VENICE WP2c

Modelling Approach & Findings

Client:	National Grid Electricity Distribution
Client Ref.:	N/A
Date:	January 2023
Classification:	FRAZER-NASH IN CONFIDENCE

Project No.:	010837	**Compiled By:**	Rob Wadsworth
Document No.:	54476R	**Verified By:**	Zoe Hodgins
Issue No.:	Draft 2	**Approved By:**	Mark Joyce
		Signed:	

Distribution

Copy	Recipient	Organisation
1	Liza Troshka	National Grid Electricity Distribution
2	Project file	Frazer-Nash Consultancy
3		
4		
5		

1 Overview

1.1 Project overview

Distribution Network Operators (DNOs) offer support to vulnerable households in the form of ongoing assistance and advise, and the provision of priority services during an interruption to supply. As DNOs do not own the relationship with customers, it is difficult to identify these households and provide them with this assistance. Frazer-Nash were tasked with determining whether it was possible to identify a vulnerable household from its smart meter data.

Smart meters are an essential digital upgrade to our energy system. Smart meters work by measuring the electrical current flow and voltage at regular intervals and adding this up to calculate the power used in a half-hour period. The uptake of smart meters across the UK will provide the foundation for a more sophisticated, green, and consumer-friendly energy system. The Smart Meter Data Communications Centre (Smart DCC) provide the communications infrastructure that handles smart meter data and is responsible for all access to smart meter data [1].

A household is considered vulnerable for a variety of reasons and each of these can represent themselves differently in smart meter data. Work completed prior to this report identified a selection of behavioural characteristics that may be present in a household's smart meter data if they have a specific vulnerability. It is thought that DNOs could use smart meters to determine if a household may be considered vulnerable. If possible, a model could be developed to analyse smart meter data from anonymised households. If a house is identified, DNOs could request access to the personal information from the Smart DCC. This would benefit the consumer by ensuring they receive the support they require. It would also benefit DNOs by reducing the burden of maintaining the PSR and it would ensure the full social benefit of the smart meter roll-out is realised. However, providing DNOs with access to consumer smart meter data would require changes to legislation put in place.

Project VENICE (Vulnerability and Energy Networks Identification and Consumption Evaluation) is a network innovation project funded through Ofgem's Network Innovation Allowance (NIA), supported by National Grid Electricity Distribution (previously known as Western Power Distribution). The project started in August 2021 and is using a three-pronged approach to support vulnerable customers in the future. The Frazer-Nash work package, WP2, explores whether smart meter data can be used to identify vulnerable households. The work to complete this was separated into WP2a, 2b and 2c. The findings from WP2b are detailed in [2], and this report presents the findings from WP2c.

1.2 WP2b findings

The previous phase of the VENICE project, WP2b, was undertaken to review three different approaches to detection of vulnerability patterns in electrical energy usage data. Each model was developed and tested in isolation using different, open source, datasets. Whilst the open source datasets were used to develop each model, none were validated using real households with known vulnerabilities. This was due to the restrictions in place which limit smart meter data access to DNOs and the public. The modelling approaches and findings are summarised as:

- **The cohort comparison model:** A model using a broader set of characteristic data to map vulnerability characteristics across neighbourhoods to the overall energy usage. The model used comparative usage to determine vulnerabilities related to expected usage, for example, if a household is using less electricity than anticipated for its building type and occupancy. The results showed that some of the variation in a household's average usage could be attributed to the vulnerability characteristics, however there were large uncertainties due to the proxy datasets used for the individual household characteristics and usages. It was concluded that to develop the model further, more testing must be undertaken with real average household usage data for homes with known characteristics.

- **The appliance disaggregation and prediction model**: A model using half hourly energy usage data, similar in format to smart meter outputs to identify appliance related vulnerabilities. For example, the use of medical equipment. The model estimated the likelihood of certain appliances being used within the household. The key model limitations are driven by high uncertainty in the power drawn by an appliance, as expected. These conclusions show that the approach has shown potential, but further development should be undertaken to reduce the uncertainties, in particular, a key limiting factor was the 30 minute aggregation of data; with higher resolution datasets the uncertainty could be greatly reduced. Once this is completed, the model should be tested with medical appliances.
- **The change detection model:** A model used to identify vulnerabilities related to a change in a household's circumstances. The distribution of energy usage for each half hour over a training period was analysed and a change was detected if this distribution changed significantly. This was completed using data from UK Power Network's Low Carbon London project [3]. The changes detected by this model could be representative of the onset of a vulnerability such as job loss or entering fuel poverty. The usage model concluded that step changes in usage could be identified and quantified and this was tested by artificially adding vulnerabilities into the smart meter data. If this model were to be developed further, an improved model of changes in energy usage caused by vulnerability would be required, because the magnitude and consistency of the expected changes in vulnerable customers is currently unknown.

Of the three models, the appliance disaggregation and prediction model and the change detection model showed the most promising results. The appliance disaggregation and prediction model showed promise for further development, however, this is dependent on obtaining a wider range of higher resolution energy usage data. Particularly, energy usage data resolved to a higher resolution than 30 minute windows for households for model development and an accurate capture of appliance energy usage, particularly for medical devices. A 10-second time frequency could be available to the energy suppliers in the future as standard [4], however the key focus of this phase of the project is to utilise data sources that are, or have a reasonable likelihood of being, readily available for DNO usage. As such, given time available for the project, further development was de-prioritised until higher resolution smart meter data was readily available to develop the appliance prediction model further.

The change detection model showed a strong ability to detect energy usage changes, and the work in WP2b concluded that there was a high degree of confidence that given a wider ranging smart meter dataset the model could be better developed. In order to improve confidence in the model performance, a more accurate characterisation of vulnerability and patterns in energy usage is required, allowing the model performance to be better tuned to energy usage patterns characteristic of vulnerable households. Following this conclusion, improving the change detection model became the focus of WP2c.

1.3 WP2c Aims

The primary aim of WP2c of the project is to further develop the change detection model to be able to recognise the patterns of energy usage expected from the onset of various vulnerability characteristics for each individual household. The scope of the work package is to improve the model performance such that lower intensity signals could be detected in the data.

Secondary aims of the phase of the project are to gain a better understanding of how vulnerability may be represented in energy usage and to provide a dashboard with a representation of how the model may be applied to improve the process of identifying individual households that could contain vulnerable residents.

1.4 Dataset

A large smart meter dataset was sourced from Hildebrand containing energy usage and associated metadata for numerous households across the United Kingdom. The dataset consists of 5,691 households with data ranging between November 2017 and the present day. This dataset is more recent than the UKPN dataset [3] used in WP2b, and should provide a more meaningful analysis given changes in energy usage patterns over the last decade. A potential complication of the dataset is the onset of the COVID 19 pandemic, which changed energy usage patterns of households in an unpredictable fashion due to the requirement to stay home and when external institutions were closed.

The energy usage data is provided in Watt Hours (Wh) aggregated over each half hour within a given day. The metadata for the majority of households includes built form, current and potential EPC ratings and local authority. Alongside this, to filter out households with exceptional energy usage patterns, a metric was developed to filter households with electrical vehicles based on excessive overnight charging usage and households with solar panels, energy from which bypassed the detection device.

Households within the dataset are self-identifying, with energy users manually signing up to the Hildebrand service for their energy data to be included. This has a side effect of having households where the occupants are likely to be interested in their energy usage, and to actively monitor their own habits. There is no metadata showing the vulnerability status of each of the households, so some households within the analysis may be exhibiting vulnerability energy usage features.

