Series*, ICDE, 2025.

Proposed Approach: CamAL

Improving CAM for Appliance Localization

A **ResNet ensemble** with **varying kernel sizes**

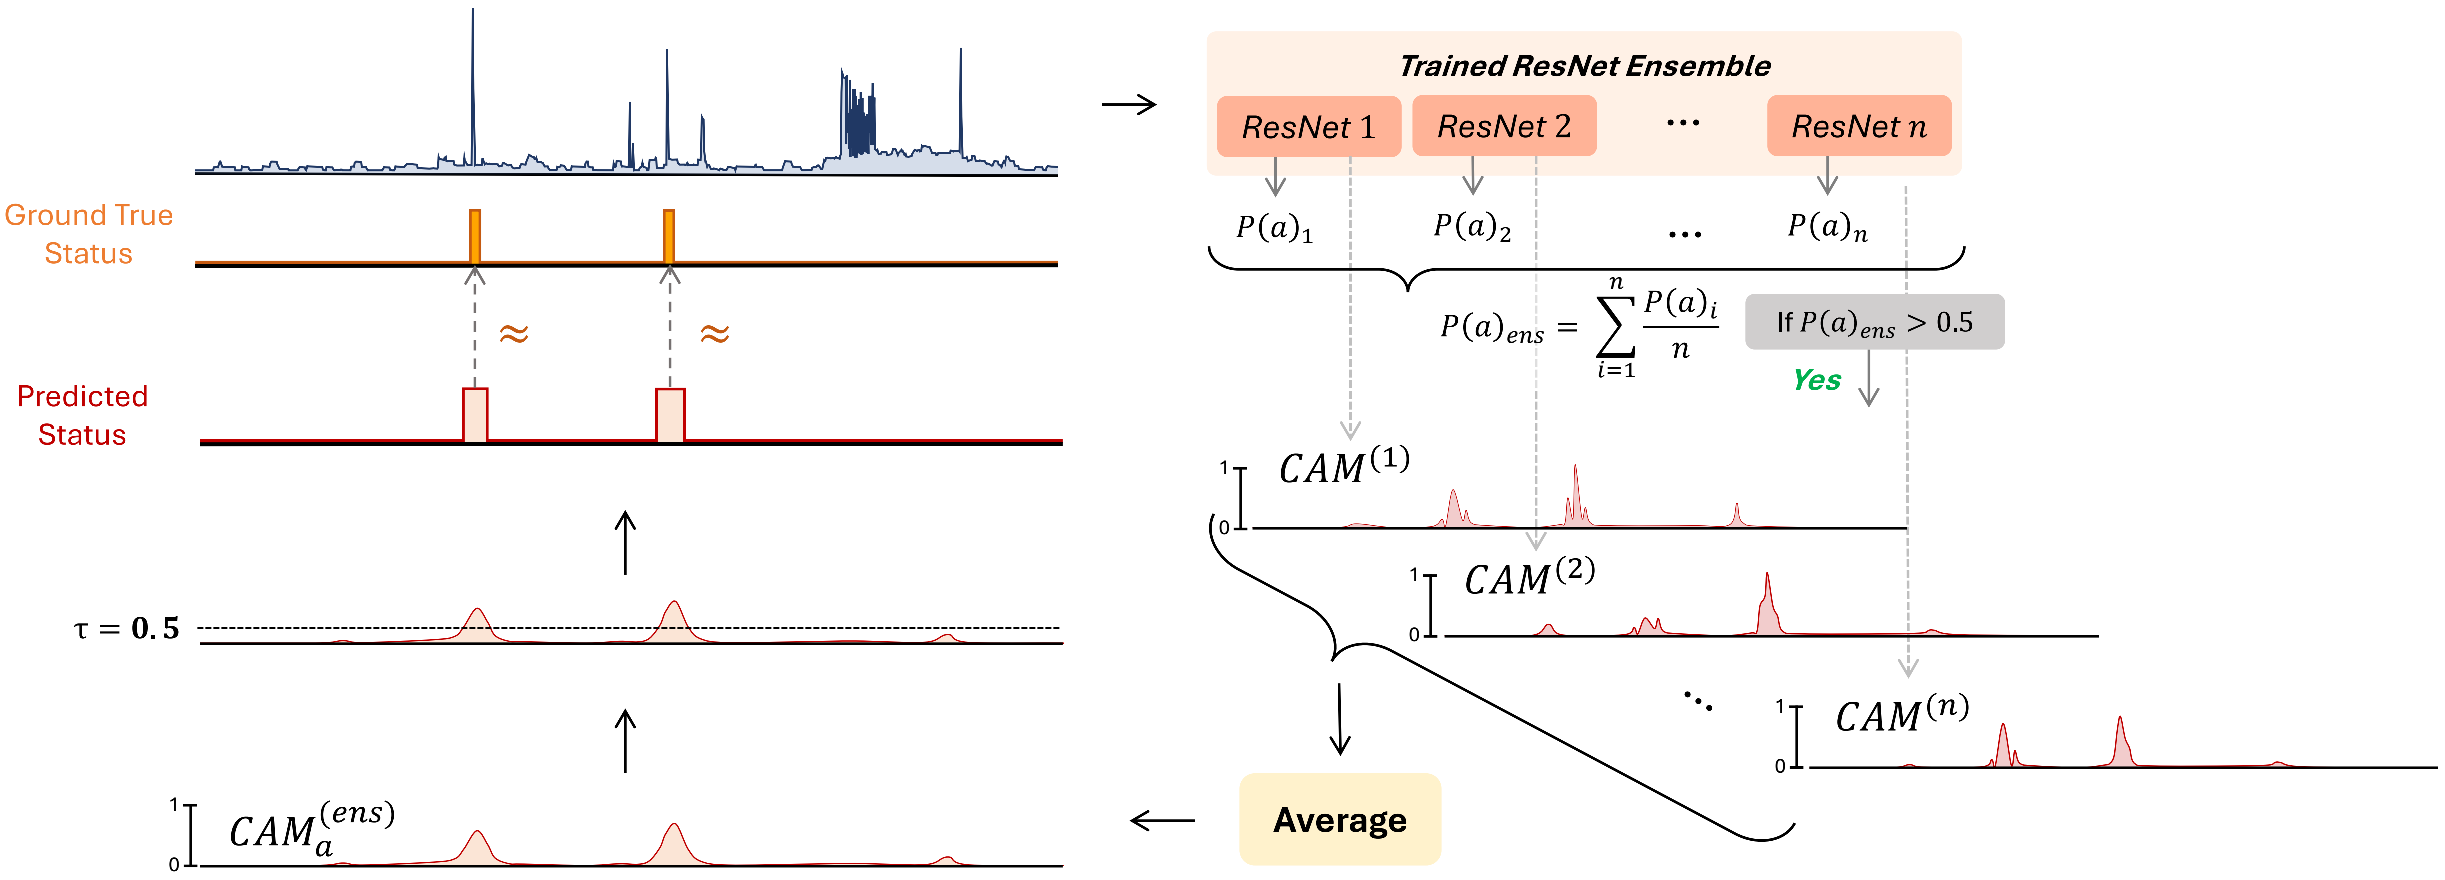

A. Petralia et al., *Few Labels are all you need: A Weakly Supervised Framework for Appliance Localization in Smart-Meter Series*, ICDE, 2025.

Proposed Approach: CamAL

Improving **CAM** for **A**ppliance **L**ocalization

Continues to depend on a **hyperparameter,** which needs to be **manually tuned** for **each scenario**

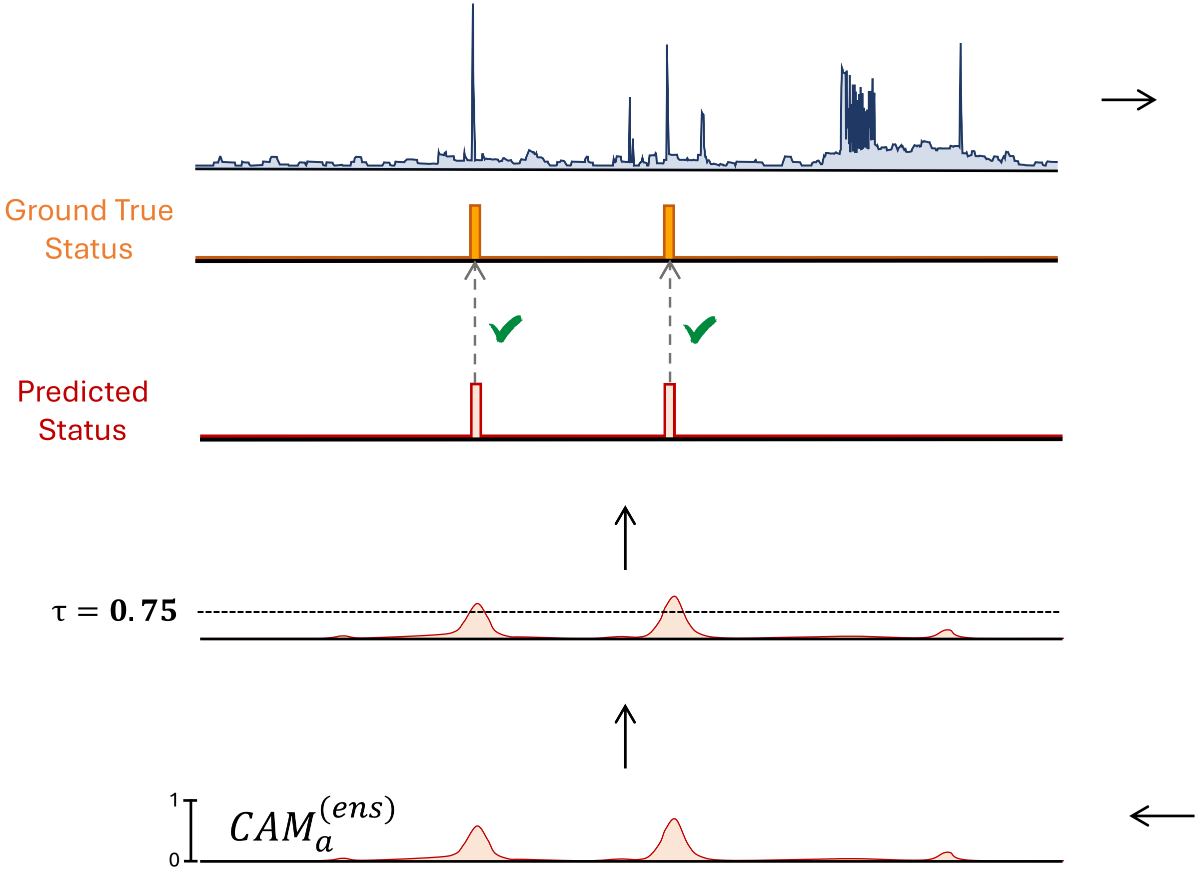

A **ResNet ensemble** with **varying kernel sizes**

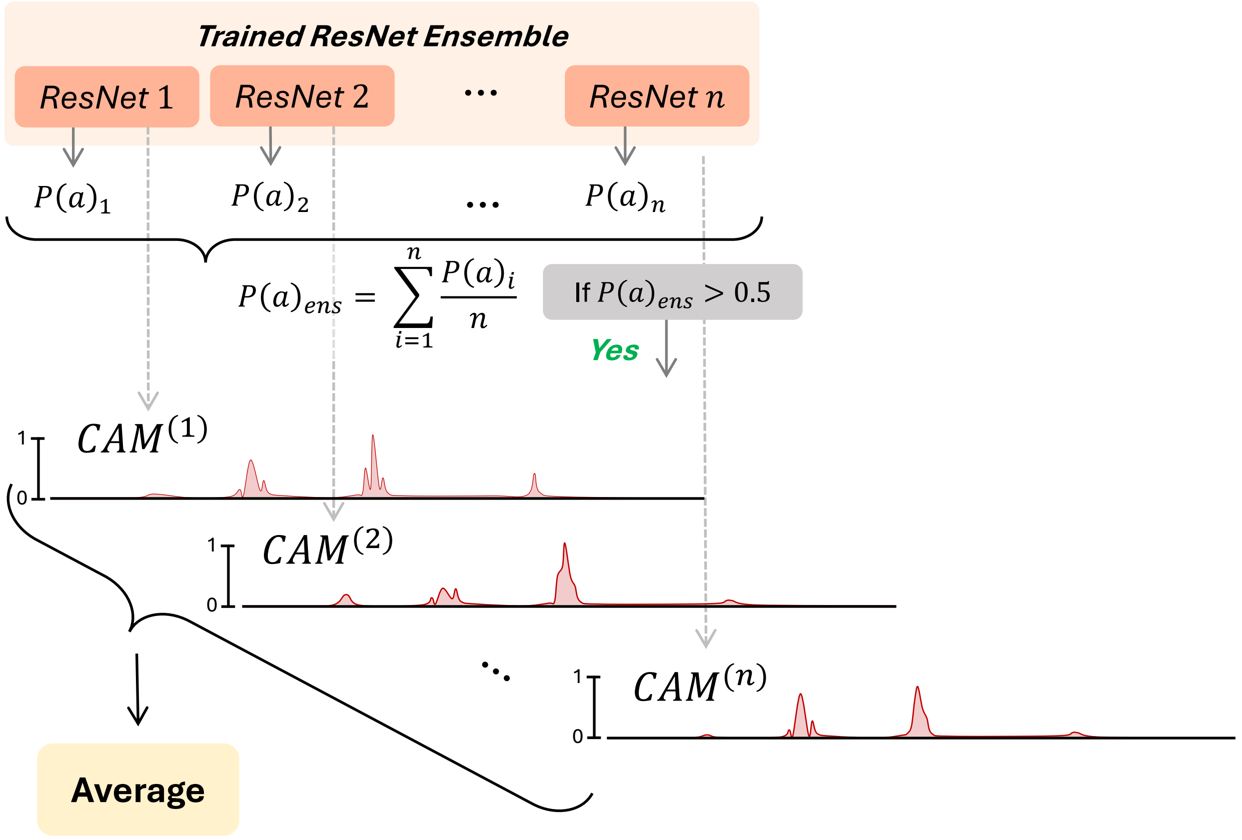

A. Petralia et al., *Few Labels are all you need: A Weakly Supervised Framework for Appliance Localization in Smart-Meter Series*, ICDE, 2025.

