SUSTAINABLE PLACES 2020

October 27-30, 2020

Measurement and verification approaches for performance-based models of energy efficiency services

HEBES Intelligence / Sotiris Papadelis

This project has received funding from the **European Union's Horizon 2020 Research** and **Innovation programme** under Grant Agreement No **847066**

CONTENT

1. The SENSEI H2020 project
2. Measurement and Verification (M&V) in SENSEI
3. The M&V related contributions of SENSEI so far
4. The M&V pre-processing stage
   i. The day typing approach
   ii. Using the day typing information
5. The M&V predictive stage
6. Sources of additional information
7. Comments or questions

sensei

The SENSEI H2020 project (1/6)

SENSEI: Smart Energy Services to Improve the Energy Efficiency of the European Building Stock

Duration: 1 September 2019 to 31 August 2022

Website: https://senseih2020.eu/

The SENSEI H2020 project (2/6)

Goal: The overall goal of the SENSEI project is to propose services that allow energy efficiency to be treated as a transactable resource, as well as business models that utilize these services in order to valorize energy efficiency as a power grid resource.

To this end, the pay-for-performance (P4P) concept has been adopted by SENSEI as a way to bridge the aforementioned services and business models, and define the transactions between two or more involved parties.

Performance-based agreements are already part of the ESCO model. SENSEI focuses on designing programs that support energy efficiency by compensating (paying) energy efficiency retrofits based on the energy savings they actually deliver (performance).

The SENSEI H2020 project (3/6)

SENSEI has identified three (3) different ways through which energy efficiency may be linked to the local power grid and its operational goals and challenges. These ways differ in the degree of involvement / endorsement that is required from the power grid's main actors: regulatory authorities for energy, power system operators and utilities. The greater the need for involvement / endorsement, the shorter the distance between an energy efficiency project and the power grid in the diagram below:

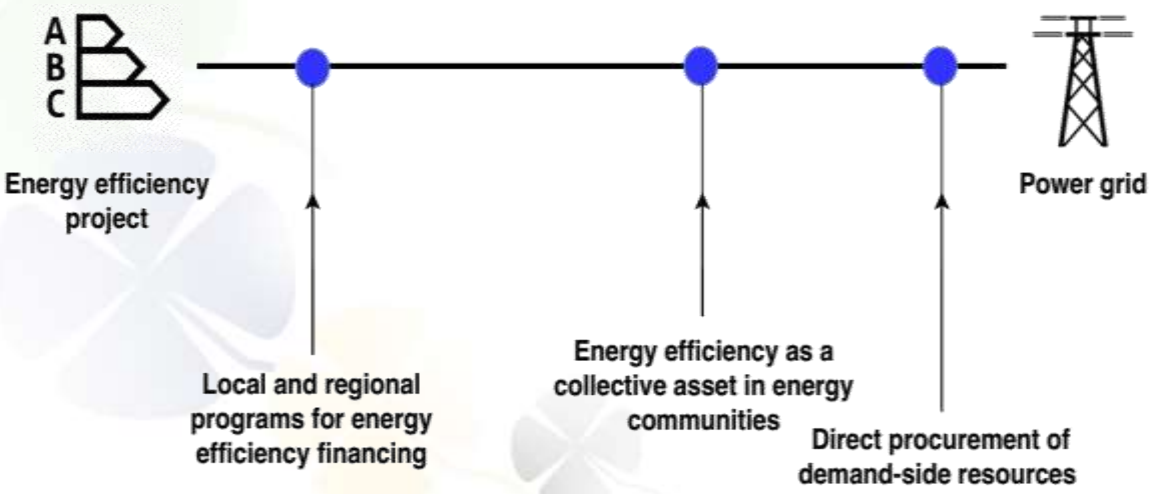

The SENSEI H2020 project (4/6)

The paradigm of direct procurement of demand-side resources: Load-modifiers are those resources or programs not seen or optimized by the energy or capacity market, but they modify the fundamental system load shape, preferably in ways that harmonize with the system operator's grid operations.

Examples of load-modifiers are dynamic rates and energy efficiency programs. An effective load modifying program helps create a flatter system load profile, attenuating high energy peaks and valleys and reducing extreme upward and downward ramps.

Load-modifiers are compensated through programs that identify and support the measures that are most effective in offsetting the need for new generating plants or transmission upgrades.

The SENSEI H2020 project (5/6)

Energy efficiency as a collective asset in energy communities: The term ‘generation’ is defined in the Electricity Market Directive (EMDII) as ‘production of electricity’. Most often, the primary activity is the production of electricity from a renewable energy plant, and the main objective of the community is the distribution among its members of the revenues accruing from the operation of a renewable generation project.

SENSEI aims at promoting the idea that energy efficiency can also be part of the operations of an energy community, and at leveling the field for energy efficiency improvements to be financed and valorized in the same way that generation assets are financed and valorized in energy communities.

sensei

The SENSEI H2020 project (6/6)

Local and regional programs for energy efficiency financing: Engage public institutions with a budget and mission to promote energy efficiency, and facilitate the design of energy efficiency programs that link compensation with specific performance indicators.