1.5 Report Structure

This report captures the design and findings of work package 2c (WP2c) of the project, a further development of the prototype change detection model to detect changes in energy usage on an individual household level. The remainder of this report is structured as follows:

- Section 2 details the breakdown of “usage features”, the set of adjustments made to energy usage to represent the characteristics of vulnerable households.
- Section 3 details the approach taken to developing the change detection model, and follows with results taken from bulk application of the model to test data.
- Section 4 captures how the model can be applied on an individual household basis and details of a series of scenarios used to demonstrate the model capability.
- Section 5 and 6 detail conclusions and recommendations for further development of capability respectively.

2 Usage features

2.1 Concept

There are many categories of vulnerability, with each having different demands on the energy usage of a household. These demands on the energy usage are designated as “usage features”, consisting of patterns within the data representing the difference in energy usage behaviour expected between two identical households, where one exhibits a resident with a vulnerability and the other does not.

The usage features take an individual expected behaviour of a vulnerable customer and describe how they are represented in energy usage data. Detecting a usage feature could indicate that a given household contains an individual with the related vulnerability, but it is not a necessary condition of the feature existing in the data. Similarly, a household can exhibit multiple usage features, each representative of the same or multiple vulnerabilities, potentially making detection difficult.

Figure 1 shows a mapping of how usage features can be utilised to develop an understanding of the vulnerability of households, and provide context of energy usage changes in the wider space of vulnerability detection.

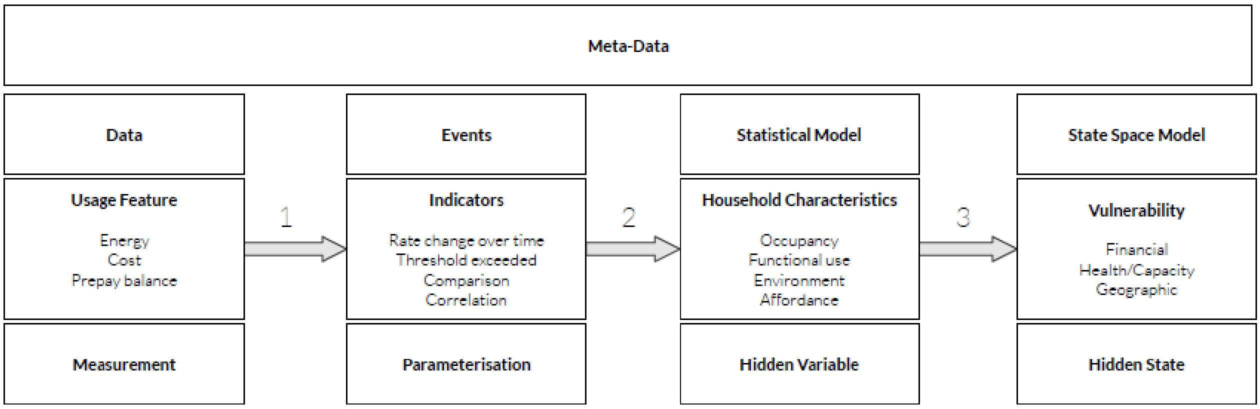

Figure 1: Usage feature diagram showing the context of how data measurement can be mapped into modelling of the unknown vulnerability state of a household.

From the left, Figure 1 shows how a usage feature in the available data (whether energy usage, prepayment details, etc.) can be measured and mapped into becoming parameterizable indicators. These indicators provide a richer image of patterns within the usage data, and provide an event-based series of data features that allow model users to observe key elements of the data. For example, an indicator may be a threshold of some parameter describing a set of usage data being exceeded.

Once indicators are raised, the underlying variables and characteristics of a household that are not outwardly visible within the data can be inferred. This may include for example household occupancy levels, functional uses of energy or the ability of the household to afford required energy usage. The hidden variables, which capture the characteristics of the households, can be mapped onto potential vulnerabilities within households, the highlighting of which captures the final goal of the model functionality.

The model developed within the scope of this work package focusses on development of the indicators from underlying energy usage features, indicated with the number one on the diagram.

2.2 Implementation

The Hildebrand dataset does not include information on whether the household is vulnerable and therefore accurate vulnerability characteristics cannot be extracted or tested on the raw household usage. In the absence of this knowledge, usage features were manually created based on assumptions around the expected change in usage. Functions were created to apply the usage feature over the existing data for a range of days, each with a "severity level" which allowed the expected change to be applied on an increasing level of severity. The usage features were applied to test ranges for each individual household, allowing a calculation of the ability for a model to detect whether the usage feature becomes a feature of the household's energy usage.

The initial usage feature developed to represent a vulnerability for the purposes of model development was the "increased daytime usage" feature. Additional usage features are listed in detail in Section 4.2 below. The "increased daytime usage" feature is intended to represent a household where at least one of the residents has lost their job, and is therefore in the house during the daytime on a more regular basis. In the case of losing a job, they are determined to be vulnerable because of a reduced ability to pay for the electricity required. The feature applied to a test day split into 48 half hour windows is shown in Figure 2.

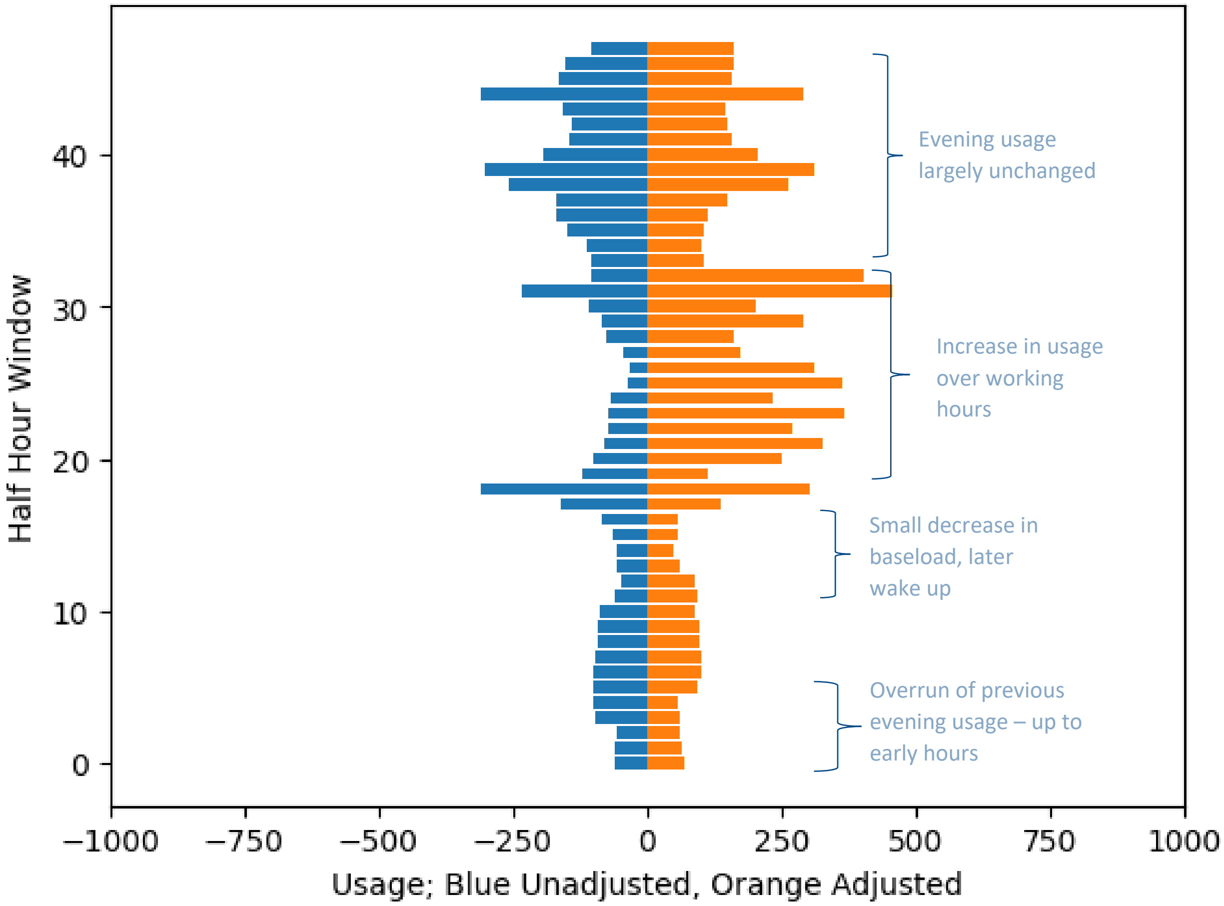

Figure 2: A randomly selected day, from a randomly selected household, showing energy usage before and after the "increased daytime usage" feature is applied.