Proposed Approach: CamAL

Improving **CAM** for **A**ppliance **L**ocalization

*Considers the **shape** of the **input aggregated signal!***

A **ResNet ensemble** with **varying kernel sizes**

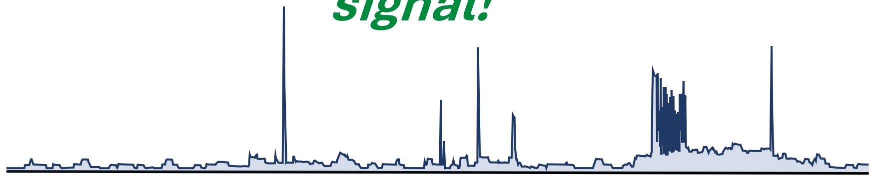

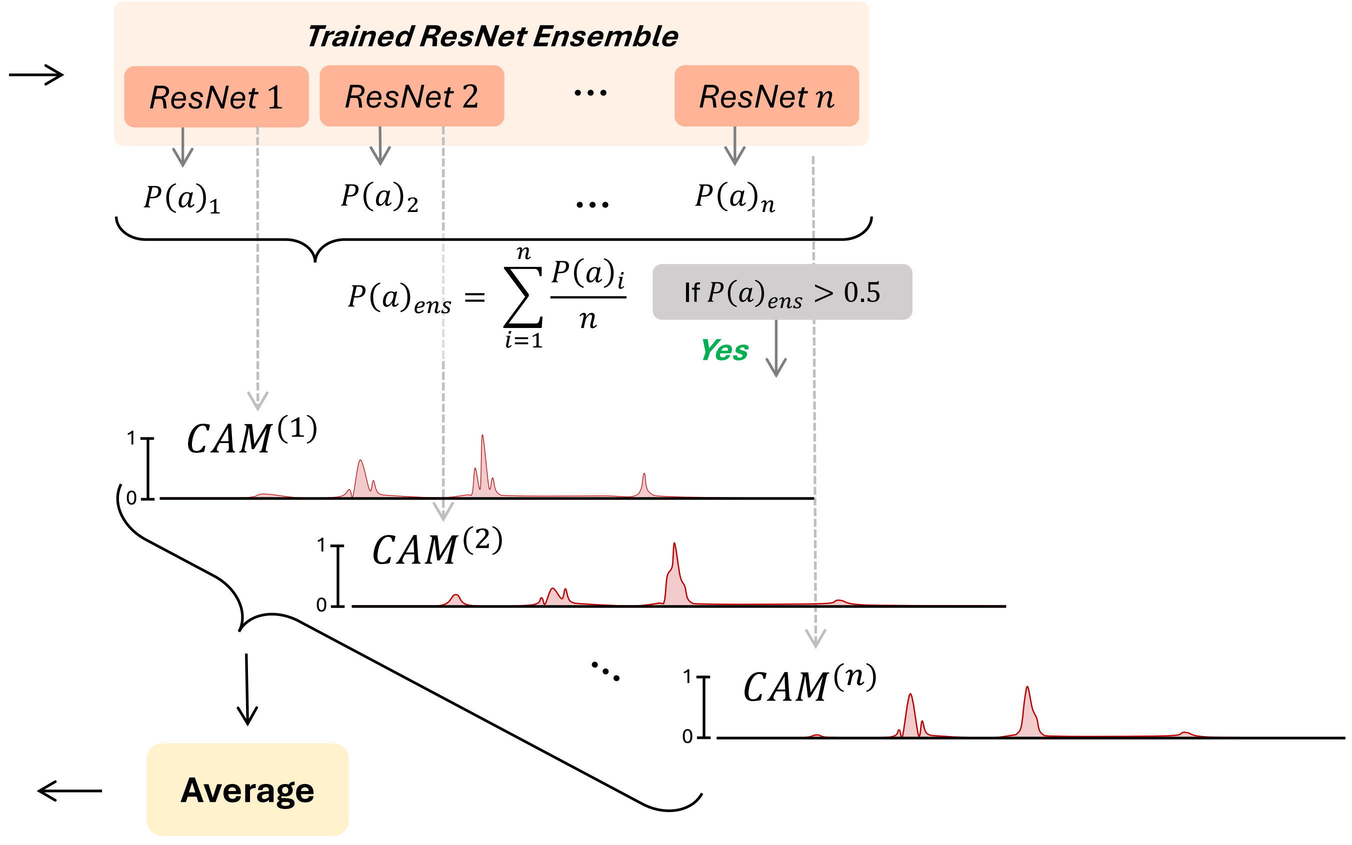

A. Petralia et al., *Few Labels are all you need: A Weakly Supervised Framework for Appliance Localization in Smart-Meter Series*, ICDE, 2025.

Improving CAM for Appliance Localization

*Considers the **shape** of the **input aggregated signal!***

A **ResNet ensemble** with **varying kernel sizes**

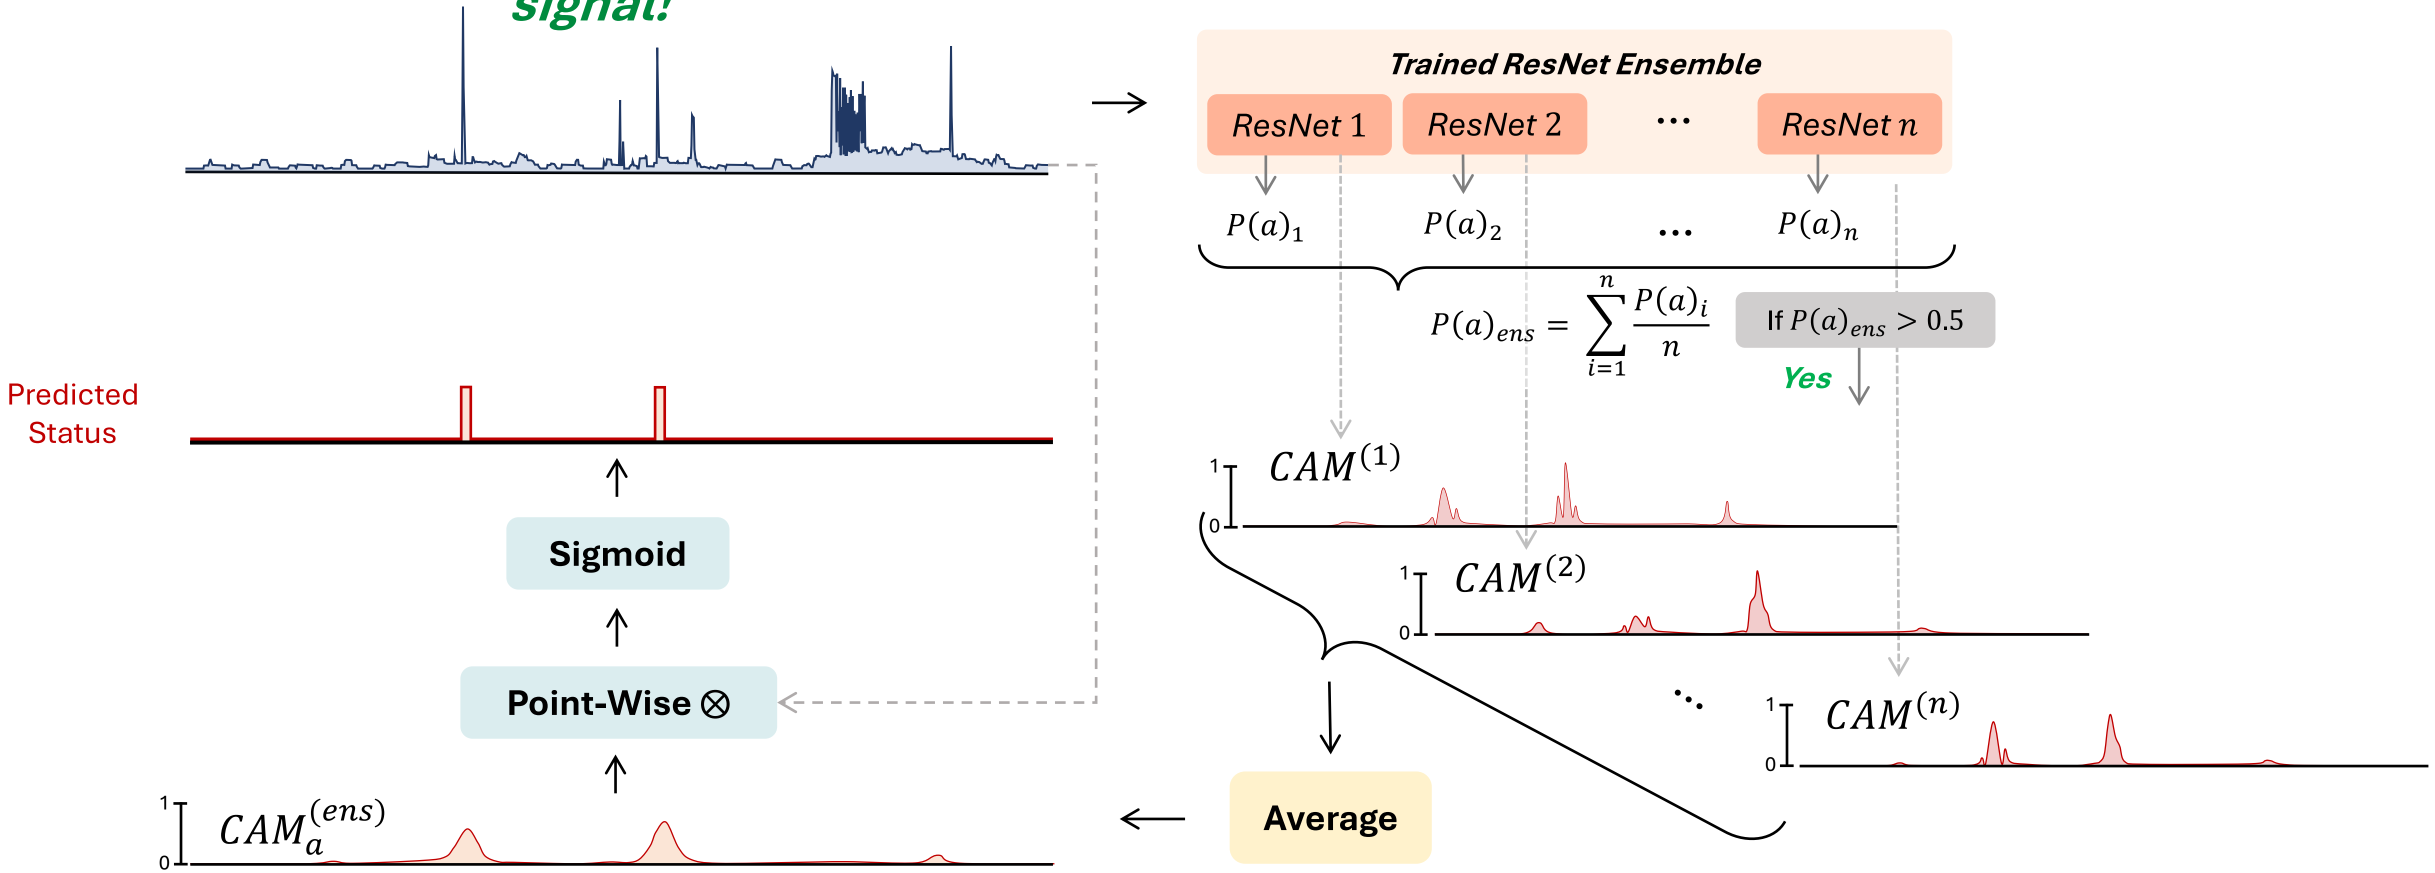

A. Petralia et al., *Few Labels are all you need: A Weakly Supervised Framework for Appliance Localization in Smart-Meter Series*, ICDE, 2025.

Proposed Approach: CamAL

Improving **CAM** for **A**ppliance **L**ocalization

*Considers the **shape** of the **input aggregated signal!***

A **ResNet ensemble** with **varying kernel sizes**

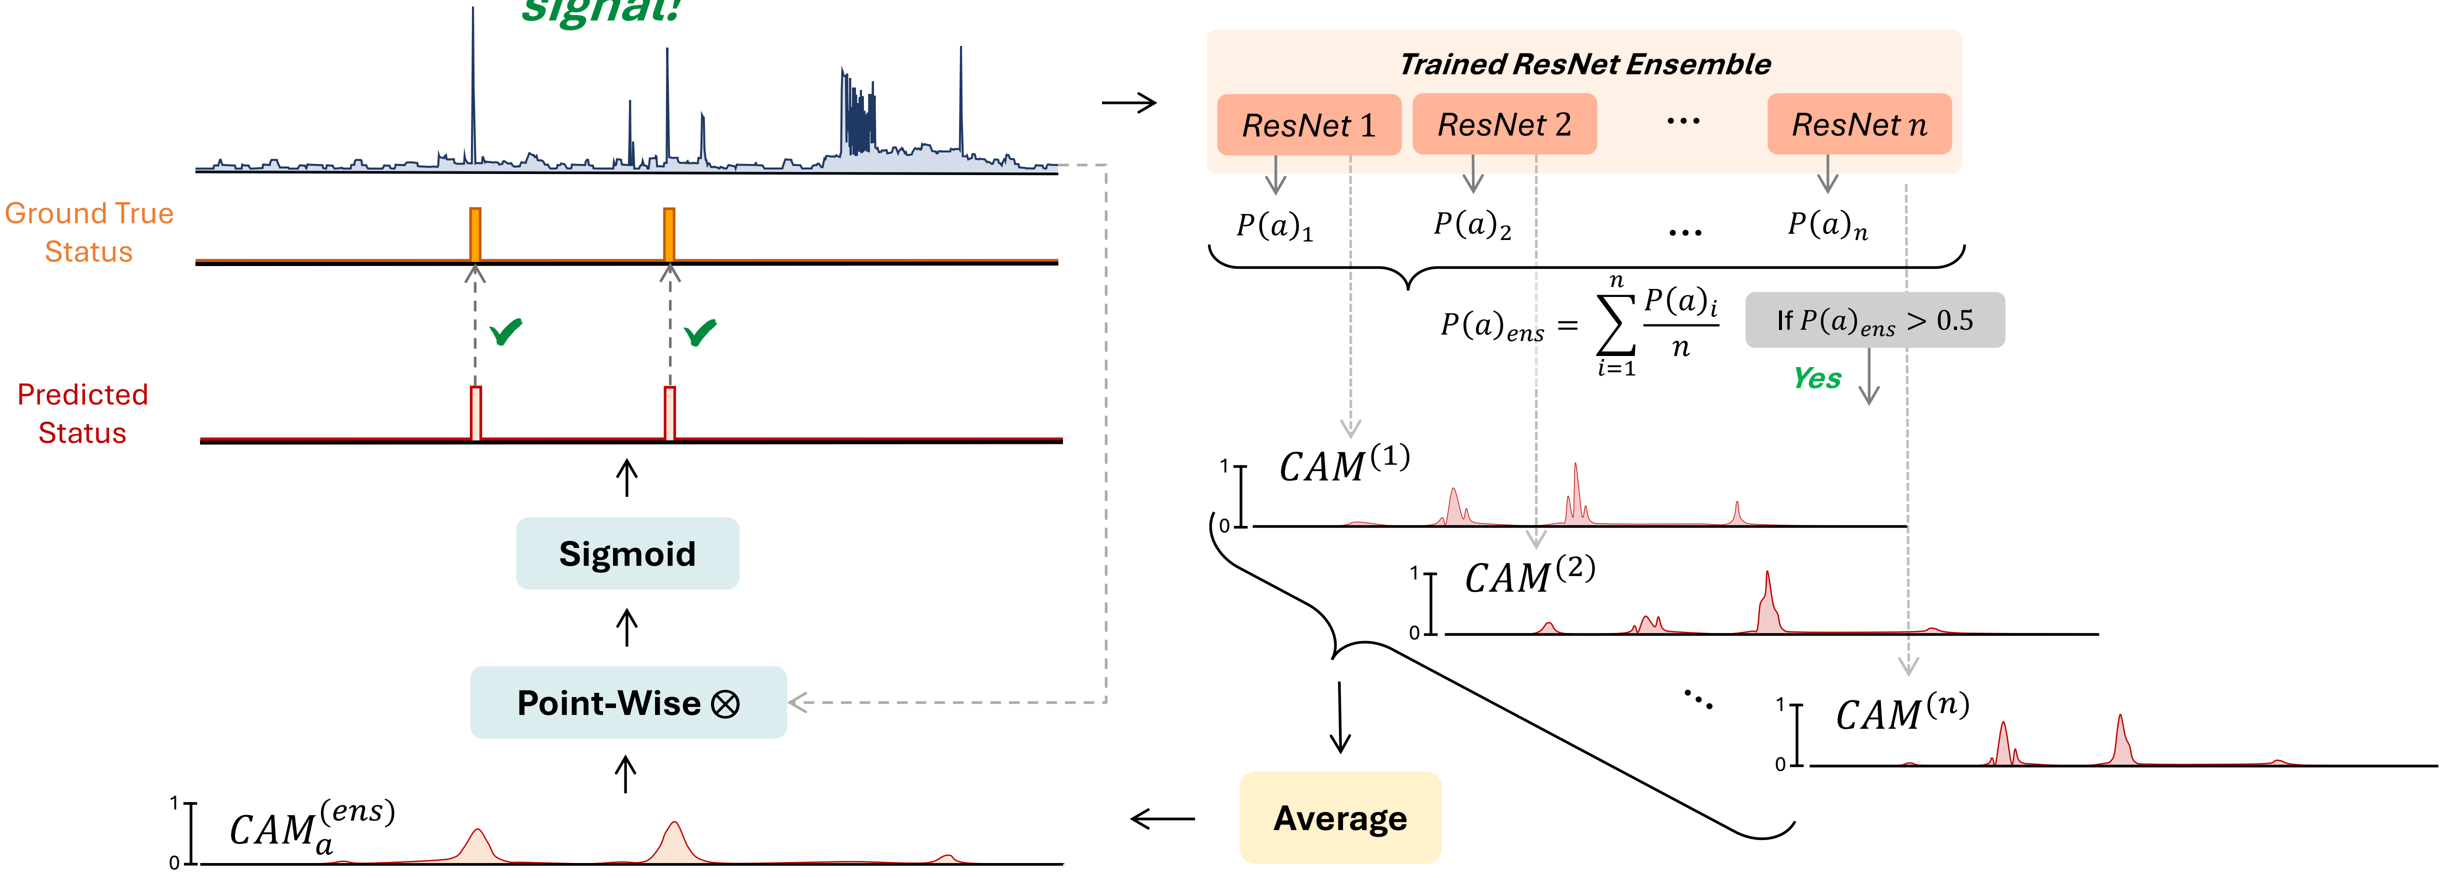

A. Petralia et al., *Few Labels are all you need: A Weakly Supervised Framework for Appliance Localization in Smart-Meter Series*, ICDE, 2025.