These indicators may incorporate a variety of goals in a P4P scheme: produce energy savings, mitigate energy poverty, promote new technologies, align energy consumption with the needs of the local energy systems (supply and distribution constraints), or increase capacity for new energy consumption trends (such as electric vehicles).

sensei

Measurement and Verification in SENSEI (1/4)

In all cases, the main aspects that are expected to define the business model(s) proposed by SENSEI include:

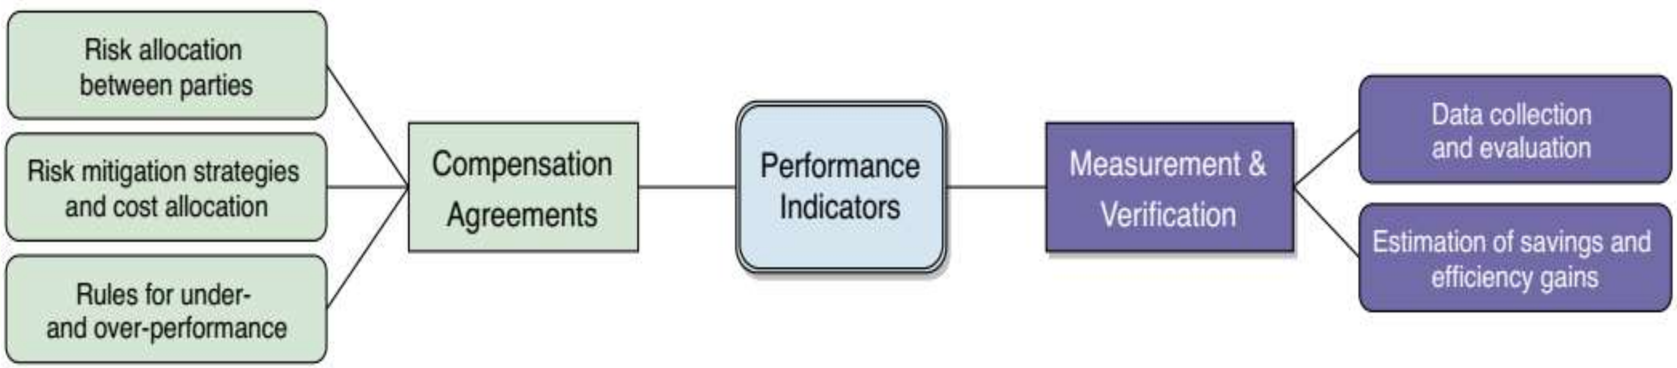

Measurement and Verification in SENSEI (2/4)

The current paradigm for measurement and verification (M&V) deviates from the way outcome-based models typically work.

In particular, the typical approach of defining a baseline consumption (before the implementation of the energy efficiency measures) and linking payments to the estimated savings means that any increase in consumption would either lead to a dispute between the consumer and the provider or be an indication of the deterioration of the installed energy efficiency measures.

In contrast, for all other outcome-based models, an increase in use is an increase in added value (more use of a service or asset).

Measurement and Verification in SENSEI (3/4)

A useful example to understand the limitations of the standard approach to M&V can be found when dealing with the impacts of long lasting non-routine events like the COVID-19 pandemic.

For a building that is used less due to measures to address the pandemic, the pre-retrofit baseline becomes irrelevant and the involved parties must renegotiate what part of the reduction in energy consumption is due to efficiency and what part is due to the change in use.

Measurement and Verification in SENSEI (4/4)

An alternative approach is to link payments to efficiency gains and the degree to which these gains are "used" by the consumer. In this case, the only thing that matters is finding a minimum surcharge that should be paid to the provider for making the efficiency gains available in the first place.

This reflects a notion that energy efficiency upgrades are an option that allows for improved monitoring, smarter control and/or reduced energy consumption. This option is associated with a cost to obtain it in the first place and an expected added value from utilizing it during the day-to-day operation of the building.

The M&V related contributions of SENSEI so far

Starting point: It is assumed that energy consumption and outdoor air temperature is the only data that is available for M&V, since this data corresponds to the minimum set of data that may be available across all buildings in a portfolio and that would allow for a consistent analysis.

Goal: Utilize the available data so that to understand the dynamics of a building's energy consumption:

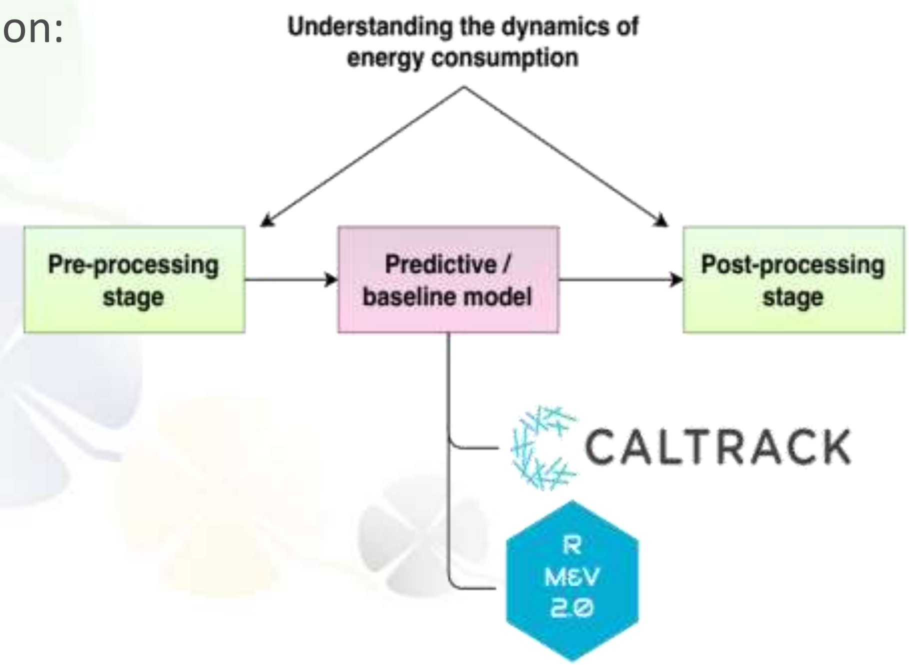

The M&V pre-processing stage (1/2)