This feature is designed with two adjustments made to a household's usage; between the hours of 10am and 4pm an increase in usage roughly commensurate with evening occupancy of the household is applied to represent the additional occupancy of the household during the day, as well as a shift in energy usage pattern later in the day to represent a change in sleeping pattern to a later bedtime and a later wake up. The large spikes in usage are unchanged, representing lesser used high energy devices being used at unchanged times.

The increase applied between the hours of 10am and 4pm is dependent on the severity level used. The severity levels are intended to measure the performance of the model, a better model would be able to detect changes at lower severity levels. In this case, the increase is by a factor of the 95th percentile of evening usage height for the given day, ranging from 50% at severity level zero to 150% at severity level five, all with random variation added.

3 Change detection model

3.1 Initial model and investigation

The change detection model created in WP2b of the project was designed as follows:

1. A lognormal distribution was fitted over each half hour usage window for a series of days known as the training set.

2. For each day within a test set of days, the probability that the energy usage for each given half hour window lay within the expected distribution for that half hour window was calculated and summed over the day.

3. A threshold value was utilised to determine whether a range of days was considered as changed behaviour if more than half the days were over a selected threshold probability of being within the expected distribution.

Investigation of the performance of this model against the increased daytime usage feature showed a lower than expected performance of the model's detection capabilities with some households. Therefore, an early phase of the work was to investigate and determine areas where the model could be improved.

High levels of energy usage, known as peaks, tended to skew results towards identifying false changes in pattern as the likelihood of the energy usage fitting the distribution tended to zero for each energy window. These low likelihood results for the windows for peaks in energy usage eclipsed the differences in small shifts in the distribution of energy usage for more general household usage. It was observable that the peaks in the data appear on a regular basis, they generally represent the usage of high-power appliances such as washing machines, but the usage does not follow a regular pattern in such a way that the usage characterisation is captured in the half hourly energy distributions.

For example, Figure 3 shows the variation in energy usage for a half hour window of 10:00 to 10:30 for a randomly selected household over 200 days. A lognormal probability distribution has been fit to the 200 usage values. There is a large peak of probability for the lower energy levels and a small rise representing peaks; it is not unexpected for the peaks to occur during each day, but the regularity of a peak being within a given half hour for sufficient days to be statistically significant is low.

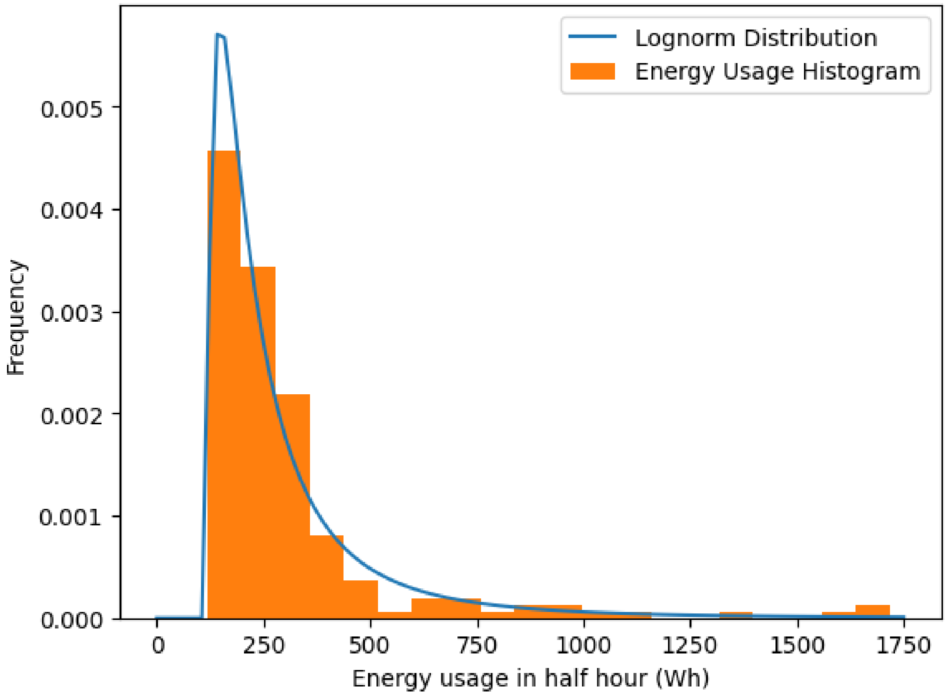

Figure 3: Distribution of energy usage over a 200 days for a half hour window 10:00 to 10:30, for a randomly selected household, including fitted lognormal distribution.

A secondary area for potential improvement of the initial model was that the households were being assumed to having a single typical energy distribution pattern, representing an average day. The distribution was described by each half hour window's distribution pattern, and in order to cover the range of days observed in the training set had a large uncertainty relative to the energy usage. Reducing the large uncertainty in the characterisation of the expected pattern would provide improvements in detecting small changes in energy usage pattern. Applying a feature to the usage such that the pattern would be noticeably altered, but not shifted outside the broad range provided by the distribution of all the days would be better detected by multiple clusters with smaller uncertainty. Investigation into the performance of the initial model suggested that the key feature of detecting energy usage patterns is not in the individual half hourly usage in isolation, but in the overall pattern of each given day; therefore, a methodology that treats each day as an input of 48 half hour windows taken as an overall pattern is likely to be more successful than treating each half hour window as separate distributions which are less variant between days.

3.2 Peak detection and baseline generation

To improve the model's detection capabilities the noise and variability in the usage data must be reduced. The more interesting patterns are those of the bulk of the energy usage, so large peaks of data provide a high level of uncertainty without being major elements in describing the characteristics of energy usage in the household. A process by which the "peaks" in the data could be isolated and removed such as to observe the ongoing trends in the pattern of the underlying "baseline" energy usage data was created. The baseline energy usage represents the underlying load of energy that is independent of standalone instances of high energy usage. Within the scope of this package of model development, the characteristics of vulnerability that will lead to baseline pattern changes are more important, and analysis of the trends in peak detection is left outside the scope of this prototype model.

The peaks are detected by finding the half hour energies that protrude significantly above the surrounding areas of energy usage, and the surrounding half hour windows between the peak and associated trough in usage pattern. The peaks are then separated from the baseline usage, and the gap in the underlying baseline is filled in using a linear interpolation. Figure 4 shows an example of a day where the peaks are identified and separated from the baseline usage. The baseline energy usage is shown in blue, with the peaks highlighted in orange.