How does **CamAL perform** compared to **strongly-supervised baselines**?

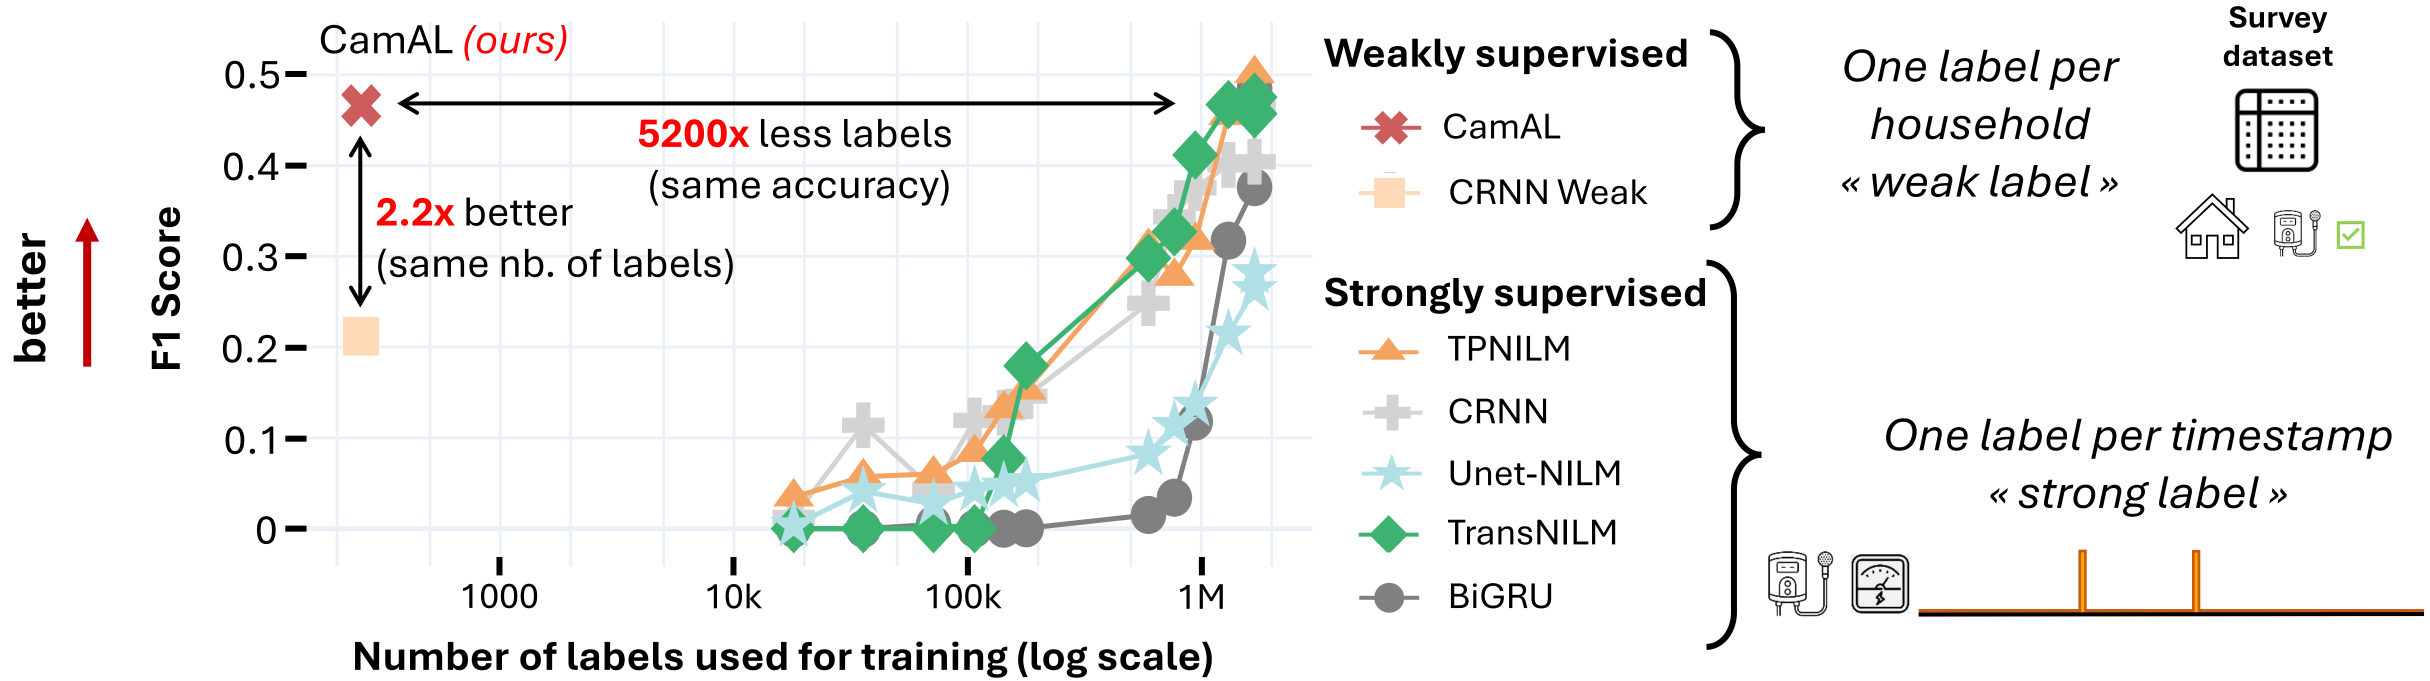

A. Petralia et al., *Few Labels are all you need:* A Weakly Supervised Framework for Appliance Localization in Smart-Meter Series, ICDE, 2025.

Results: Label Costs Comparison

How do **label-collection costs** vary between **approaches**? dollars$ gCO2

Cost *for training*
CamAL

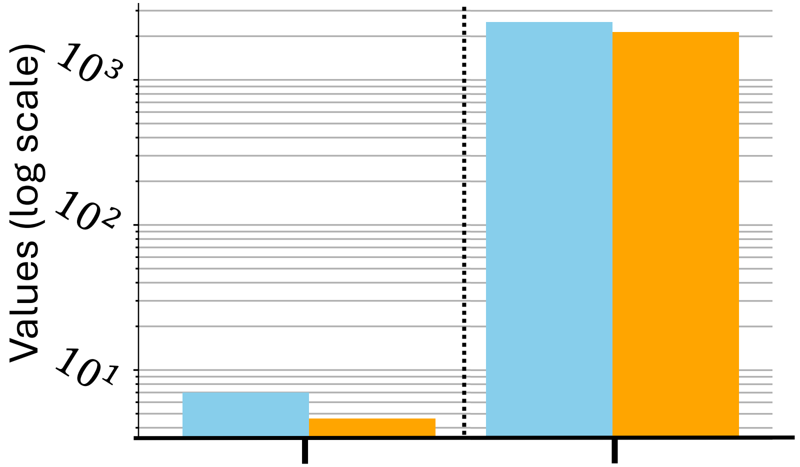

Cost *for training*
strongly supervised
NILM methods

One label per: **Household** **Timestamp**

Survey dataset

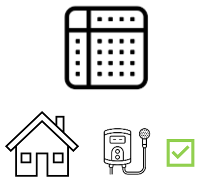

Asking customers to
fill out a questionnaire

Instrumenting households with
dedicated submeters per appliance

Reduce *collection* ***cost*** *(gCO2 and cash)* ***by up*** *to* ***2 magnitude orders!***

A. Petralia et al., *Few Labels are all you need: A Weakly Supervised Framework for Appliance Localization in Smart-Meter Series*, ICDE, 2025.

Problem

*Can we tackle the **Appliance Pattern Localization** problem using **minimal supervision**?*

Challenge

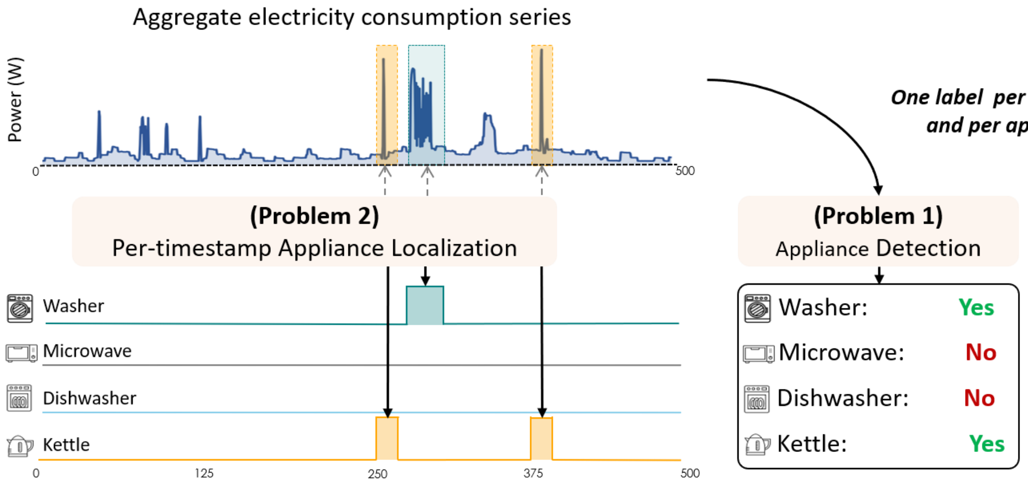

Solution

✓ **CamAL** (**C**lass **A**ctivation **M**ap based **A**ppliance **L**ocalization)

- Combine **explainable AI** with **weak labels** to tackle **appliance-pattern localization**
- Achieve near–**strongly supervised** method's accuracy while drastically **reducing labeling costs**

I. Introduction

II. Contributions

1. Appliance Detection Presence in Consumers Household

2. Appliance Pattern Localization

3. Energy Disaggregation

III. Conclusions

1. Appliance Detection - *Time Series Classification*

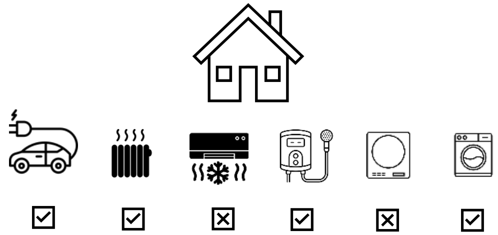

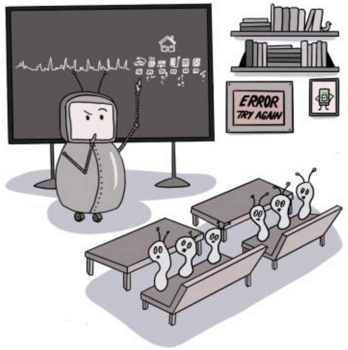

2. Appliance Pattern Localization – *Pattern Identification*

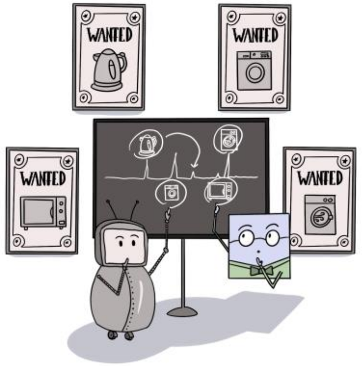

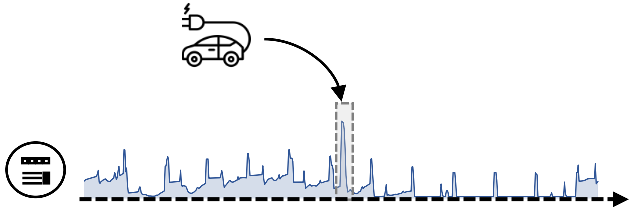

3. Energy Disaggregation – *Time Series Regression*

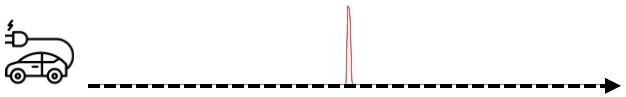

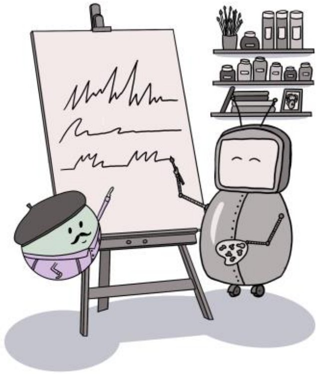

Background: EDF's monitoring solution (Mon Suivi Conso)

EDF's Appliance-Level Feedback Solution

- 📈 **2015** - Launch of *Mon Suivi Conso* (web + app)
- 🛠️ **2018** - **Annual appliances estimate** using a **semi-supervised** statistics approach[1]
- 📊 **2023** - **Deep-Learning based** approach → monthly estimation reduced error **by** $\approx -70\ \%$

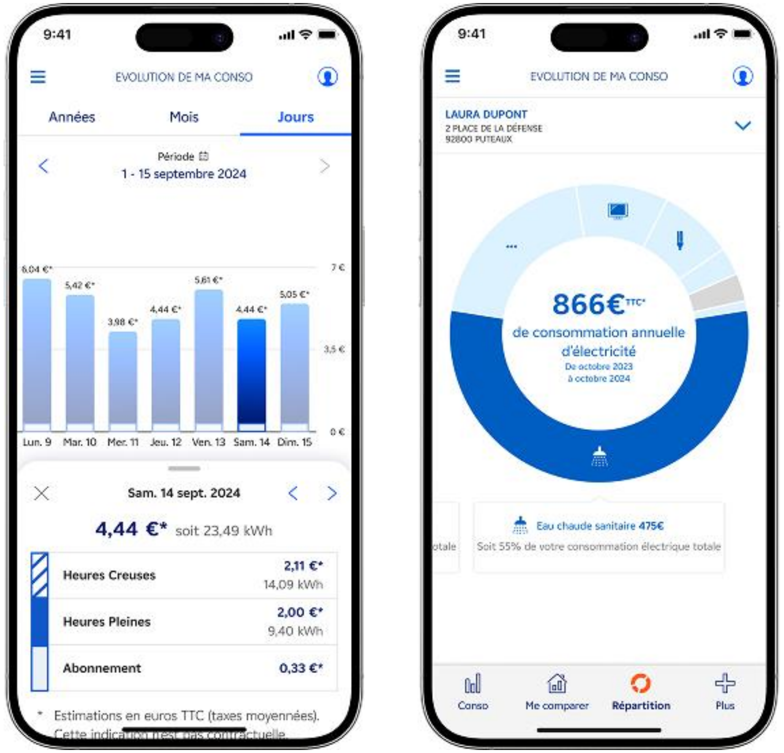

Room for improvement: Monthly estimation is **still coarse, and users recently requested** daily appliance-level **insights**

[1] L. Bozzi et al., Individual electricity consumption of a given appliance from a set of electrical equipment, French Patent, 2014.