The preprocessing stage is responsible for the following tasks:

1. Evaluate the adequacy of the available energy consumption and outdoor air temperature data.
2. Define a benchmark for the baseline predictive model. Having a benchmark for the baseline model allows us to distinguish early on between easy and difficult to predict observations in the dataset. The benefits from this distinction are twofold:
   a) We have a target performance for the baseline model.
   b) The difficult to predict observations are the most interesting instances for auditing the baseline model. Auditing the model means identifying in a transparent and interpretable way what the model has learned from the data.

The M&V pre-processing stage (2/2)

3. Identify common and uncommon patterns in daily energy consumption (day typing). The general rule is to not allow the baseline predictive model to learn patterns, the persistence of which cannot be verified with the available data.
4. Define the locality boundaries for each instance in the dataset. Auditing the baseline model requires the development of an interpretable local model around a specific observation. Accordingly, we need a way to define the neighborhood of each observation in the dataset.
5. Develop the predictive features for the baseline model.

The day typing approach (1/12)

Example dataset: The building with *SiteId=50* in the open dataset of building electricity consumption that Schneider Electric has made available on their Data Exchange:

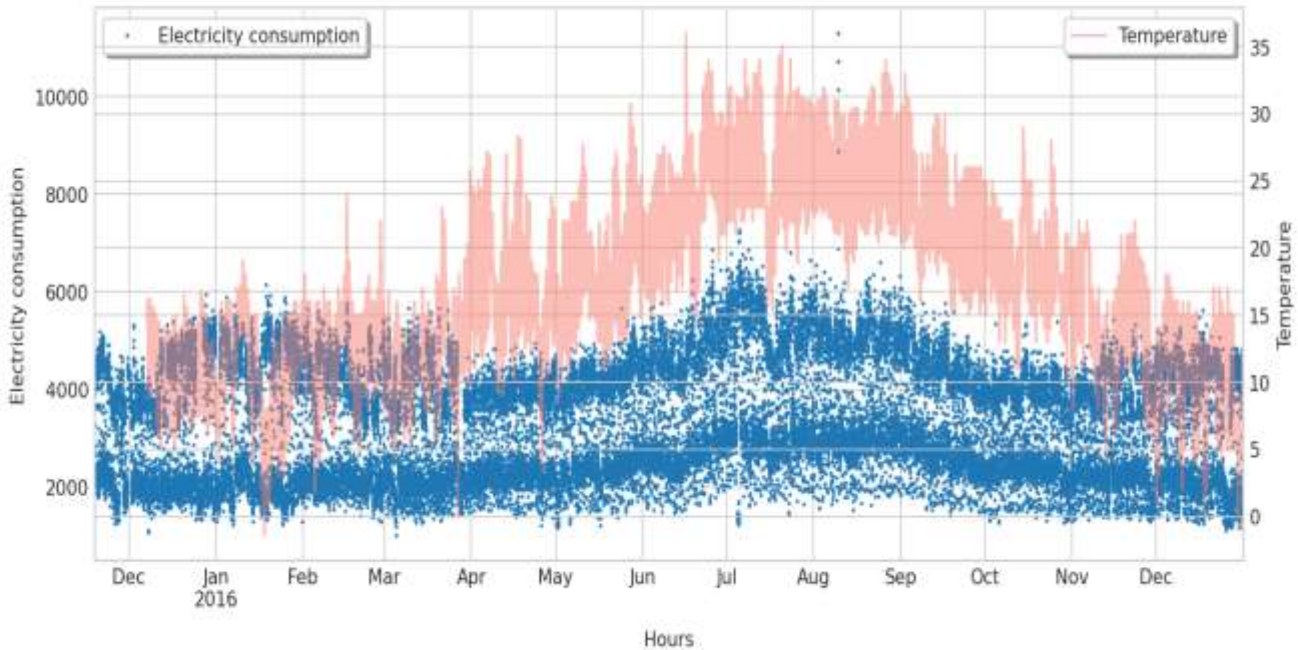

The day typing approach (2/12)

The proposed method builds upon the distance and matrix profile data structures (https://matrixprofile.org/).

A **distance profile** between a time series and a specific subsequence of length *m* is another time series that stores the (normalized) Euclidean distance between the query subsequence and each possible subsequence of length *m* in the original time series. By definition, the distance profile takes values close to zero at the location of other subsequences that are very similar to the query one.

A **matrix profile** of a time series is another time series that stores at each position the distance between the subsequence that starts at that position and its nearest neighbor. If a subsequence has a matrix profile value far greater than zero, it is unlike any other subsequence in the dataset, whereas if it has a close to zero value, it is a pattern.