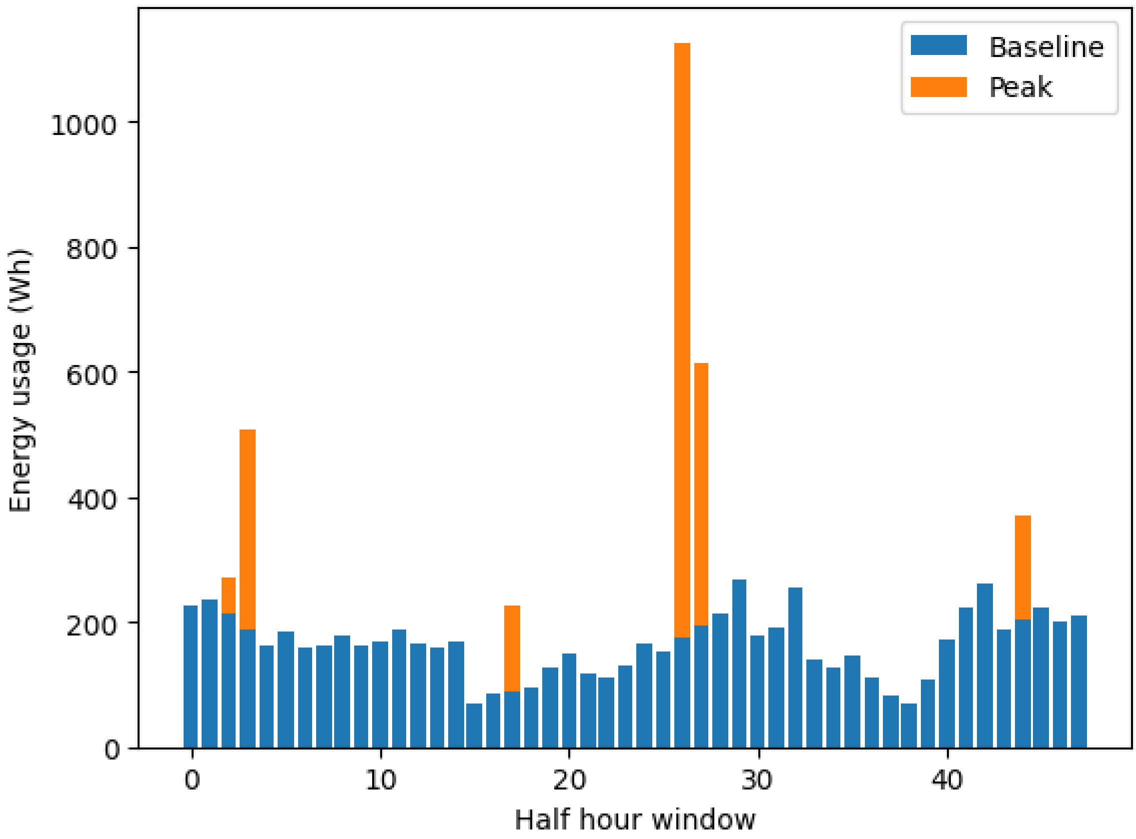

Figure 4: Plot showing the separation of the peaks and the baseline energy usage for a random day from a randomly selected household.

The separation of the two components of the energy usage allows the model to observe the adjustments in trends of general usage of electricity by more consistently used appliances, such as lights, heating and small kitchen appliances without being skewed by the more sporadic, higher energy appliance usage as seen in the previous model iteration.

3.3 Day clustering

A new change detection model was developed using KMeans clustering[1] to cluster the typical day patterns for each household into distinct patterns. The clusters consist of a pattern of 48 half hour windows of energy usage as seen in the smart meter input data. Like the previous iteration, the model is trained on each individual household over a training period of days, and tested over a separate series of days. The model is trained and tested only on the baseline usage with the peaks removed as discussed in the previous section.

KMeans clustering creates a set of representative energy usage patterns and adjusts these usage patterns, known as cluster centres, to minimise difference between the inputted training days and the cluster centre that most closely resembles each day's energy usage. The number of clusters is a parameter for development of the model, and represents a balance between: forming too few clusters where individual "types" of household day will not be

[1] 2.3. Clustering — scikit-learn 1.2.0 documentation

captured; and too many clusters where changes in pattern are not highlighted as too broad a range of "expected" behaviour is identified due to cluster overfitting.

The clustering was run for a randomly selected household, and Figure 5 shows bar plots of the half-hourly usage of a randomly selected day that the model attributed to each of the four clusters. The orange line represents the cluster centres for that cluster, which is the energy usage pattern over a day used as a reference against which each day is compared.

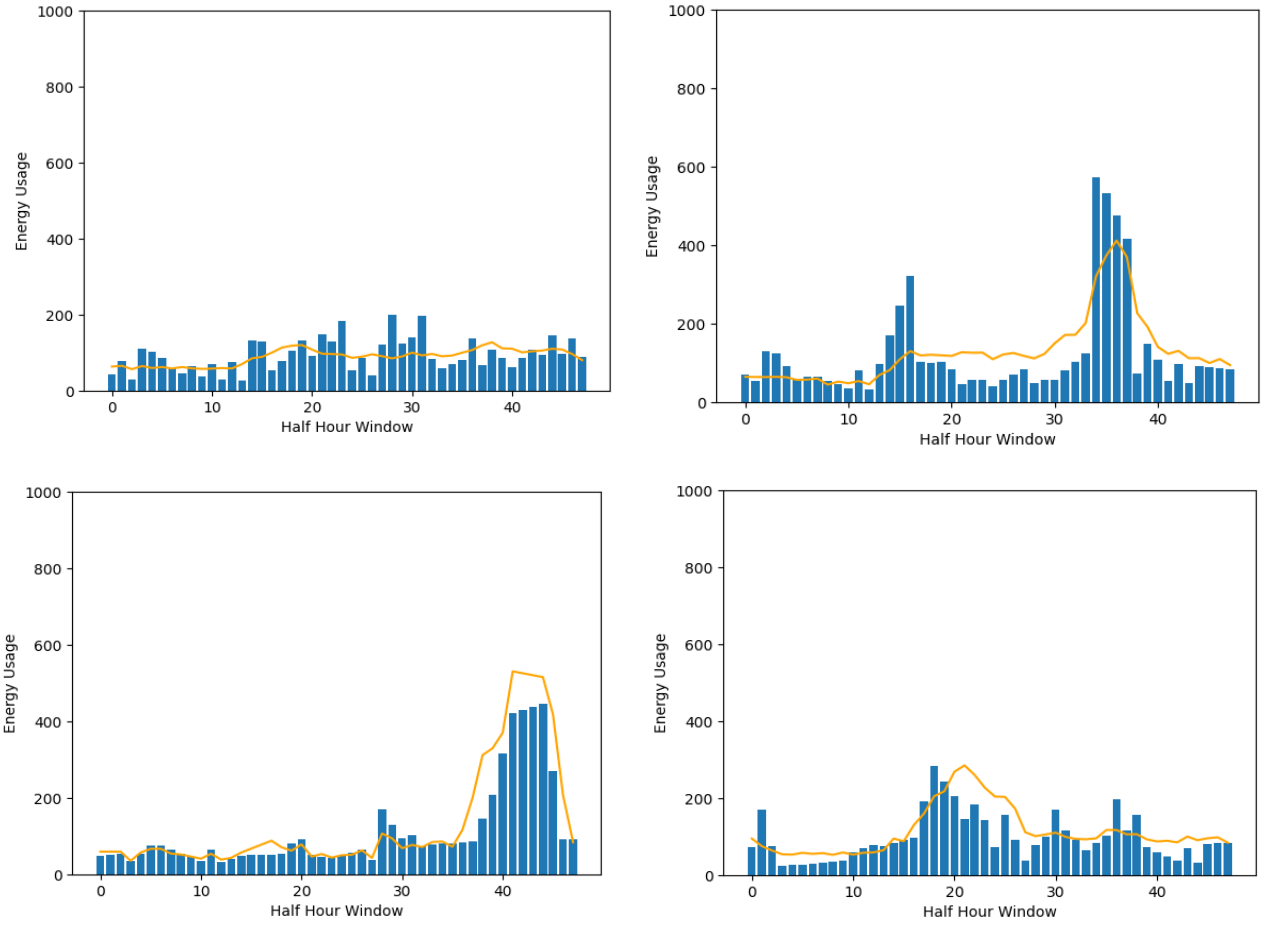

Figure 5: The half hourly usage for a randomly selected day from each cluster. The orange line represents the cluster centre for each cluster. The clustering was run on a randomly selected household.