Background: SotA NILM Approaches

Energy Disaggregation

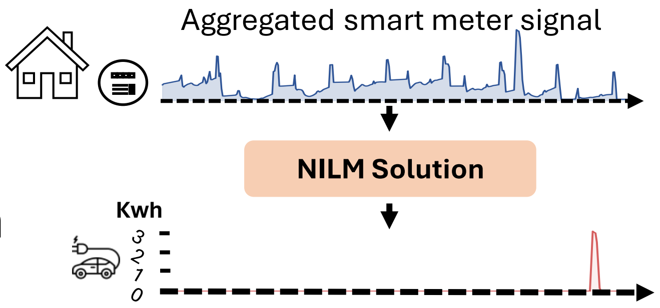

- Current **SotA methods** are based on **deep-learning**
- **Operates** on **subsequences** of an entire electricity consumption series

Energy Disaggregation

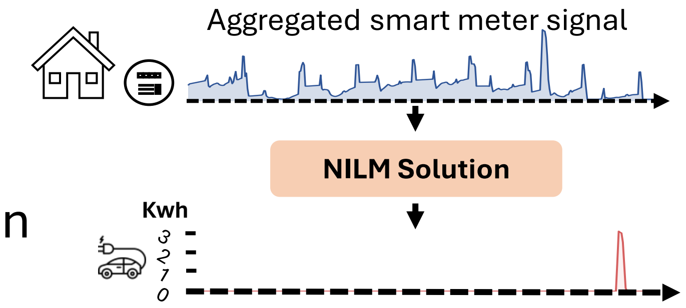

Current **SotA methods** are based on **deep-learning**

Operates on **subsequences** of an entire electricity consumption series

The *Sequence-To-Sequence* paradigm

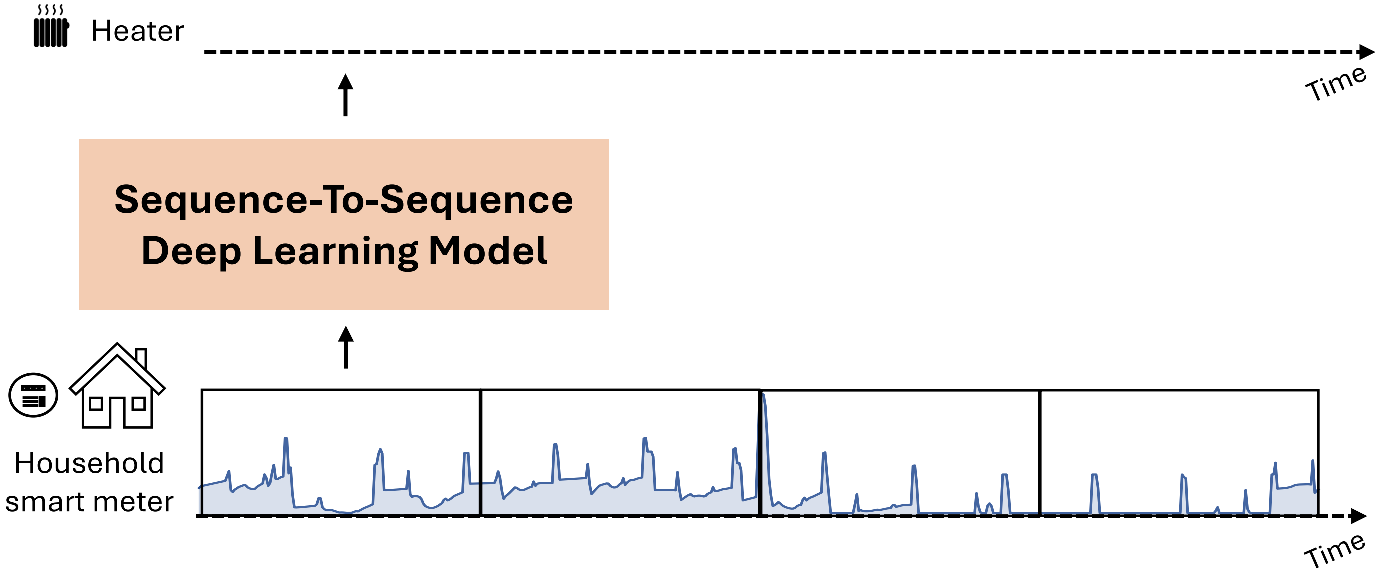

Problem

Non-Stationarity Aspect of Electricity Consumption Data

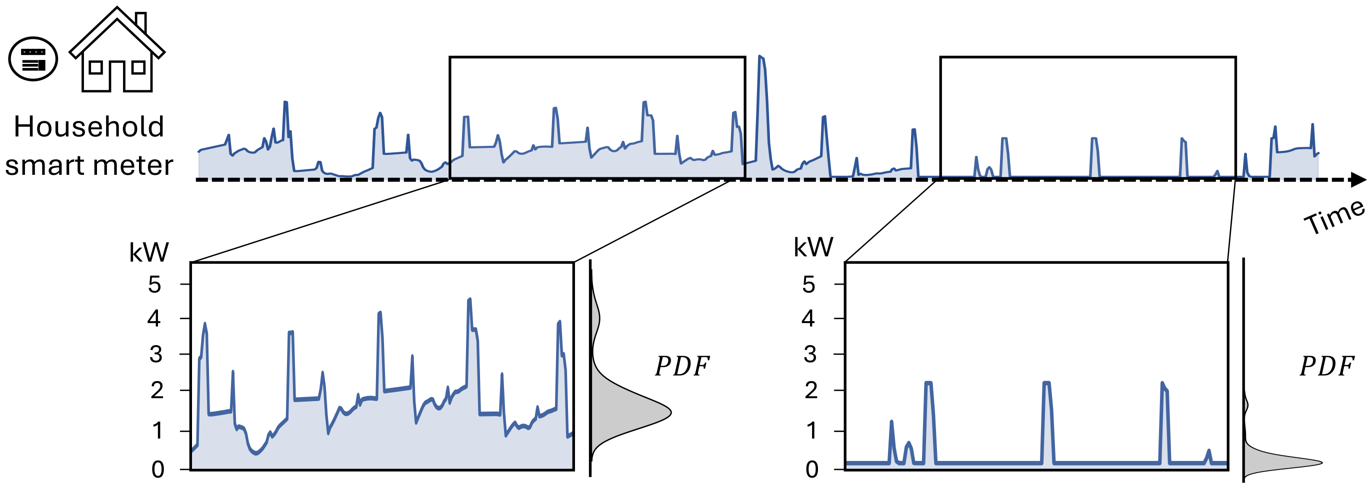

Accounting for **non-stationarity** in deep learning **significantly improves** time series **forecasting accuracy** ! [1, 2]

[1] T. Kim et al., Reversible Instance Normalization for Accurate Time-Series Forecasting against Distribution Shift, ICLR, 2021
[2] Y. Liu et al., Non-stationary Transformers: Exploring the Stationarity in Time Series Forecasting, NIPS, 2022

Problem

*How to provide detailed and accurate **fine-grained** appliance consumption **feedback** to customers?*

Challenges

1. **Considering non-stationary**
   Mitigating the data drift within each subsequence

2. **Delivering granular, actionable feedback to customers**
   Per-timestamp, daily, weekly and monthly

Problem

*How to provide detailed and accurate **fine-grained** appliance consumption **feedback** to customers?*

Challenges

1. **Considering non-stationary**
   Mitigating the data drift within each subsequence

2. **Delivering granular, actionable feedback to customers**
   Per-timestamp, daily, weekly and monthly

Solutions

✓ **NILMFormer**

✓ **Deployment** in "***Mon Suivi Conso***"

Proposed Approach: NILMFormer

How to **mitigate** the subsequence **data drift aspect?**

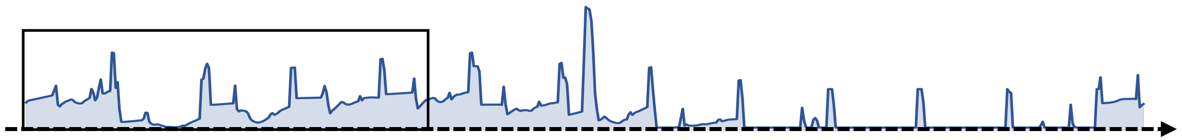

Entire aggregate consumption series

A. Petralia et al., NILMFormer: Non-Intrusive Load Monitoring that Accounts for Non-Stationarity. Under review for KDD 2025.

Proposed Approach: NILMFormer

How to **mitigate** the subsequence **data drift aspect?**

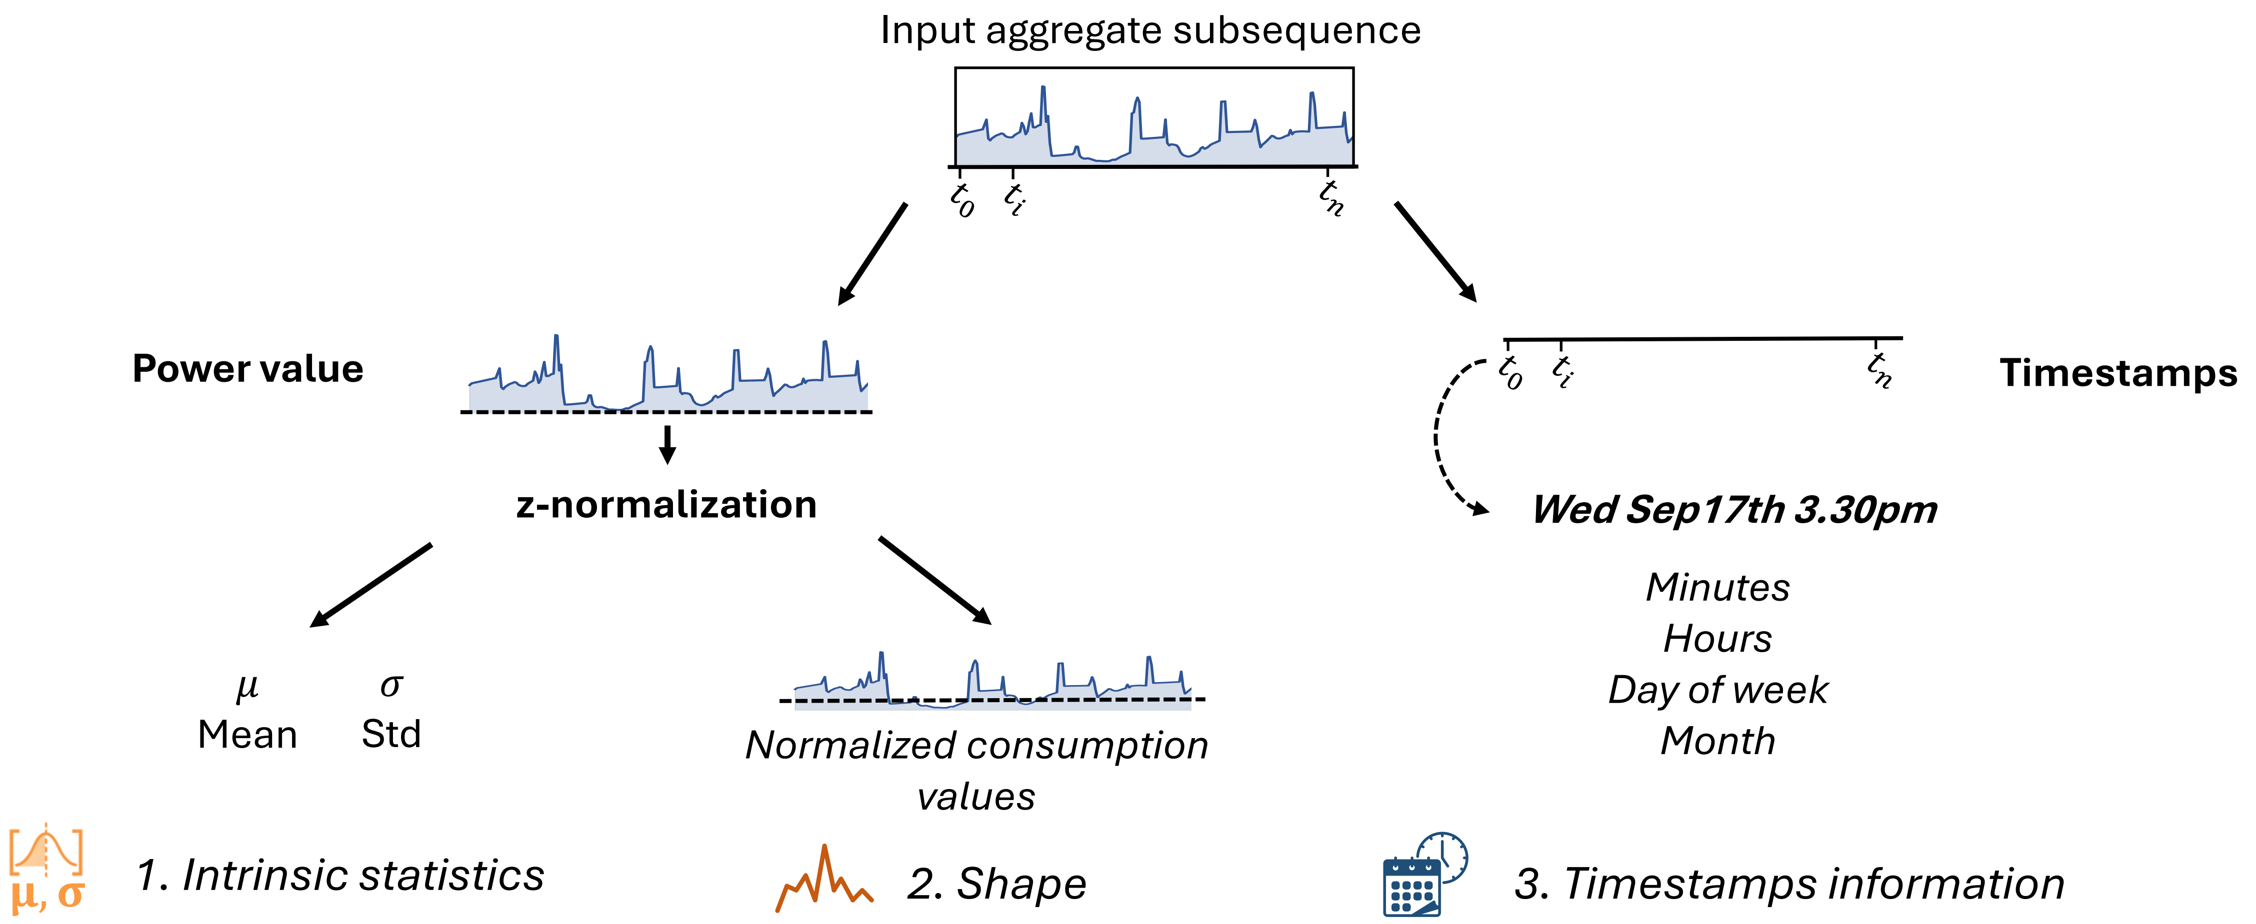

A. Petralia et al., NILMFormer: Non-Intrusive Load Monitoring that Accounts for Non-Stationarity. Under review for KDD 2025.

Proposed Approach: NILMFormer

NILMFormer: A Non-Stationarity Aware Transformer for Non-Intrusive Load Monitoring

I. Distinct encoding modules (*tokenization*)

- *1. Intrinsic statistics*
- *2. Shape*
- *3. Timestamps information*

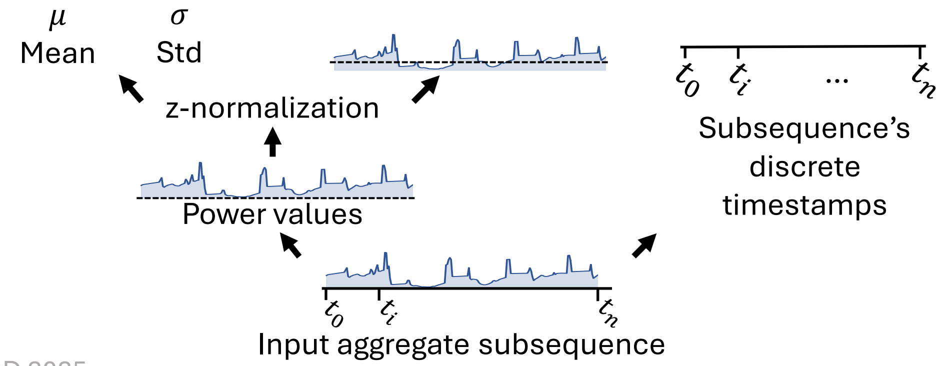

A. Petralia et al., NILMFormer: Non-Intrusive Load Monitoring that Accounts for Non-Stationarity. Under review for KDD 2025.