The day typing approach (3/12)

What we are searching for:

1. **Discords**. A discord is a daily profile that is maximally different to all the other daily profiles in the dataset. In some cases, it is because of outliers in metered data. In other cases, the profiles correspond to national holidays.
2. **Infrequently recurring profiles**. This category concerns profiles that appear more than once (making them a pattern), but only for a small number of times. Because they are recurring, these patterns can be learned by the baseline model. However, it is not always desirable to allow this to happen.
3. **Frequently recurring profiles**. This category concerns profiles that appear often and, as a result, we want the baseline model to exploit them for its predictions.

The day typing approach (4/12)

Identification of discord days based on their electricity consumption profiles: The matrix profile allows us to turn the problem of identifying unusual daily profiles into indentifying outlier points in a series of distance values. The following diagram presents the matrix profile values of all the subsequences in the building's consumption data that: (a) have length *m* that corresponds to one day (this means that *m=24* for hourly data or *m=96* for 15-min data), and (b) start at 00:00 hours, so that to only compare subsequences that span a full day's period.

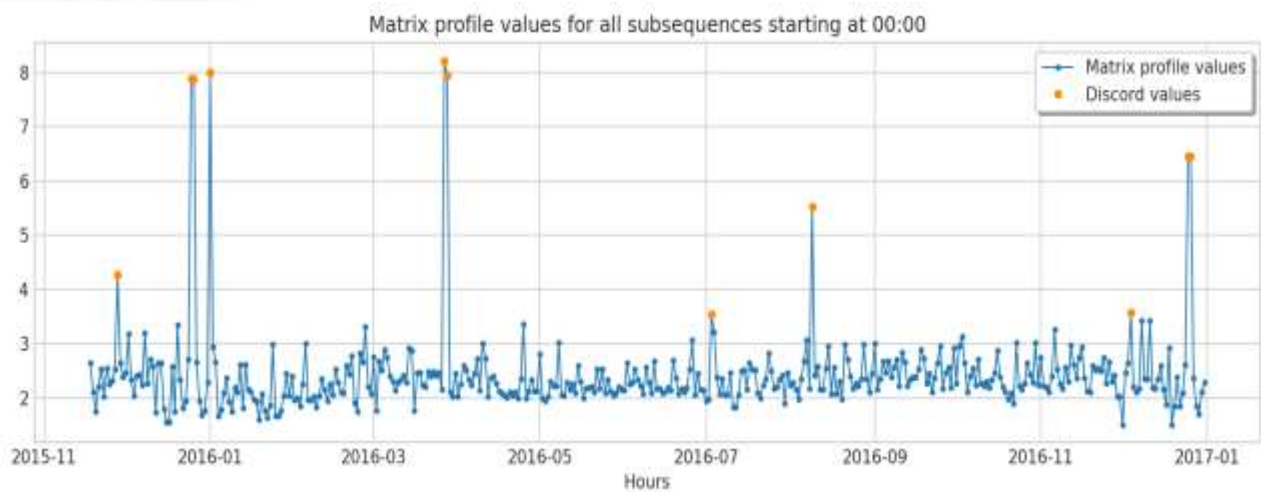

The day typing approach (5/12)

The daily consumption profiles of all the identified discord days:

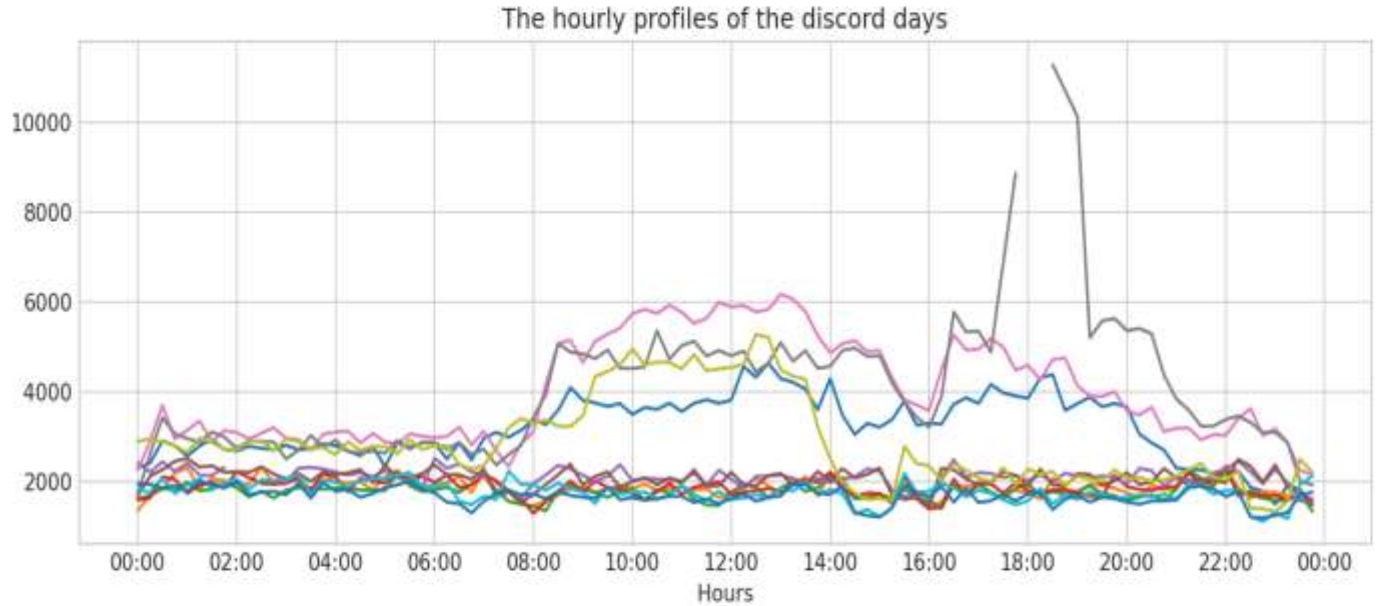

The day typing approach (6/12)