The figure shows how each day's energy usage can be fitted roughly into the usage distribution for one of the clusters, with the orange lines representing the energy usage of the cluster centres and the blue bars the actual baseline usage. New days of energy usage can be measured against each of the clusters and fit into the cluster that has the most accurate pattern based on minimising the difference between each half hour in the usage and the cluster centres.

3.4 Change Detection

There are two ways a change in behaviour can be detected. Either by a consistent 'new' type of day is observed, or the relative proportion of each day 'type' changes. Therefore, two detection methods were developed.

3.4.1 Outlier detection

From these clusters, the next step for the model is to calculate a threshold indicator to determine whether the usage in the given test period of days sits outside of the expected clusters, thereby indicating a change in household energy usage. The outlier detection method calculates the absolute sum of the difference between the energy usage for each half hour window and the centre of the cluster the day was predicted to be from. This value is normalised by the overall energy usage for the day, and is known as the "distance to cluster centre". The distance to cluster centre is also calculated for each of the days within the training set, providing an expected distribution of days energy usage relative to the fitted clusters.

The mean of the expected distribution of distances is taken as a threshold value for each day in the test set. If the distance between the energy usage in the test day is greater than the mean distance of the training set, then that day does not fit any of the known patterns for the given household's usage and is considered an outlier day.

To reduce the noise of natural variation of household usage, the household is evaluated over a set of test days and the proportion of days highlighted as outliers is used as threshold for detection. The thresholding method is intended to show a sustained change in days to becoming outliers relative to the household clusters, thereby indicating behaviour changes rather than day-to-day variation.

3.4.2 Cluster proportions

The cluster proportions detection method takes a proportion of days in each of the clusters as a rolling window over the test period. Usage changes in a household may alter energy usage to sit within different clusters to those expected as opposed to being energy usage that is an outlier for all the clusters generated in the training set. For example, a household using less energy for heating to avoid costs during the winter will have an increase in days that are clustered into the usage cluster that represents the summer rather than being total outliers to the household usage. In this case, this may include self disconnection, an example of which is included in the scenarios presented in Section 4.2.

A first pass detection method for analysis to detect reductions in energy usage is to set a threshold of proportion of days within a rolling window that are classified into the lowest energy cluster. Figure 6 shows the four energy usage cluster centres for an example household. Figure 7 shows the relative cluster assignment proportions of the same household over the range of one calendar year for a rolling 28-day window. The figures show a general seasonal pattern to the clustering of energy levels. The lowest energy usage cluster, shown in blue, is the predominant usage cluster within the summer months, representing low heating and lighting usage, and may also represent minimal occupancy for holidays. The highest energy cluster shown in red is the most common during the colder winter months at the end of the year. The cluster in green shows an interesting peak in the mid-morning, and is most common during the spring and autumn of this year, perhaps highlighting areas where there is an increased occupancy in the household due to some unexpected behaviour change.

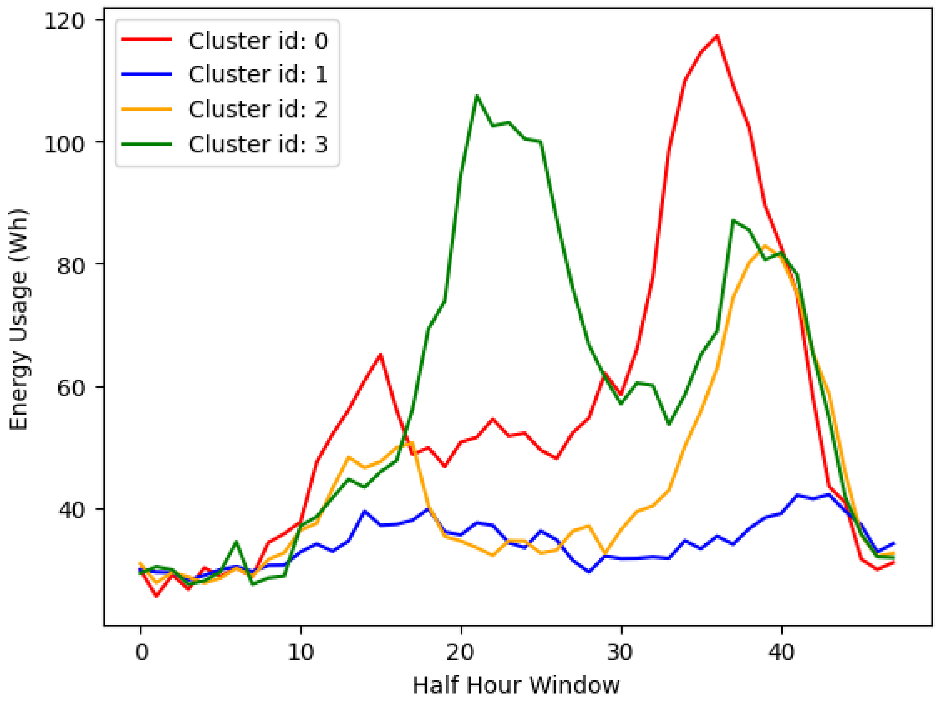

Figure 6: Cluster centre energy usage for an example household trained using the KMeans clustering model.

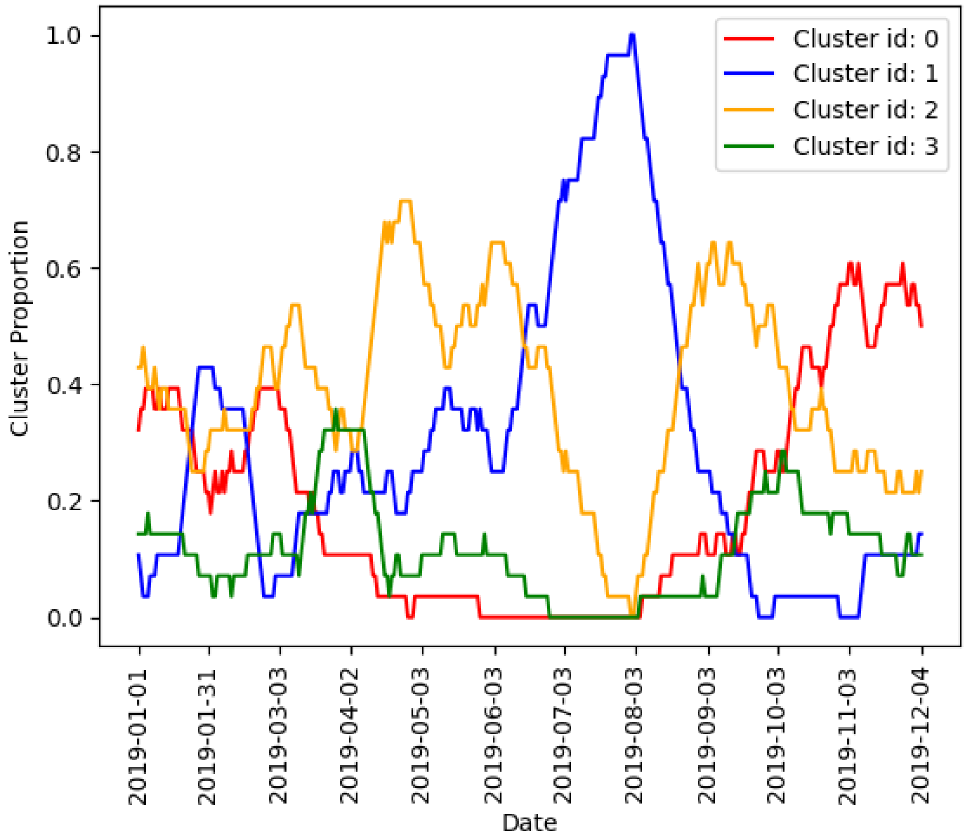

Figure 7: Proportion of days in a rolling 28 day window attributed to each cluster as shown in Figure 6.