Proposed Approach: NILMFormer

NILMFormer: A Non-Stationarity Aware Transformer for Non-Intrusive Load Monitoring

I. Distinct encoding modules (*tokenization*)

1. *Intrinsic statistics*
2. *Shape*
3. *Timestamps information*

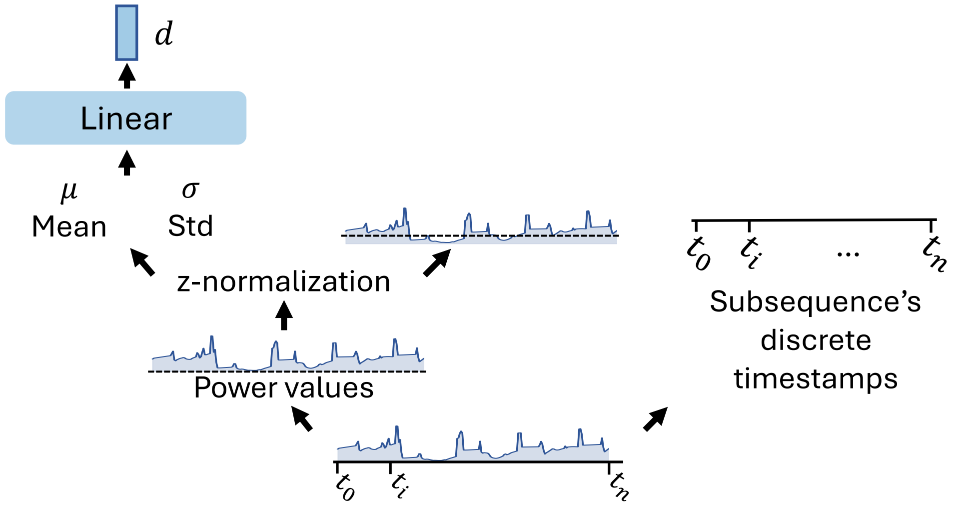

A. Petralia et al., NILMFormer: Non-Intrusive Load Monitoring that Accounts for Non-Stationarity. Under review for KDD 2025.

Proposed Approach: NILMFormer

NILMFormer: A Non-Stationarity Aware Transformer for Non-Intrusive Load Monitoring

I. Distinct encoding modules (*tokenization*)

1. Intrinsic statistics

2. Shape

3. Timestamps information

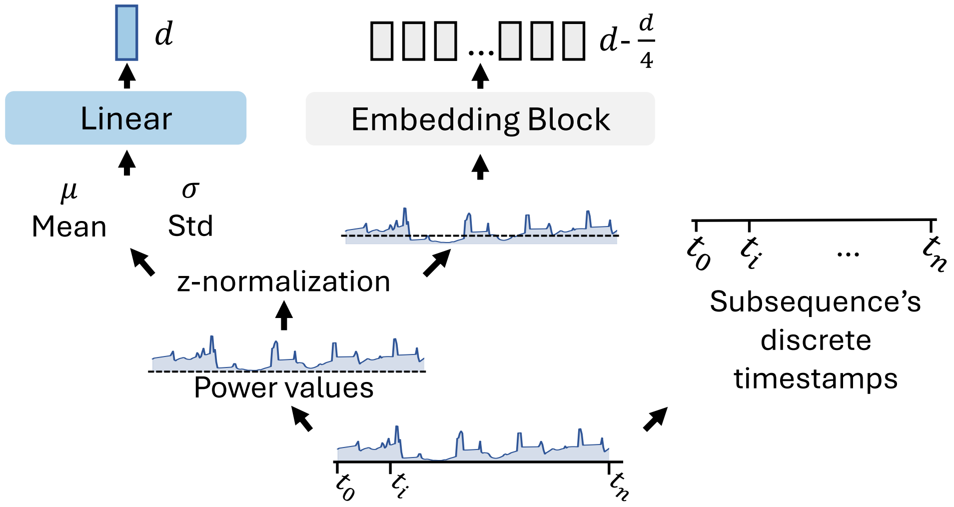

A. Petralia et al., NILMFormer: Non-Intrusive Load Monitoring that Accounts for Non-Stationarity. Under review for KDD 2025.

Proposed Approach: NILMFormer

NILMFormer: A Non-Stationarity Aware Transformer for Non-Intrusive Load Monitoring

I. Distinct encoding modules (*tokenization*)

μ, σ *1. Intrinsic statistics*

2. Shape

3. Timestamps information

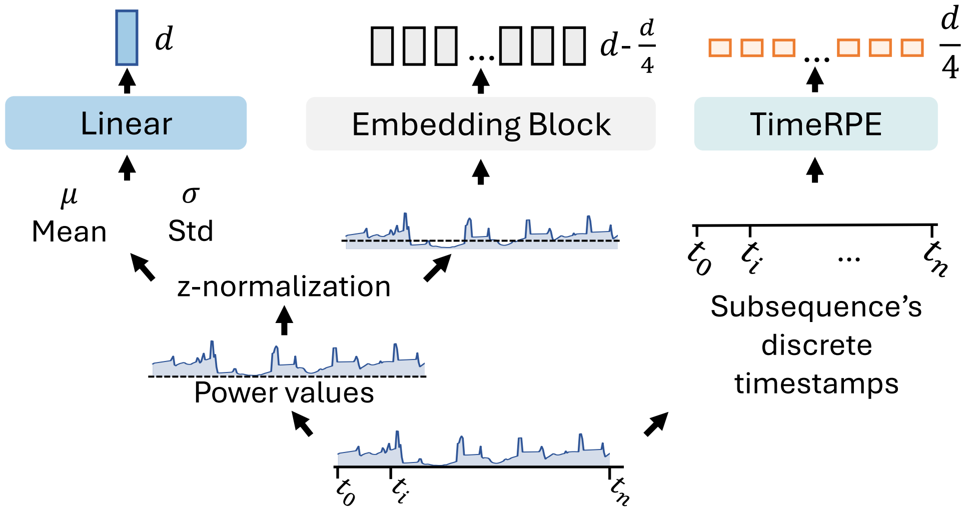

A. Petralia et al., NILMFormer: Non-Intrusive Load Monitoring that Accounts for Non-Stationarity. Under review for KDD 2025.

Proposed Approach: NILMFormer

NILMForm… …r for
Non-Intrus…

I. Distinct e…

μ, σ

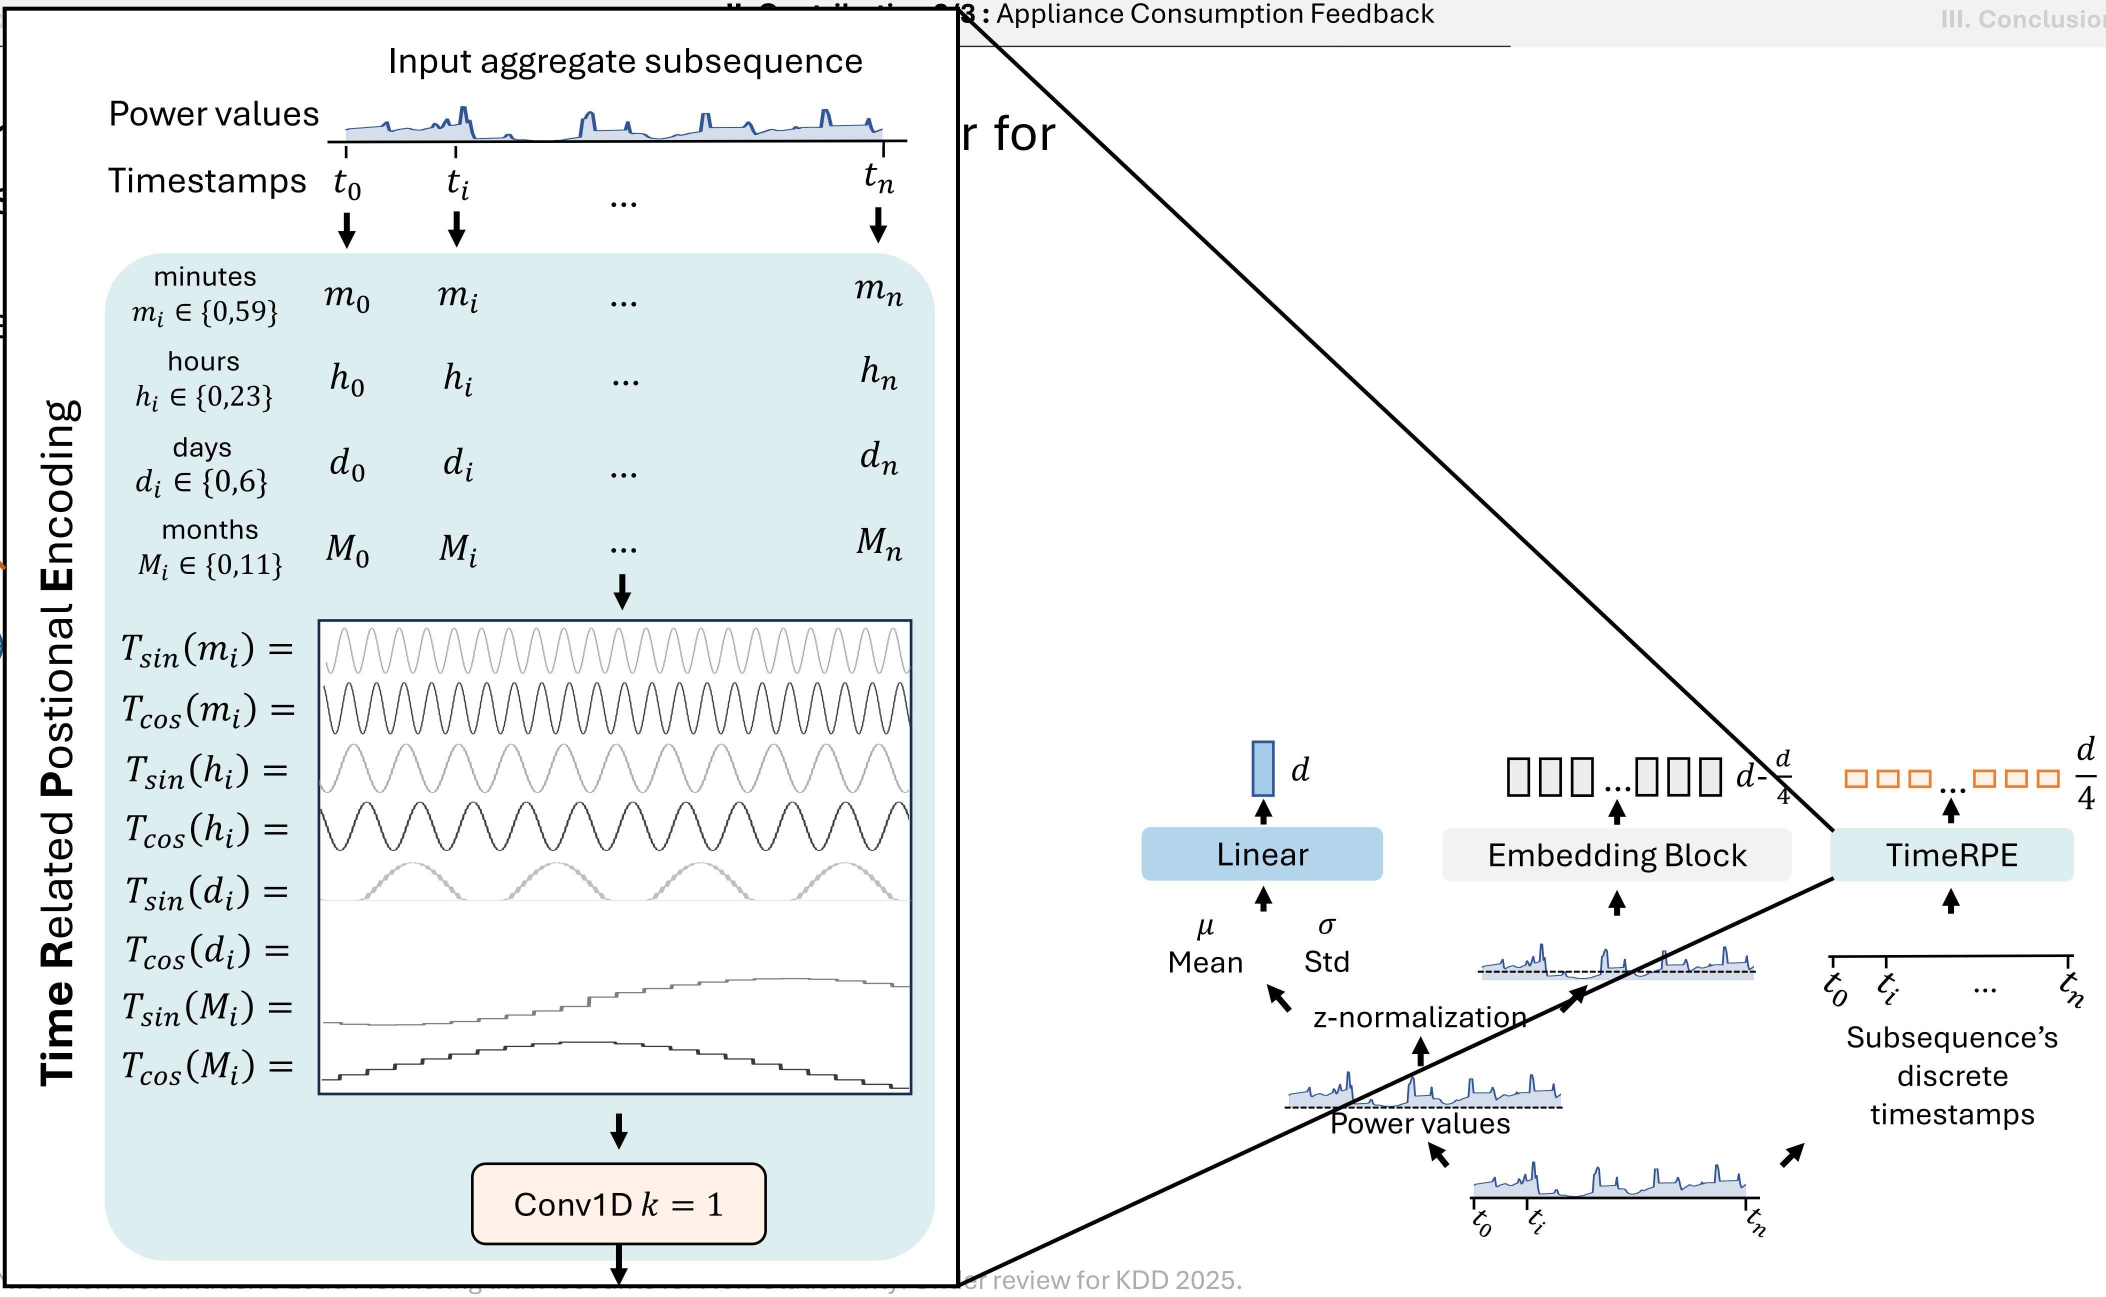

NILMFormer: A Non-Stationarity Aware Transformer for Non-Intrusive Load Monitoring

I. Distinct encoding modules (*tokenization*)

1. *Intrinsic statistics*
2. *Shape*
3. *Timestamps information*

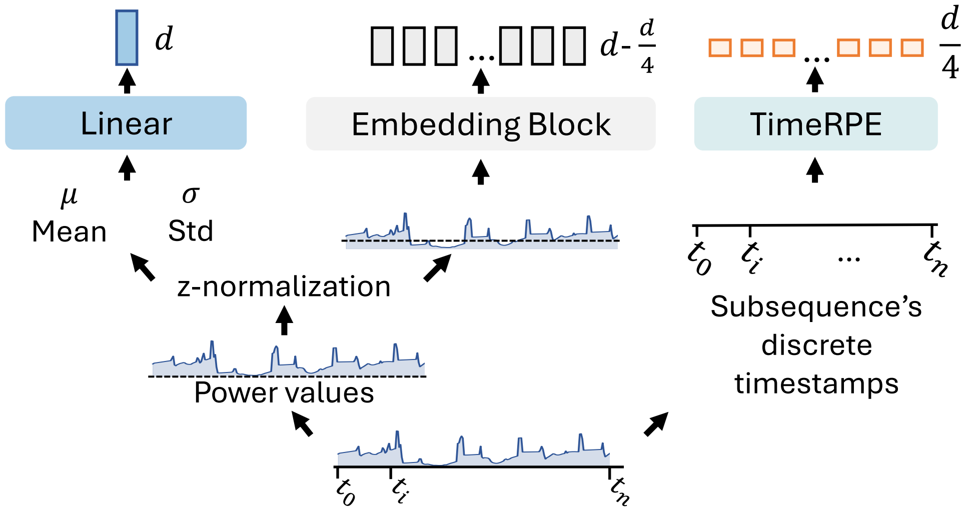

A. Petralia et al., NILMFormer: Non-Intrusive Load Monitoring that Accounts for Non-Stationarity. Under review for KDD 2025.