The consumption profiles of all the days that haven't been identified as discords:

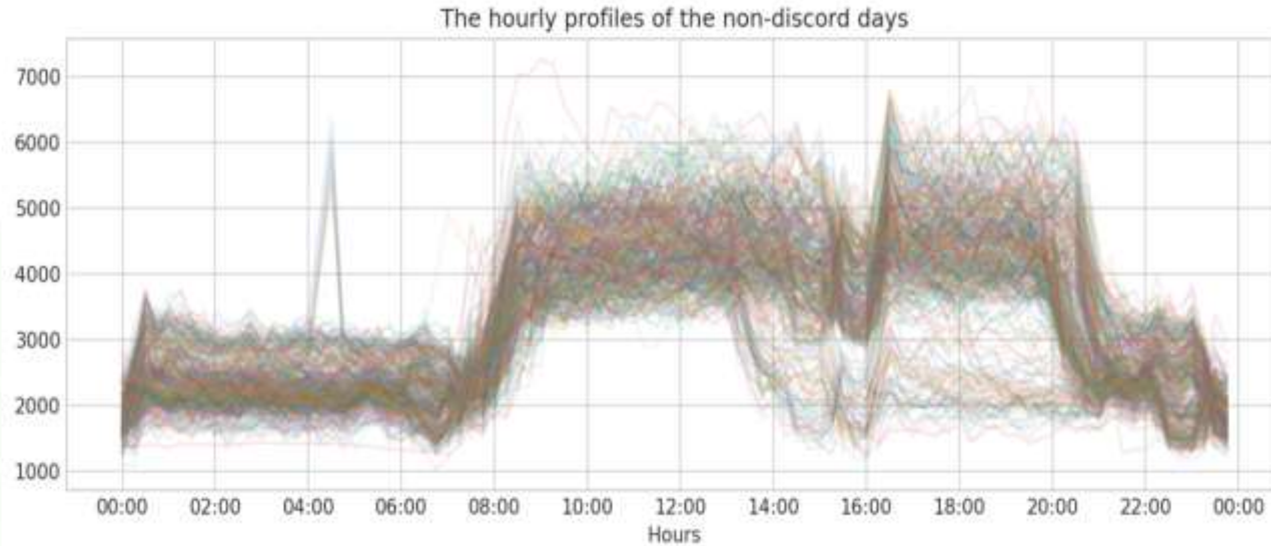

It is possible to visually spot at least three different daily patterns: one where the high-load period spans from 09:00 to 19:00, one where the high-load period spans from 09:00 to 13:00, and one that includes an unusual peak during the early morning hours.

The day typing approach (7/12)

Categorization of the non-discord days based on their electricity consumption profiles: The key idea behind the proposed method is to identify a small number of recurring patterns that are very dissimilar to each other and can be used as points of reference for comparing all the daily consumption profiles found in the dataset.

Instead of working on the complete daily profiles, we first split the electricity consumption of the non-discord days into three (3) non-overlapping 8-hour intervals: 00:00-08:00, 08:00-16:00 and 16:00-00:00, and then apply the proposed method on each and every interval.

sensei

The day typing approach (8/12)

The process begins with calculating the matrix profile of the electricity consumption data for a subsequence length *m* that corresponds to 8 hours, and isolating the matrix profile values of all the subsequences that start at 00:00 hours. Then, the reference patterns are identified through the sequence outlined next:

1. **Find the minimum value in the matrix profile**. This is the starting index of the first reference pattern.

2. **Compute the distance profile of the selected reference pattern with respect to all the subsequences in the time interval** under study. The result measures how similar each 00:00-08:00 subsequence in the time series is to the currently selected pattern. The fact that we only care about subsequences that start at 00:00 hours protects us from trivial matching.

The day typing approach (9/12)

3. **Store the distance profile** of the first selected pattern into an array (let’s call it *Z*).

4. **Compute a relative matrix profile** by dividing the original matrix profile data with the distance profile of the previous step.

5. **Find the minimum value in the relative matrix profile**. This is the starting index of the second reference pattern. The intuition is that if a subsequence has a very close nearest neighbor (i.e. a very low matrix profile value), but is very different from the selected reference pattern, then its value in the relative matrix profile should be also very low because we divide an already small number by a large one.

The day typing approach (10/12)

6. **Compute the distance profile** of the new reference pattern with respect to the electricity consumption time series.

7. **Update the *Z* array** with the element-wise minimum between *Z* and the distance profile of the previous step. In other words, we use *Z* to store the distance between every subsequence and its closest match among all the patterns selected so far.

8. **Recalculate the relative matrix profile** by dividing the original matrix profile data with the updated *Z* array.

The day typing approach (11/12)

In order to control the number of iterations, we monitor how the minimum value of the relative matrix profile evolves, and terminate the process when it fails to increase by more than 25% (a number that we have found empirically that works well for filtering out redundant patterns).