3.5 Testing model performance

3.5.1 Receiver Operating Characteristic Curve

The outlier detection model was tested utilising a train-test split of 100 training days and 30 test days over 200 households within the dataset. The training and test days were taken consecutively, in the data presented here starting on the 1st May 2019. For each test set the "increased daytime" usage feature was applied for a range of severity levels, including a level with no change. The model was used to predict whether each day is considered an outlier and a range of thresholds between zero and one, each representing a proportion of days within the test period which have been identified as outliers.

A positive result, representing a detected change in behaviour, is defined as a proportion of days in the test period greater than a detection threshold, where outliers are defined by distance to cluster centre being greater than the mean of the distance to cluster centre for the training set. This is a true positive detection if the change had been applied to the household, and it is a false positive detection if no change had been applied[2]. The calculation of the false and true positive rates is made to review the capabilities of the model to detect the usage features applied above the natural variation in household energy usage patterns. A well performing model minimises the false positives, thereby reducing overload of verification checks against the outputted detections, whilst maximising the true positive rate to ensure as many households as possible that are exhibiting vulnerable characteristics can be identified.

The testing results are presented in a Receiver Operating Characteristic (ROC) curve, which shows the variation in detection ability between the usage feature severity levels and detection threshold values. The Area Under the Curve (AUC) can be calculated as a metric for performance – an ideal curve has an area of one where 100% of true changed households are identified for all thresholds above zero and 0% of unchanged households are identified as such. In a real model, the area under the curve is unlikely to be one, but a higher value indicates that the model can be more effectively used without having excessive missed positive identifications or falsely identified positives. Figure 8 shows the ROC curve for the outlier detection model for the increased daytime usage test. The severity levels represent an increase in magnitude of the signal applied to the data, in this case the increase in energy applied to the working hours in the day and the shift in pattern to later hours is a greater shift in time.

[2] It is noted that there may be some variation in the underlying usage of the household, so a small proportion of false positive results will be due to detection of actual change in household behaviour that has not been artificially applied.

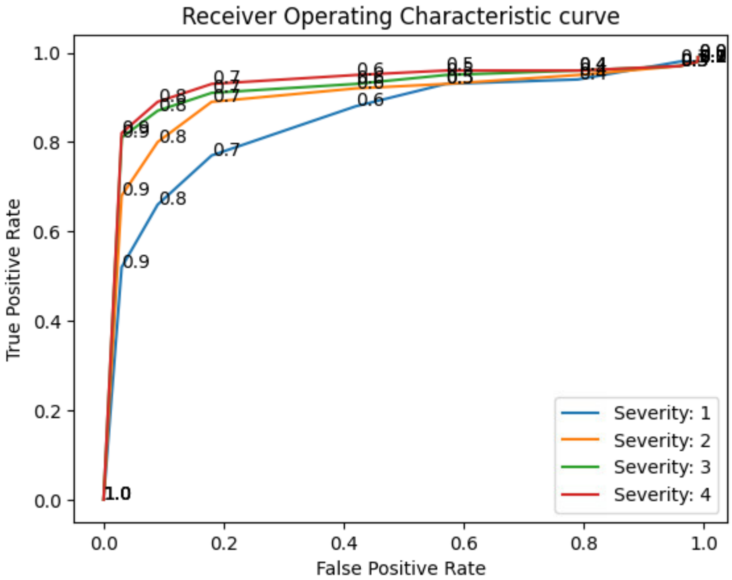

Figure 8: A set of Receiver Operating Characteristic curves for 200 households with the increased daytime usage feature applied on a range of four severities.

The figure shows how for increasing severity of change, the model performs better at identifying days which have had the usage feature applied. A diagonal line, where true positive rate equals false positive rate, would be representative of selecting whether a household has changed at complete random. Figure 8 indicates that the model is performing well, even for the smallest change (severity 1.0), at detecting true positives with a low false positive rate.

3.5.2 Severity vs Detection rates

One of the key expectations when developing the model was that higher severity of applied change would be more easily detected by the model. Performance of the model in detecting false positives is heavily dependent on establishing the signal of a behaviour change from the noise of expected household variation. Figure 9 shows the true positive rate generated over a range of severity levels for the increased daytime usage feature model for a detection threshold of 0.8.

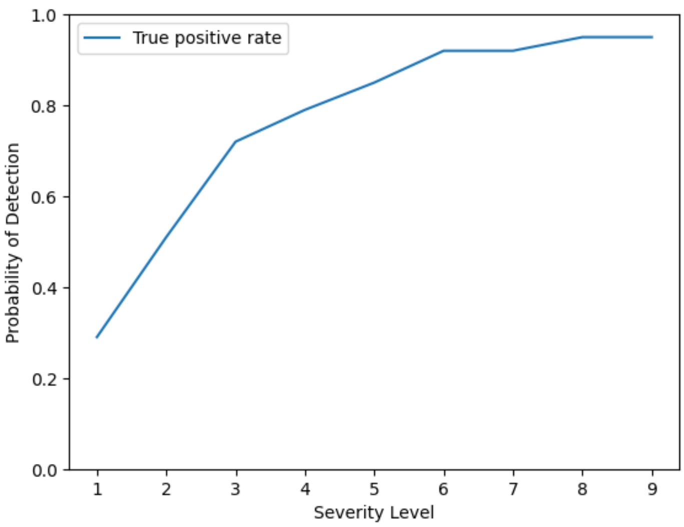

Figure 9: Plot showing the true positive rate for a series of test households with the increased daytime usage feature applied with nine severity levels.

The figure shows how as severity increases, the model becomes more capable of detecting the change above the noise of the household usage, with the true positive rate tending towards one where the change becomes highly severe.

3.5.3 Investigation into number of clusters and dates of test period

To improve the performance of the model, an investigation into two variable parameters was made:

- Utilising a different range of days as test and train sets to ensure any seasonal skews are removed from the model performance.
- Utilising a different number of clusters to apply within the model.

Three start dates were selected and applied to separate models trained to produce between four and seven clusters. The area under the curve for model performance with the increased daytime usage feature at severity level two was calculated. Figure 10 shows the different area under curves for each model run.

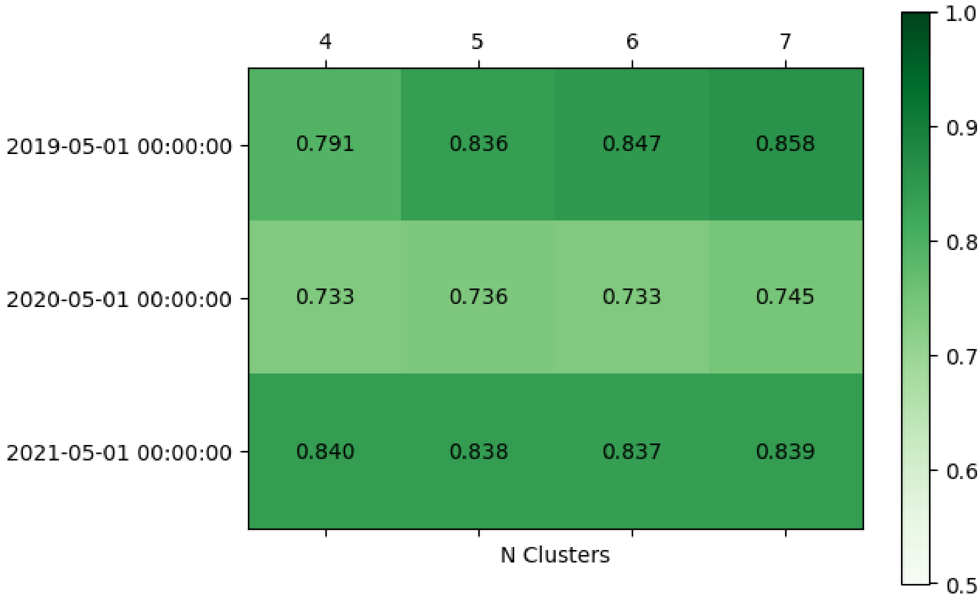

Figure 10: Heatmap showing area under the ROC curve for variation in number of clusters and start date of training data where the model is applied to the increased daytime usage feature.