NILMFormer: A Non-Stationarity Aware Transformer for Non-Intrusive Load Monitoring

I. Distinct encoding modules (*tokenization*)

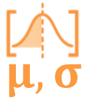

1. Intrinsic statistics

2. Shape

3. Timestamps information

II. Embedding parts **concatenation**

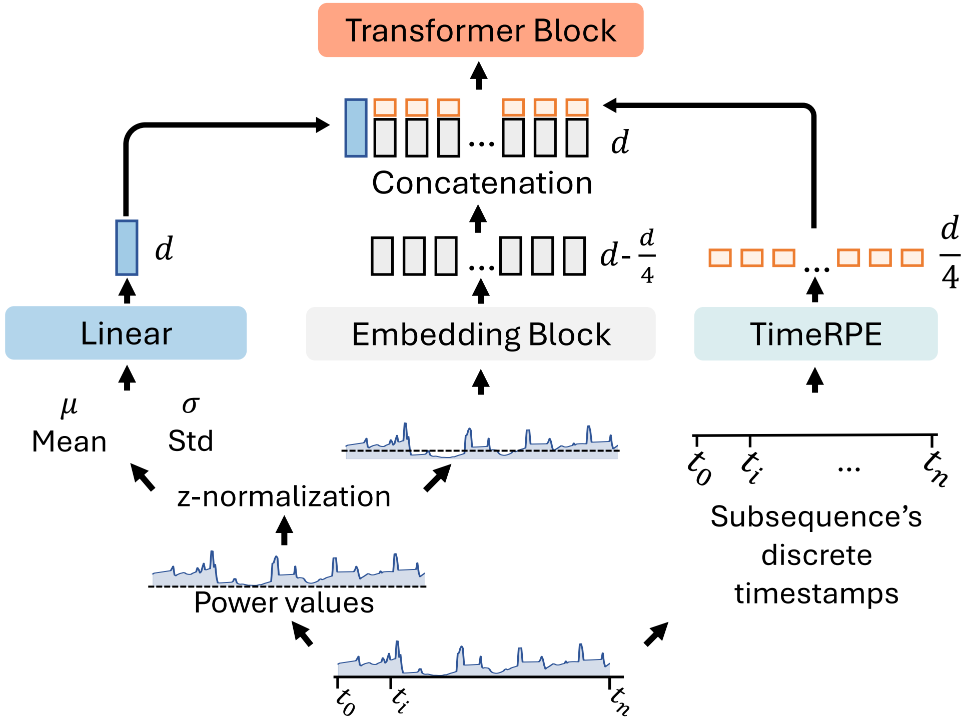

A. Petralia et al., NILMFormer: Non-Intrusive Load Monitoring that Accounts for Non-Stationarity. Under review for KDD 2025.

Proposed Approach: NILMFormer

NILMFormer: A Non-Stationarity Aware Transformer for Non-Intrusive Load Monitoring

I. Distinct encoding modules (*tokenization*)

- μ, σ *1. Intrinsic statistics*
- *2. Shape*
- *3. Timestamps information*

II. Embedding parts **concatenation**

III. Subsequence's individual **appliance power** and **statistics** prediction

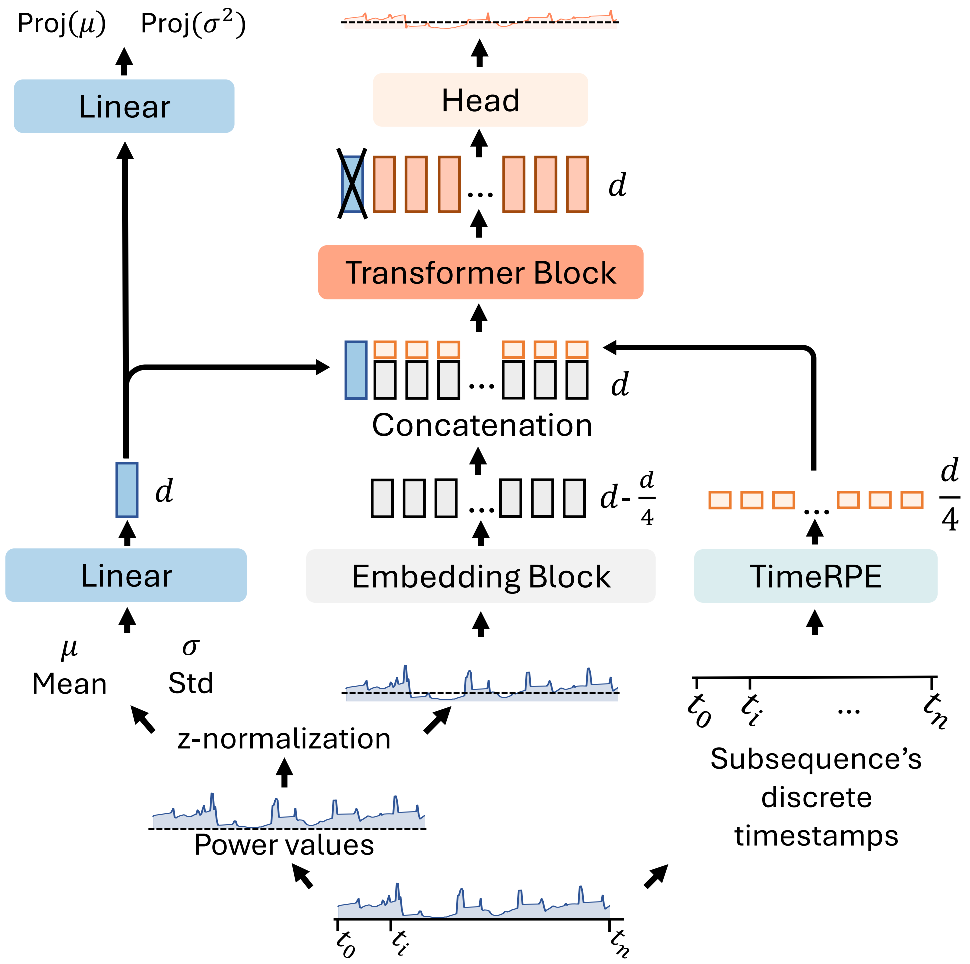

A. Petralia et al., NILMFormer: Non-Intrusive Load Monitoring that Accounts for Non-Stationarity. Under review for KDD 2025.

Proposed Approach: NILMFormer

NILMFormer: A Non-Stationarity Aware Transformer for Non-Intrusive Load Monitoring

I. Distinct encoding modules (*tokenization*)

1. Intrinsic statistics

2. Shape

3. Timestamps information

II. Embedding parts **concatenation**

III. Subsequence's individual **appliance power** and **statistics** prediction

IV. Output **de-normalization**

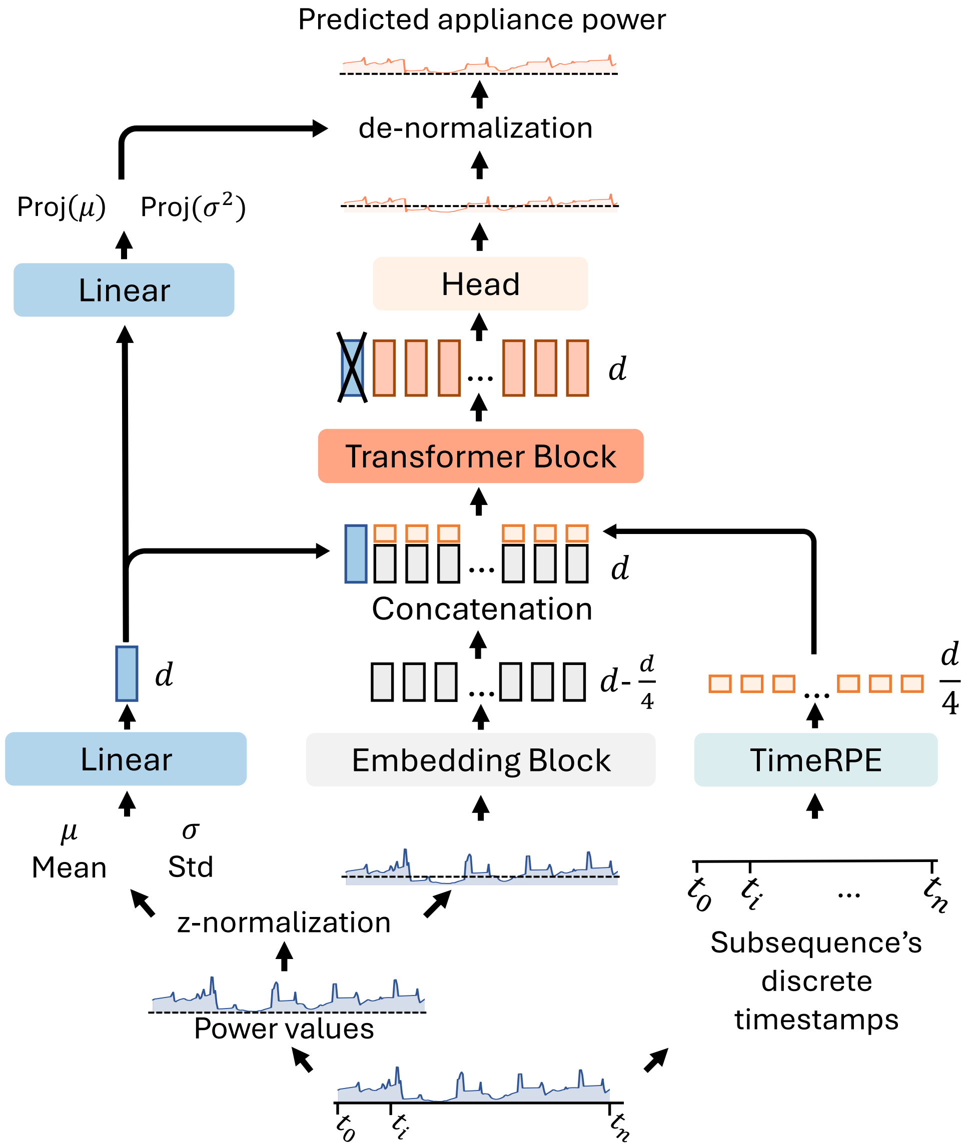

A. Petralia et al., NILMFormer: Non-Intrusive Load Monitoring that Accounts for Non-Stationarity. Under review for KDD 2025.

Results: Per-timestamp Energy Disaggregation

Performance comparison with **10** SotA NILM baselines

Avg. rank on two metrics

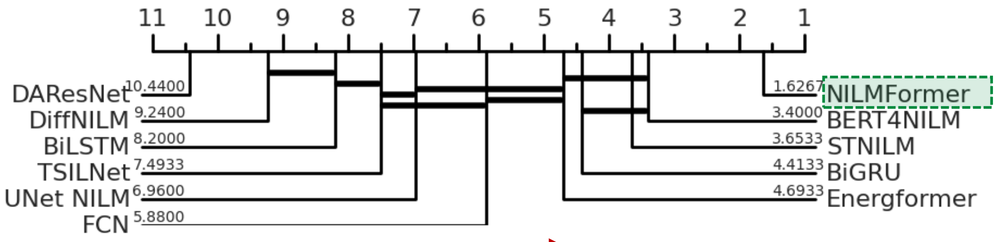

better

Averaged results over 4 datasets and 14 appliances disaggregation scenarios

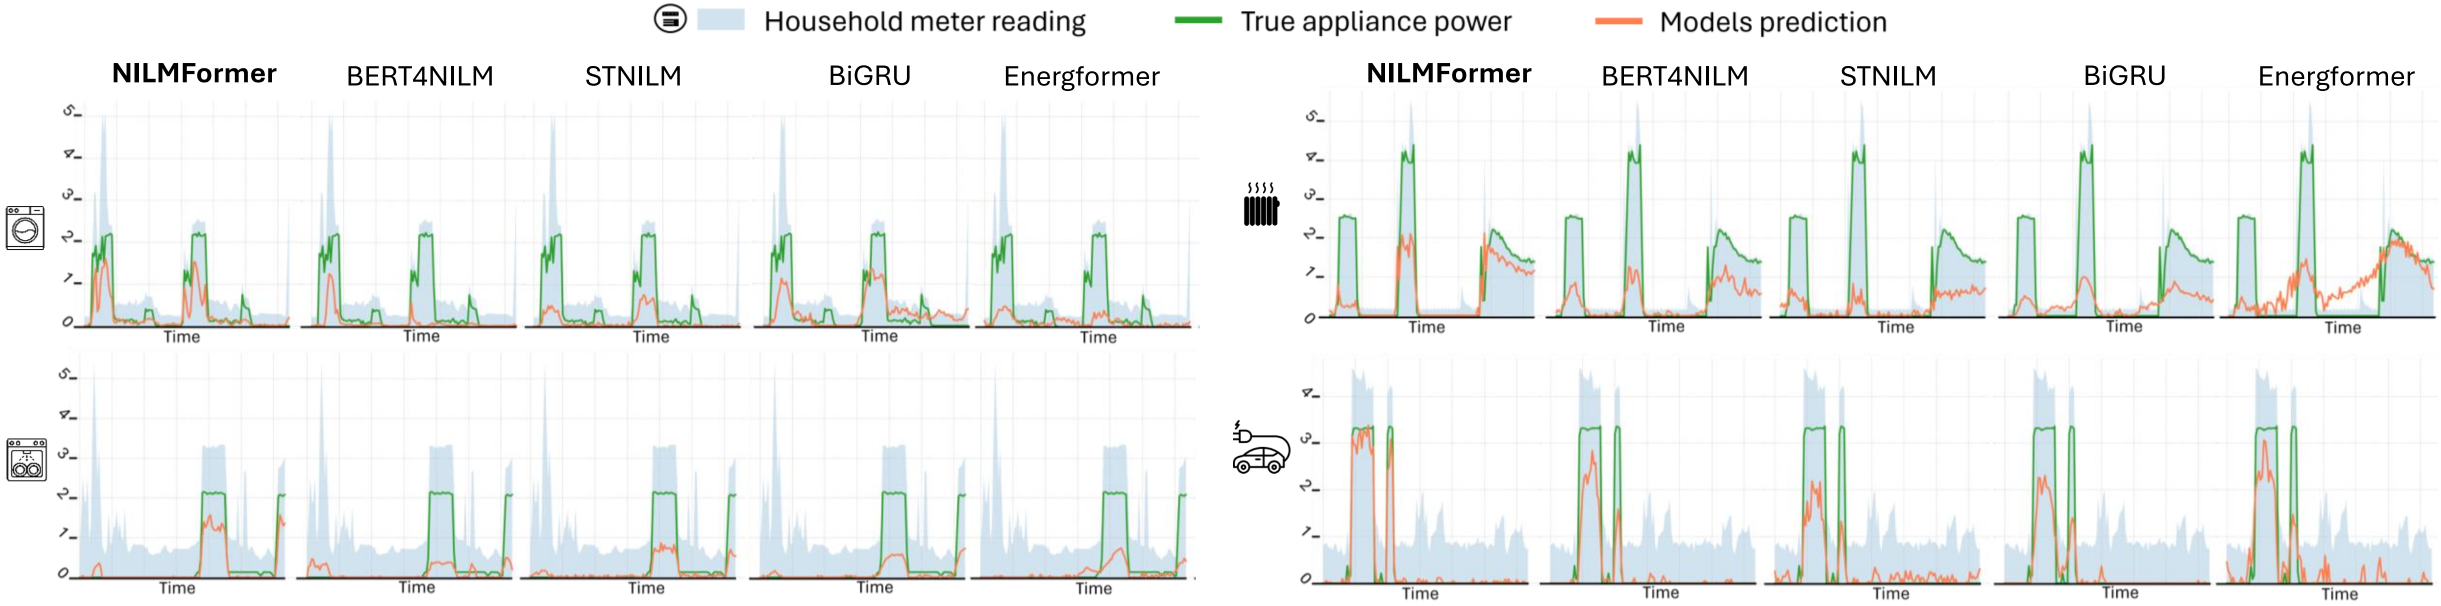

Public benchmark **EDF data**

A. Petralia et al., NILMFormer: Non-Intrusive Load Monitoring that Accounts for Non-Stationarity. Under review for KDD 2025.