This would result into the following reference patterns:

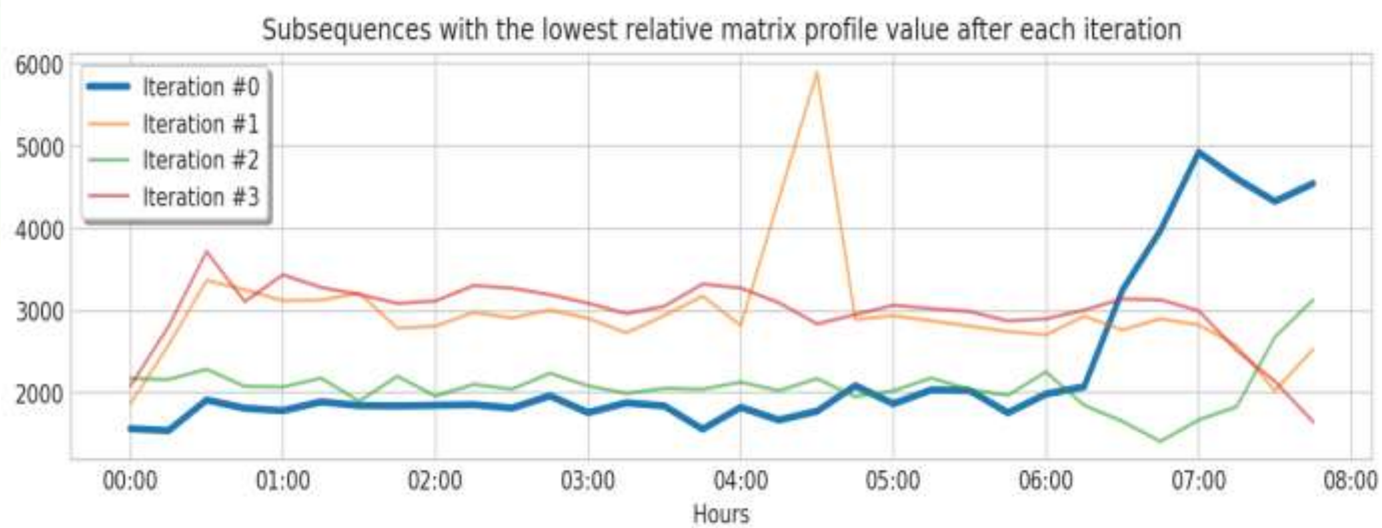

sensei

The day typing approach (12/12)

Finally, we classify all subsequences based on their similarity to the reference patterns:

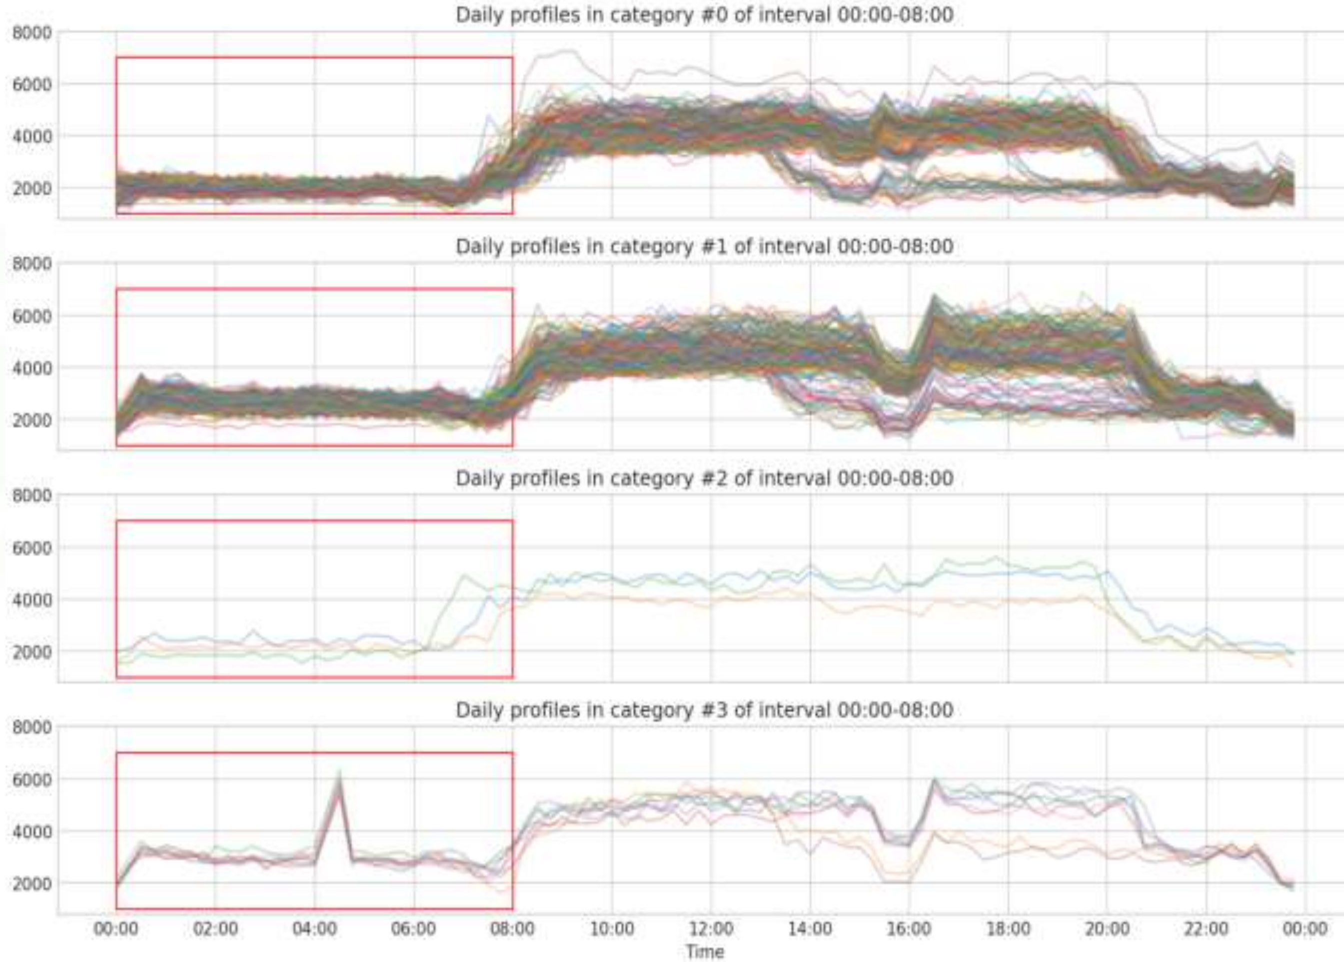

Using the day typing information (1/4)

Decide on what the baseline model should see: The whole pre-processing stage is about tagging the available data.

The predictive model fitting part is concerned with minimizing regret by comparing the effect of hiding some data when the model should in fact see it with the effect of allowing the model to see this data when it should not.

Defining regret is very much related to the periods during which we care most about the predictive model's accuracy. If the savings are equally valuable during all hours of the year, the average performance of a model is all we need because we can sum up all over- and under-estimations.

Using the day typing information (2/4)

However, SENSEI is concerned with the value of an energy retrofit project to the grid, and this value depends on the operational characteristics of the grid, as well as on the time periods during which the consumption changes from the project occur.

Although different power grids face different challenges, in most cases a system operator would prefer energy retrofit projects that specifically target the hours when the probability of load loss is high and/or the hours when known and persistent variability in the net load leads to ramping events, i.e. large changes in the magnitude of the net load lasting for a few hours.

sensei

Using the day typing information (3/4)

Find the neighborhood around each observation: Day typing helps define the neighborhood of an hourly observation across days (same hour, different days).

This is supplemented by the fact that the benchmark baseline model derives its predictions from grouping together hours that share the same behavior in terms of how consumption varies with outdoor temperature:

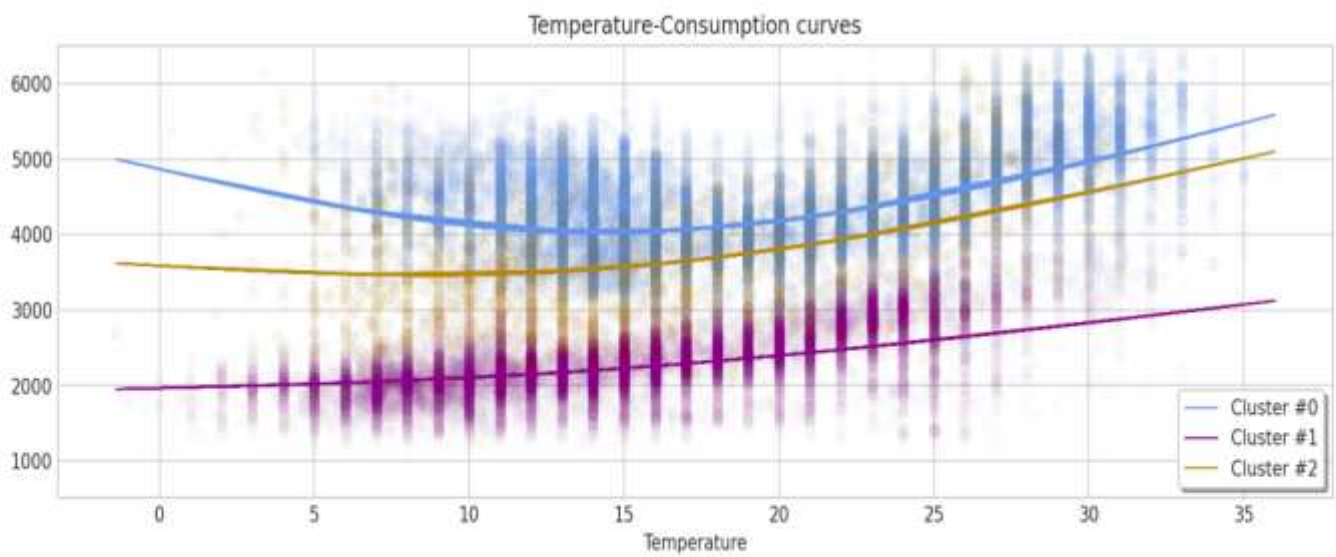

Using the day typing information (4/4)

This helps define the neighborhood of an hourly observation across both hours and days:

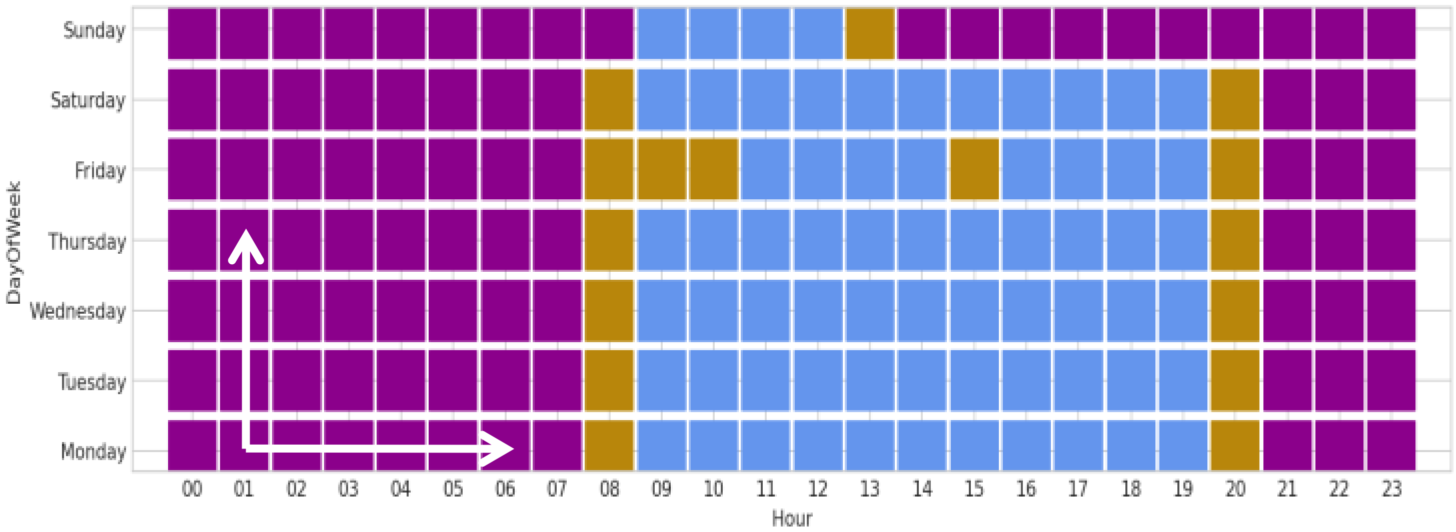

The M&V predictive stage

While the existing M&V protocols do not offer specific guidelines on how to deal with the aforementioned sources of uncertainty, one can find practical solutions in the machine learning literature. In particular, bootstrap aggregation (commonly referred to as bagging), is a technique that is well established in the machine learning community as a way to deal with model uncertainty (Jackknife+-after-bootstrap).

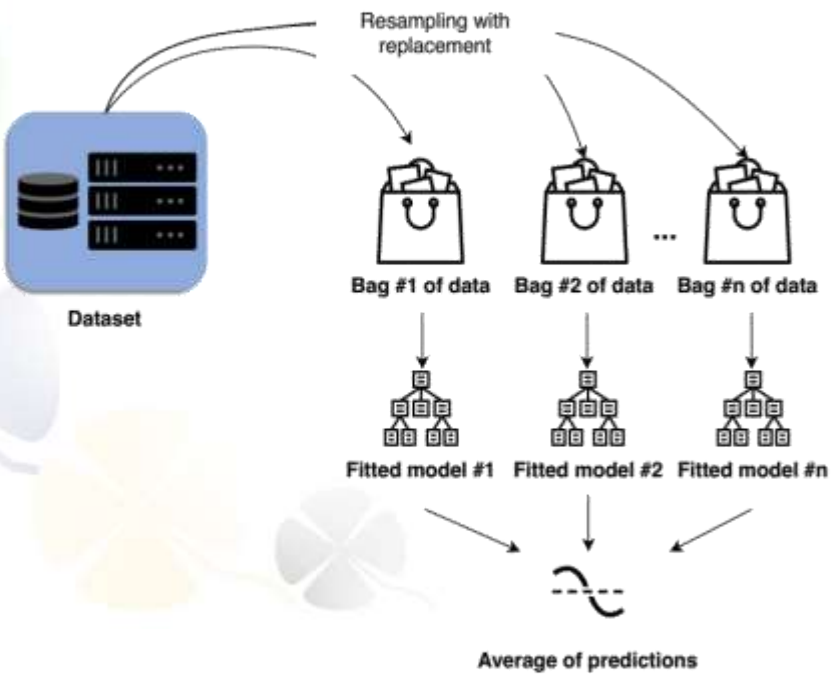

Sources of additional information

The SENSEI **deliverable D7.1** “Methods for the dynamic measurement and verification of energy savings” is expected to published to the project’s website early November 2020.

The deliverable describes all workflows for the pre- and post-processing stages in detail.

The deliverable will be accompanied by an open-source **software library** that implements all workflows for the pre- and post-processing stages.

The work on M&V in SENSEI is ongoing and new material will be published even after D7.1 gets published. Join our **Stakeholder Community** and stay informed!

sensei

Any questions or comments?

https://senseih2020.eu/

Sotiris Papadelis

spapadelis@hebes.io

HEBES Intelligence

Join our Stakeholder Community!
https://senseih2020.eu/stakeholders/

This project has received funding from the **European Union's Horizon 2020 Research** and **Innovation programme** under Grant Agreement No **847066**