The figure shows that an increase in the number of clusters could provide some performance improvements for the test day ranges, however further investigation showed that there was significant overlap between clusters as the number increased, and performance was largely unchanged when a wider set of test and training days was selected.

The performance was significantly reduced when trained on data from May-August 2020, which could be indicative of the effects of COVID 19 lockdowns on the underlying household behaviour. Where the model is trained on exceptional days but tested on “normal” behaviour a well performing model would highlight the normal days as being changed. In this case, the unchanged days are being detected as having changed, which in this test harness are indicated as being false positive detections, and not as the true positive results of an underlying change of data being detected. This difference in performance between start dates provides some confidence that the model can detect differences in household behaviour, as changes in household behaviour because of the underlying data, rather than the manually applied usage feature, would be labelled as false positives in this circumstance, thereby reducing the AUC as observed.

4 Results

4.1 Application of model to individual household

Section 3.5 captures the overall performance of the model when run on a bulk of houses for a 30 day test period of days. The aim of the project is to convert the usage data of individual households into a series of indicators highlighting where features of interest exist in the underlying data. The indicator in the case of the change detection model, is the application of a threshold alerting system to highlight where a household may have changed behaviour. An investigation of the usage of the model on individual households demonstrates the ability to apply the model to detecting individual vulnerabilities, and potential uses of the outputted threshold indicator are shown here. The foremost of the use cases of the indicator is highlighting to a human operator of the model who may be able to investigate the feature change detection further.

Demonstration households are split between a training period of 185 consecutive days and tested on the following 280 days, with the increased daytime usage feature applied during periods within the test range. The model is trained and applied over each day within the test period to calculate a distance to cluster centre value, and days are determined to be outliers if the distance is above the mean distance in the training set. A rolling window of 14 days is applied to the test period, and if more than 70% of the days within each window are indicated by the model as outliers then the window is highlighted as being changed behaviour.

Figure 11 shows an output of these results, with the changed days shown in blue and days of detection highlighted in green, below. There are two periods of usage feature applied, the initial period starting 15th April 2020 being applied only to weekdays and the latter period starting 1st July 2020. The orange line shows the distance between the usage for the day and the nearest cluster centre, with the blue sections of line showing the differences for the days that have the usage feature applied.

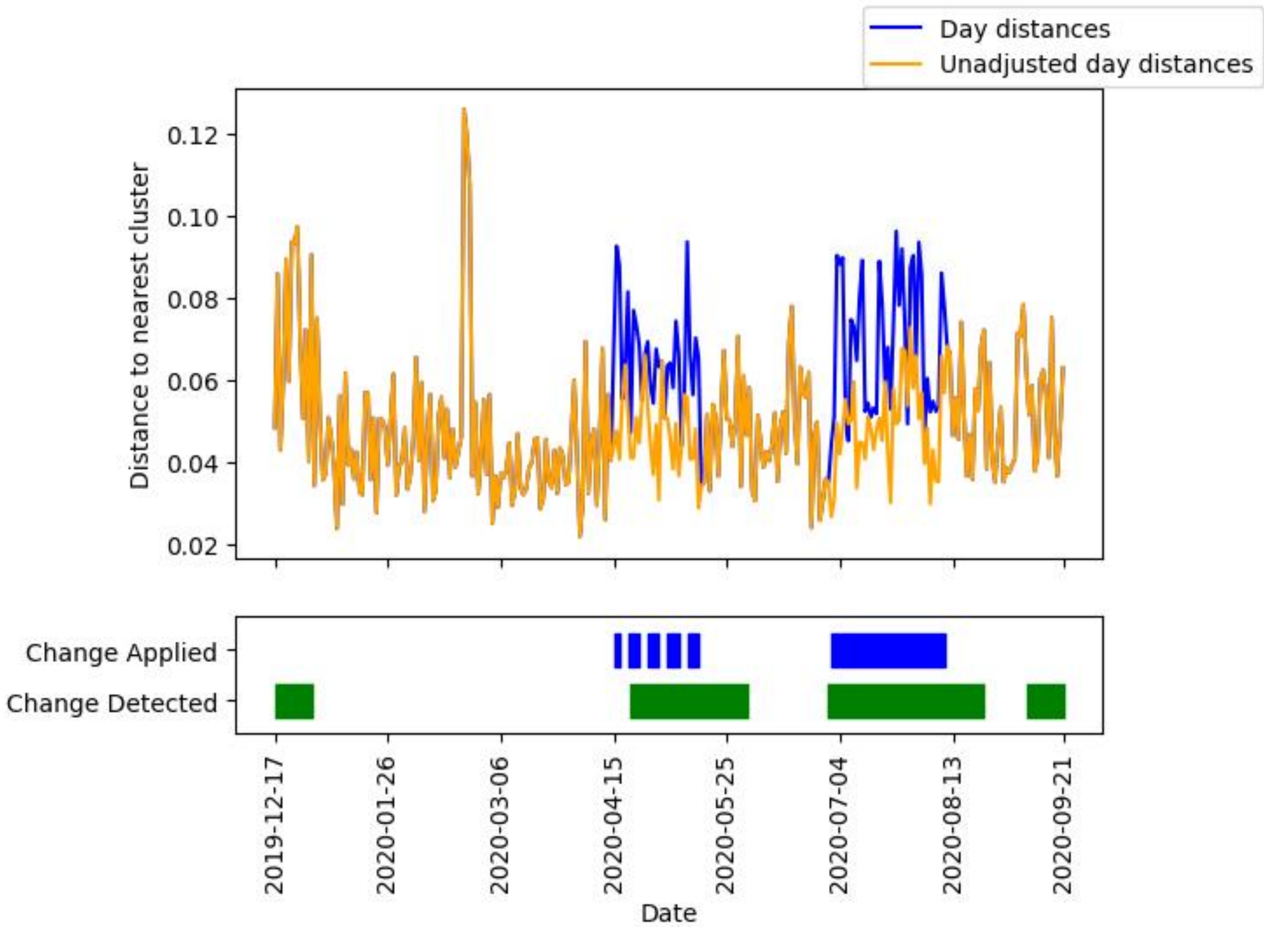

Figure 11: A plot showing the application of the model for a randomly selected household, where the increased daytime usage feature is applied for short periods in the test range. Shown are the areas of applied change, change detection, and the calculated distance to cluster centre for each day.

The lower plot shows the model performance on windows in the test period, where the days in which the change is applied and the windows in which change is detected are shown in blue and green, respectively. The overlap between the bars show where the model detected the applied changes during the two periods of usage change.

Also highlighted in green are two false positive detections; one at the start of the test period, and one at the end of the test period. The former window covers the Christmas period where the expectation is that household energy usage is likely to sit outside the known pattern, so a false positive result is a representation of the model being able to detect outliers in daily energy usage. The latter period of false positive detection is less well explained, and raises the need for better understanding for any operator using the model. In this case, further contextual information is required to determine the cause of the change.