Results: Per-timestamp Energy Disaggregation

Effects of proposed **Non-Stationary Mechanisms** on **NILMFormer Performance**

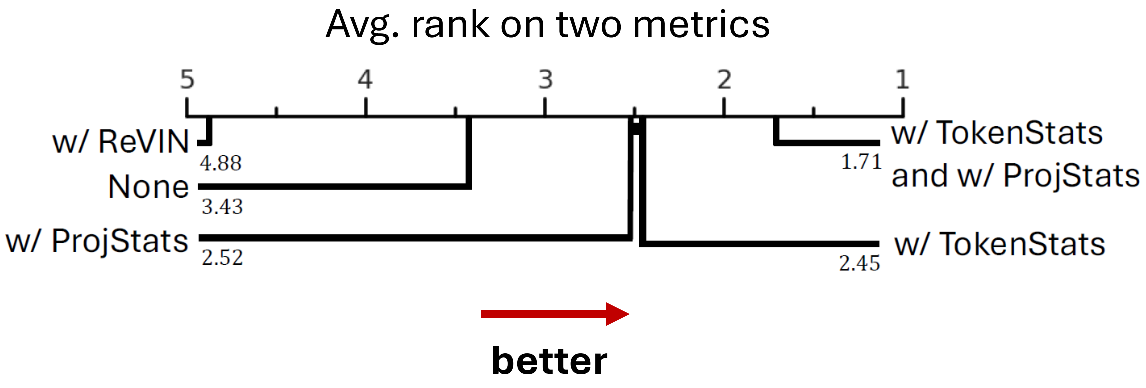

Averaged results over 4 datasets and 14 appliances disaggregation scenarios

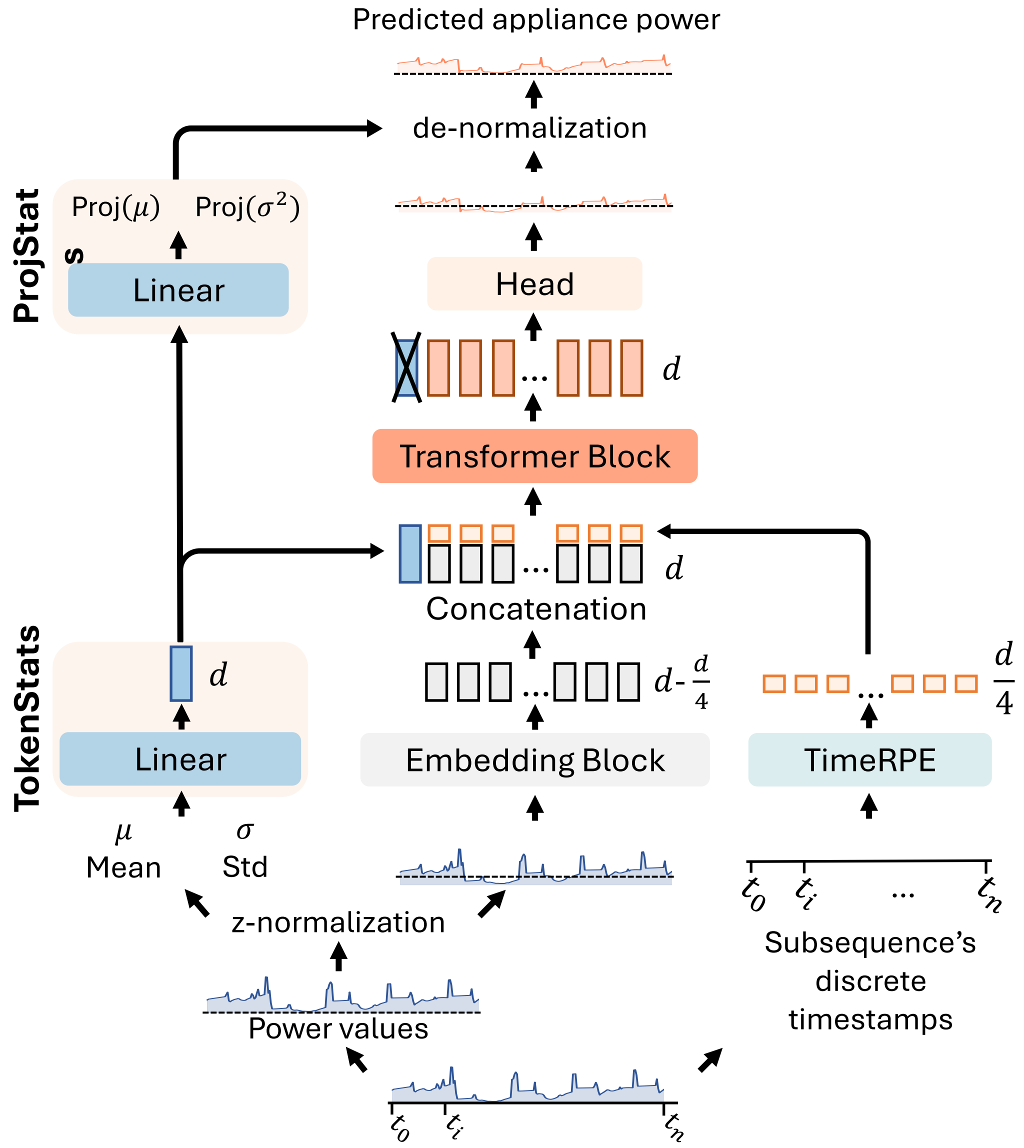

A. Petralia et al., NILMFormer: Non-Intrusive Load Monitoring that Accounts for Non-Stationarity. Under review for KDD 2025.

Results: Per-timestamp Energy Disaggregation

Effects of proposed **Non-Stationary Mechanisms** on **NILMFormer Performance**

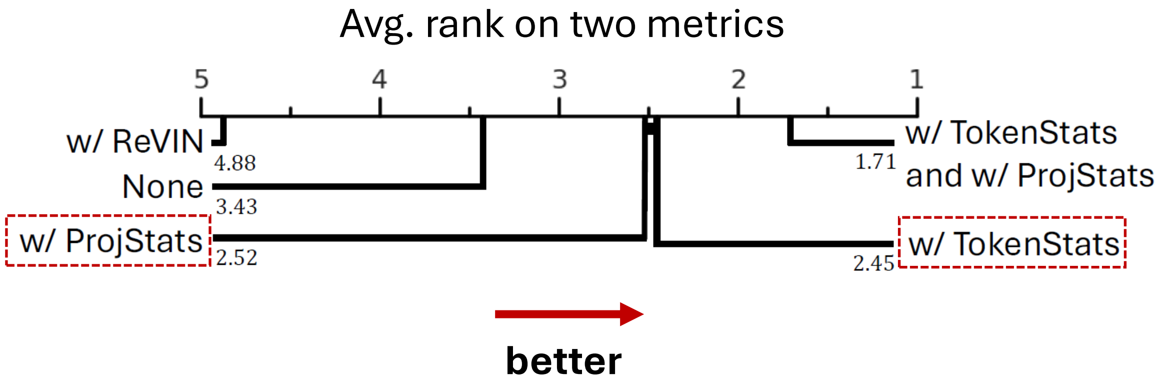

Averaged results over 4 datasets and 14 appliances disaggregation scenarios

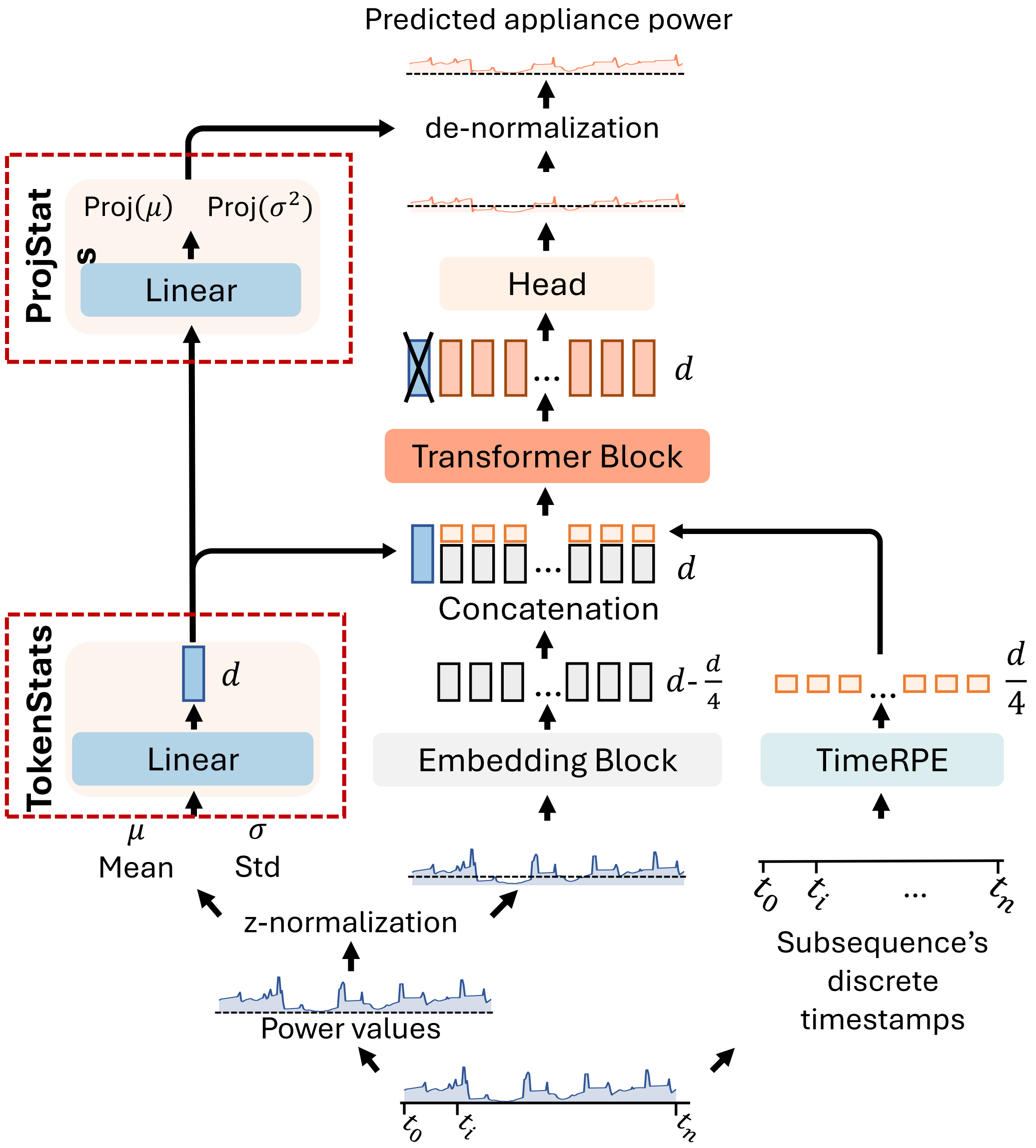

A. Petralia et al., NILMFormer: Non-Intrusive Load Monitoring that Accounts for Non-Stationarity. Under review for KDD 2025.

Deployed Solution: NILMFormer for Appliance Consumption Feedback

A Straightforward **Framework** for **Per-Period Energy Estimation**

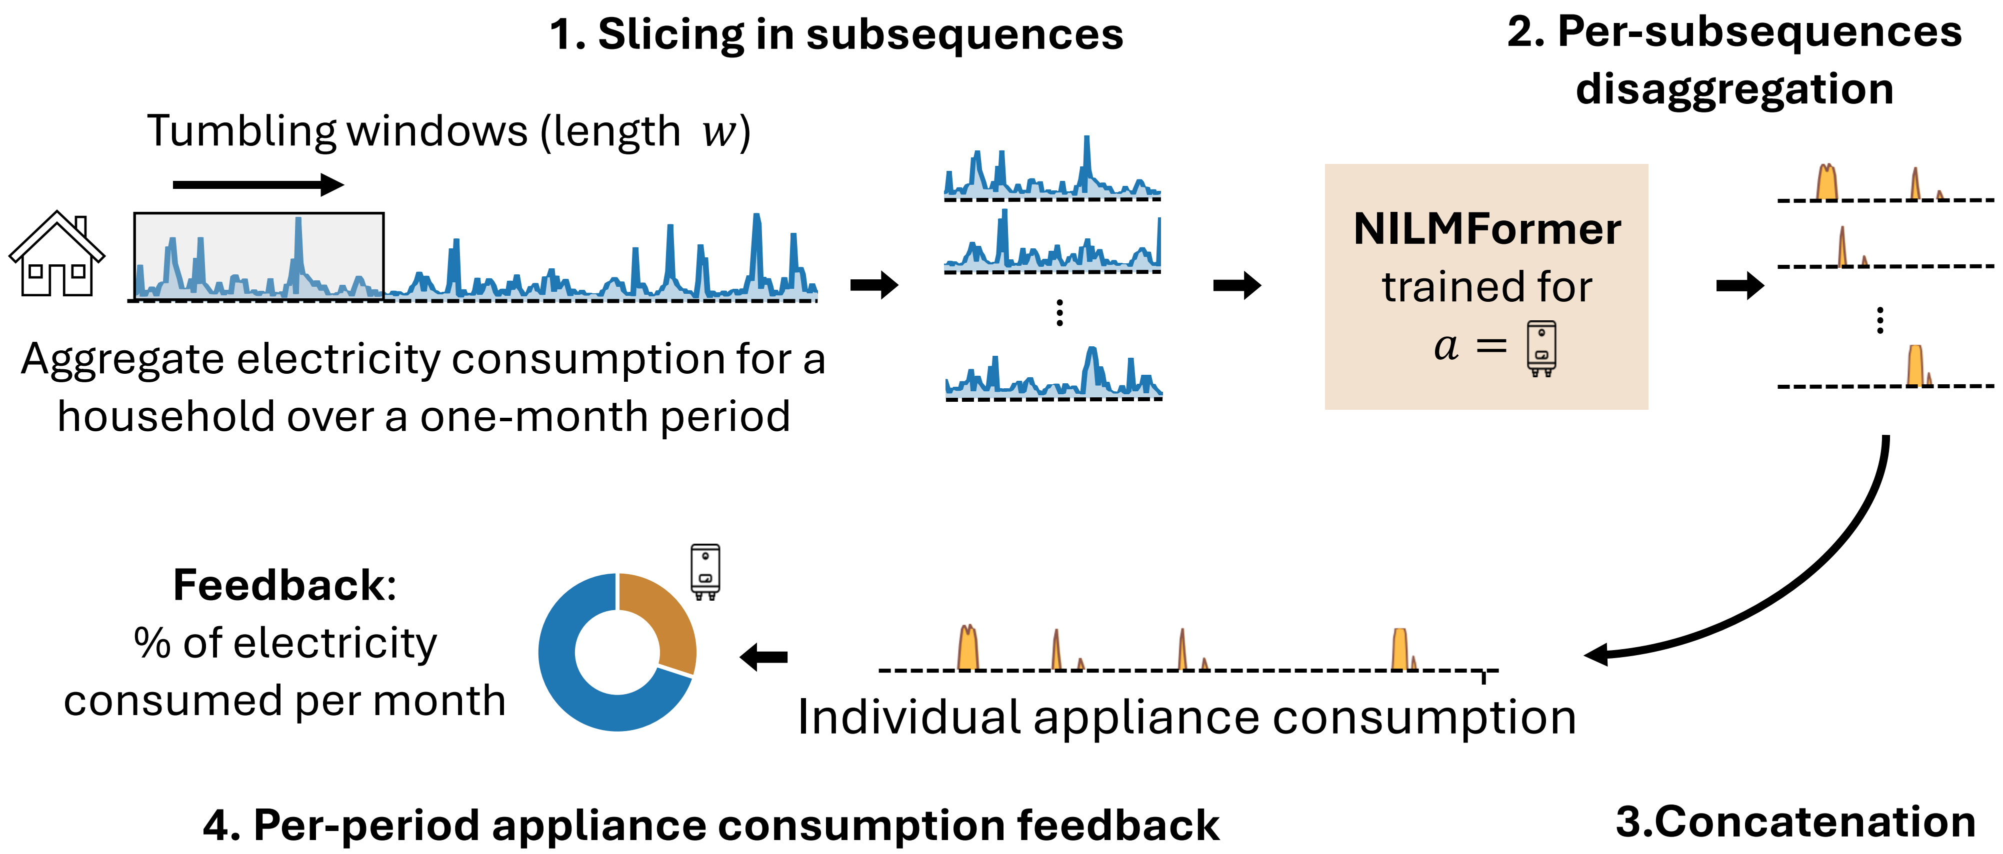

A. Petralia et al., NILMFormer: Non-Intrusive Load Monitoring that Accounts for Non-Stationarity. Under review for KDD 2025.

Deployed Solution: NILMFormer for Detailed Appliance Feedback

Performance Benchmark Against **TSER State-of-the-Art** (**EDF**'s Investigated Solution for *Mon Suivi Conso*)

Daily Power Appliance Estimation

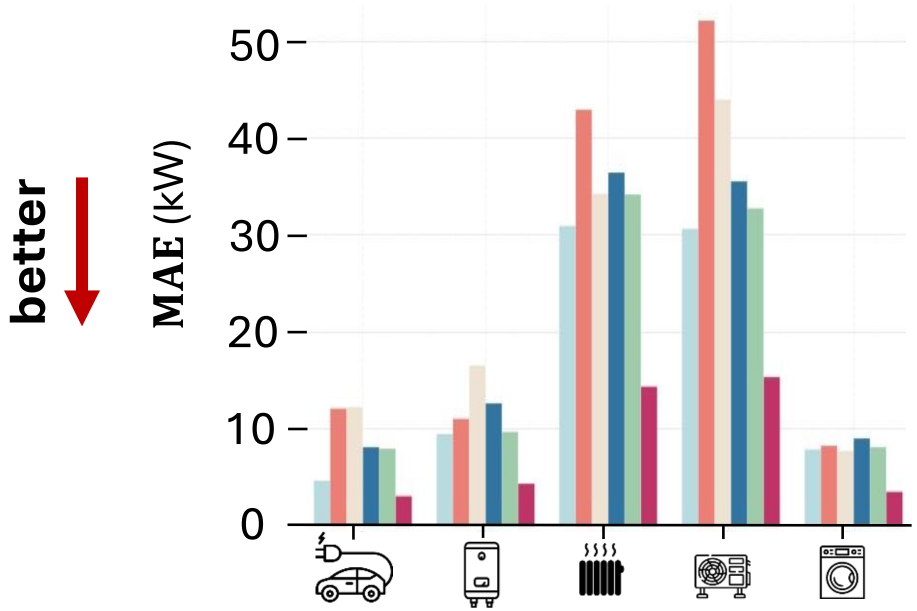

Monthly Power Appliance Estimation

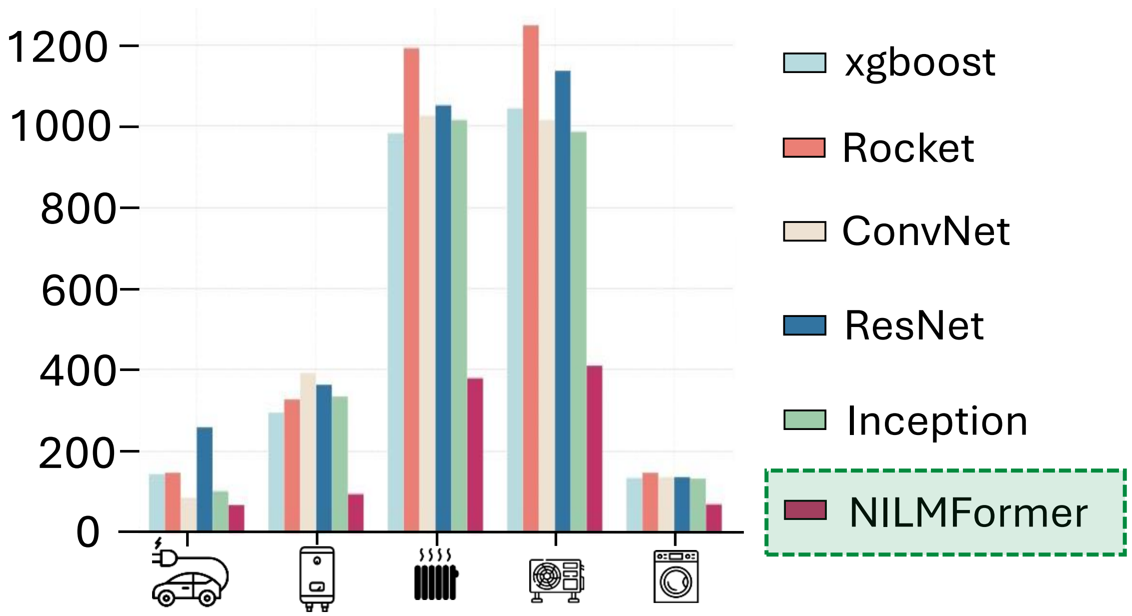

Achieves up to 52% lower error
than the 2nd-best baseline(XGBoost)

Achieves up to 151% lower error
than the 2nd-best baseline (Inception)

EV Water Heater Heater Heatpump White Appliances

A. Petralia et al., NILMFormer: Non-Intrusive Load Monitoring that Accounts for Non-Stationarity. Under review for KDD 2025.

Deployed Solution: NILMFormer for Detailed Appliance Feedback

Deployment of **NILMFormer** in *Mon Suivi Conso*

- **Already deployed** in two EDF subsidiaries: *Électricité de Strasbourg* and *EDF Solutions Solaires*
- **Progressive rollout** underway for the entire **EDF customer base** (*4M users*)
- Capable of processing the full EDF database in **11h**

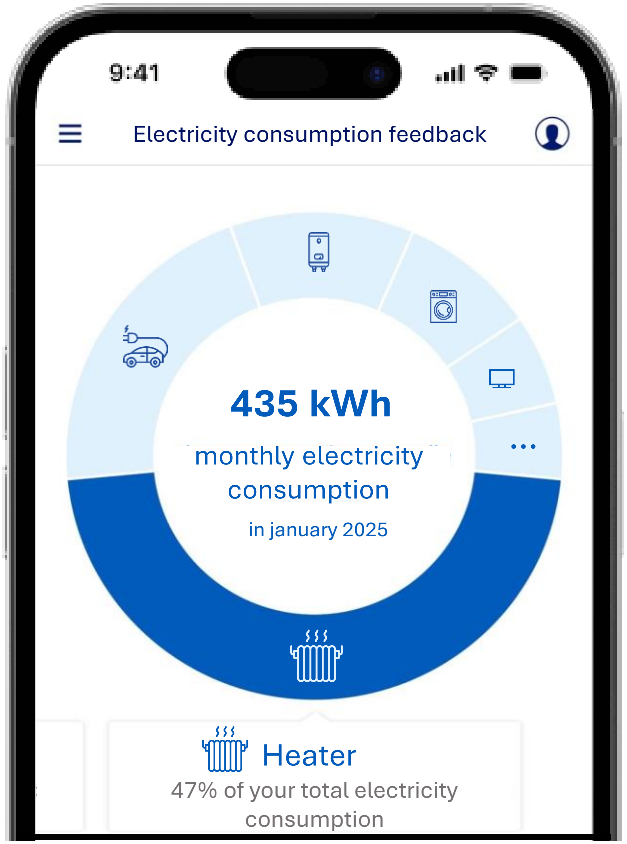

A. Petralia et al., NILMFormer: Non-Intrusive Load Monitoring that Accounts for Non-Stationarity. Under review for KDD 2025.

*How to provide detailed and accurate **fine-grained** appliance consumption **feedback** to customers?*

Challenges

1. **Considering non-stationary**
   Mitigating the data drift within each subsequence

2. **Delivering granular, actionable feedback to customers**
   Per-timestamp, daily, weekly and monthly

Solutions

- ✓ **NILMFormer**
  - ➢ Effectively takes into account the non-stationarity aspect of the data

- ✓ **Deployment** in ***Mon Suivi Conso***
  - ➢ Backbone algorithm for appliance feedback
  - ➢ Fine grain delivering insights

III. Conclusions

1. Appliance Detection

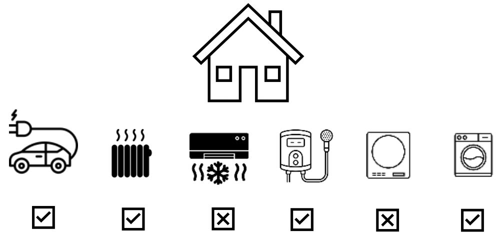

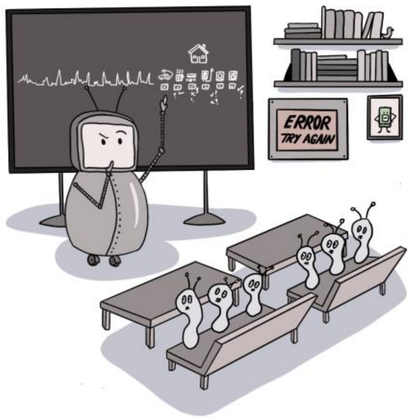

2. Appliance Pattern Localization

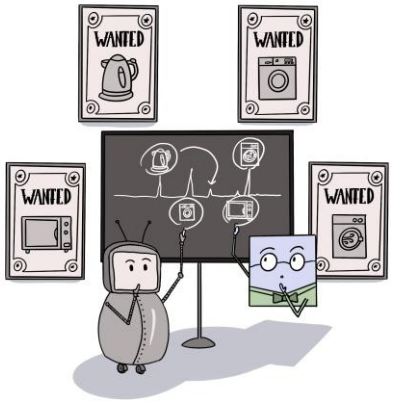

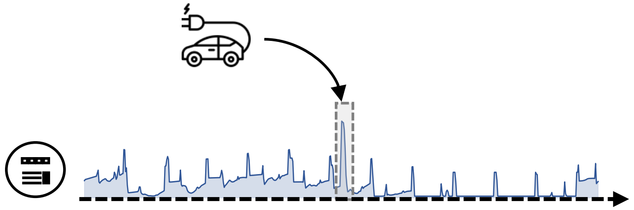

3. Energy Disaggregation

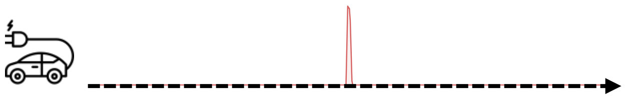

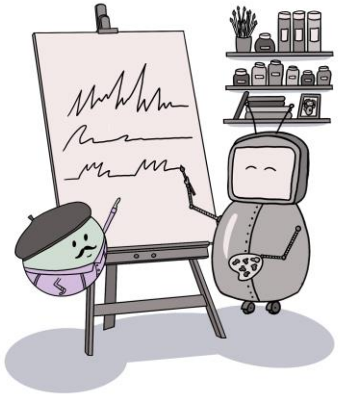

Conclusions

*Can we extract **relevant information** from **electricity consumption time series** collected by **common smart meter** at a **very low frequency**?*

Conclusions

*Can we extract **relevant information** from **electricity consumption time series** collected by **common smart meter** at a **very low frequency**?*

Yes

Conclusions

*Can we extract **relevant information** from **electricity consumption time series** collected by **common smart meter** at a **very low frequency**?*

Yes

Contributions

ADF&TransApp
for
Appliance Detection in Consumer Household

PVLDB, 2024
ACM e-Energy, 2023

CamAL
for
Weakly Supervised **Appliance Pattern Localization**

ICDE, 2025
(2 papers)

NILMFormer
for
Energy Disaggregation and Detailed Appliance Consumption **Feedback**

KDD, 2025

Open Research Directions

Toward General-Purpose **Foundation Models** for **Electricity Consumption** Analytics

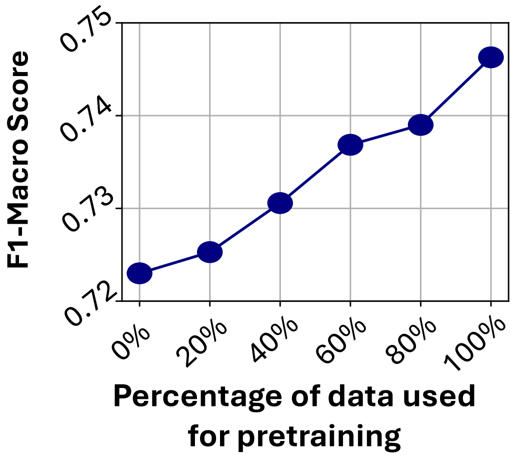

Pretext tasks and large-scale data ***significantly improve*** *TransApp's appliance detection accuracy*

Foundation Model
(pretrained on large amount of data)

Appliance Detection

NILM

Forecasting [1]

Clustering

Anomaly Detection

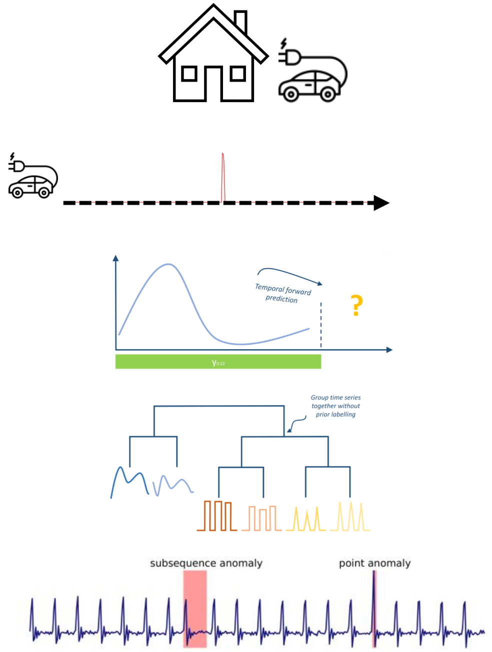

- Integration of **exogenous variables** (e.g., temperature, contractual metadata) into the learning process
- Model architecture resilient to **heterogeneous sampling frequencies**[2]

[1] G. Woo et al., Unified Training of Universal Time Series Forecasting Transformers, ICML, 2024
[2] E. Le Naour et al., Time Series Continuous Modeling for Imputation and Forecasting with Implicit Neural Representations, TMLR, 2024

Publications

List of publications related to this thesis

1. Adrien Petralia, et al. *NILMFormer: Non-Intrusive Load Monitoring that Accounts for Non-Stationarity.* **KDD, 2025**.
2. Adrien Petralia, et al. *Few Labels are all you need: A Weakly Supervised Framework for Appliance Localization in Smart-Meter Series.* **ICDE, 2025.**
3. Adrien Petralia, et al., *DeviceScope: An Interactive App to Detect and Localize Appliance Patterns in Electricity Consumption Time Series.* **ICDE, 2025.**
4. Adrien Petralia**.** *Time Series Analytics for Electricity Consumption Data.* **VLDB PhD Workshop, 2024.**
5. Adrien Petralia, et al. *ADF&TransApp: A Transformer-Based Framework for Appliance Detection Using Smart Meter Consumption Series.* **PVLDB, 2024**.
6. Adrien Petralia, et al. *Détection d'appareils dans les séries temporelles de compteurs intelligents très basse fréquence.* **BDA, 2023.**
7. Adrien Petralia, et al. *Appliance Detection Using Very Low-Frequency Smart Meter Time Series.* **ACM e-Energy, 2023.**

Patents

1. **Adrien Petralia**, Paul Boniol, Themis Palpanas, Philippe Charpentier. *Détermination d'une activation au cours du temps d'un équipement donné au sein d'un ensemble d'équipements à partir de données collectées.* **French Patent FR2504769**, **2025**.

2. **Adrien Petralia**, Themis Palpanas, Philippe Charpentier, Claire Lambert. *Extraction de la consommation électrique d'un équipement individuel au sein d'un ensemble d'équipements connectes a un réseau électrique.* **French Patent FR2410061**, **2024**.

3. Luc Dufour, Pascal Chaussumier, **Adrien Petralia,** Philippe Charpentier, Themis Palpanas, Justin Capik. *Caractérisation pour l'optimisation de la consommation électrique d'un ensemble d'équipements connectes a un réseau électrique.* **French Patent FR2314217**, **2023**.

Thank you for your attention!

Deep Learning for Electricity Consumption Time Series Analytics

Adrien Petralia

supervised by *Prof. Themis Palpanas* and *Philippe Charpentier*

May 7th, 2025. UPCité & Online.