The Christmas period highlights the need to reduce the levels of false positives shown to users, particularly those around irregular days. Further investigation of the distribution of distances to minimum cluster was largely uncorrelated with weekday or weekend status and a weak correlation with bank holidays. It is feasible that removing bank holidays from the analysis to prevent false flags is an option, however the results are heavily dependent on household; for example, many households have work patterns that are not affected by bank holidays.

4.2 Usage Feature Scenarios

To further demonstrate the capabilities of the model, several scenarios were created and applied to sample households. These scenarios alter the energy usage data from a specified date for each household based on the expected changes for a given set of vulnerabilities.

Table 1 shows each of the vulnerability scenarios created and how they have been implemented. The detection method used varies depending on the scenario. Reductions is usage generally are better detected by the cluster proportions method and changes or increases in usage are better detected by the distance to closest cluster method.

Table 1 Usage Feature Scenarios

Scenario Name	Description of vulnerability	Impact on energy usage	Implementation of usage feature	Detection method used
Job loss	An individual in the household loses their job, and therefore may have a loss of income.	Person is at home more during the day and they move to a later wake and sleep pattern.	Increased in daytime usage commensurate with the typical evening usage, and shift of pattern later to represent later wake-up and bedtimes. Applied to all weekdays.	Distance to closest cluster
In-home carer	An individual in the household requires daytime in-home care, and therefore may be medically vulnerable in case of power delivery issues.	An additional person is using appliances in the daytime.	Small increase in daytime usage commensurate with increase in usage of household devices such as using a vacuum cleaner, cooking meals, medical device usage or using other kitchen appliances.	Distance to closest cluster
Struggle to pay	The household does not have the required income to pay for electricity and therefore may be vulnerable to spikes in energy prices.	Person is consciously reducing their energy usage.	General decrease in baseline energy usage, and baseline is reduced to lower bulges to represent less use of appliances and lower heating settings. Peaks in data are shrunk to represent deliberate reduction in high power appliance usage.	Proportion of minimum energy cluster
Cannot pay – monthly pay check	The household has insufficient income to pay for electricity and is paid on a monthly basis. The residents of the household may be vulnerable to high energy prices.	Person significantly reduces their usage at the end of the month.	General decrease in baseline energy usage and baseline is reduce to lower bulges. This is applied over the last few days of the month with increasing severity.	Proportion of minimum energy cluster
Cannot pay – weekly pay check	The household has insufficient income to pay for electricity and is paid on a weekly basis. The residents of the household may be	Person significantly reduces their usage at the end of the week, and potentially gets	General decrease in baseline energy usage and baseline is reduce to lower bulges. This is applied over the last few days of the week with increasing severity.	Proportion of minimum energy cluster

	vulnerable to high energy prices.	cut off from electricity.	In addition, periodic windows of energy usage within the last days of each week taken to zero to represent being on a pre-payment meter with insufficient funds.	
Reduction in storage heater usage	The household utilises storage heaters and the residents of the household may not be able to sufficiently afford to heat their house in inclement weather.	Person deliberately reduces their heating usage in the winter when compared to previous years.	The storage heater energy usage between 12:30am and 7am is reduced by a high proportion for all days, with an increased reduction near the end of the month.	Proportion of minimum energy cluster.
Post hospital	An individual in the household has recently returned from hospital treatment, and may be vulnerable to power outages, especially if medical devices are required.	Person is at home more during the day and they move to a later wake and sleep pattern.	Increased in daytime usage commensurate with the typical evening usage, and shift of pattern later to represent later wake-up and bedtimes. Applied to all days.	Distance to closest cluster
Overnight medical device	An individual using an electrically powered overnight medical device may need additional support in case of power supply issues.	Person has increased usage at night-time due to appliance.	Small increase in energy usage overnight between 10pm and 7am, commensurate with the energy usage of a medical device such as a CPAP machine.	Distance to closest cluster
New child	Households with young children are considered vulnerable with additional requirements in case of outage. In this household a new, young child becomes present.	Significant change in usage pattern throughout the day.	Irregular increases of energy usage in line with small appliances through the night, morning peaks in usage shifted earlier in the day to represent an earlier wakeup time and a small increase in daytime usage to represent additional household occupancy.	Distance to closest cluster
Dementia	An individual within the household has dementia, and therefore is using appliances in an irregular fashion. Patients with dementia will be unlikely to manage with power outages.	Higher usage in the daytime, inconsistent with existing usage patterns.	Additional bulges in baseline energy, commensurate with small to medium appliance usage sporadically added to each day in the household.	Distance to closest cluster

4.3 Individual results from a scenario

The following section presents and discusses the results from the overnight medical device usage scenario being applied to an example household. The results from all 10 scenarios are presented in Appendix Scenario FiguresA.1.

The change is an increase in night-time usage to represent the usage of a medium energy medical device overnight. The model is trained over 185 consecutive days starting on the 15th June 2019, and applied to a test range of 280 days. The medical device usage feature is applied during this test period on the 15th April 2020.

Figure 12 shows the feature as applied to an example day for the household. The orange bars show the energy usage after the medical device usage feature is applied, the blue bars show the same day without the change. The overnight increase in energy usage is visible here, and the remainder of the day remains unchanged. The night-time usage increase is in line with existing periods of higher usage within the household.

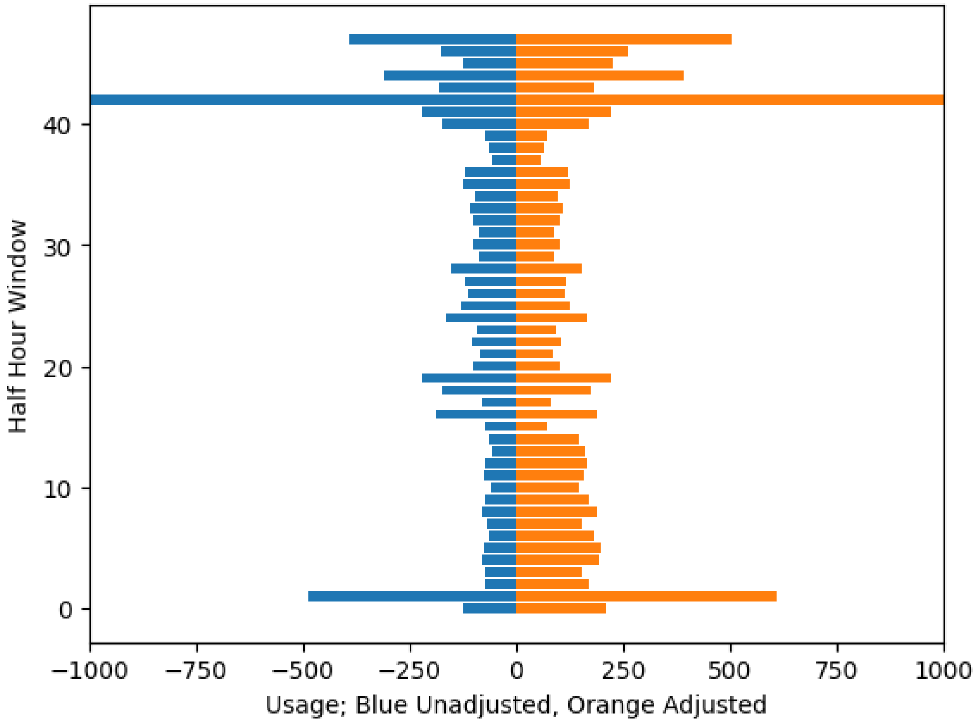

Figure 12: A random sample day with the overnight medical device usage feature applied showing energy usage before and after the change is applied.

The model is then applied to the 280 day test period and the change detected. Figure 13 shows the daily energy usage over the period, with the orange line capturing the unchanged usage and the blue line capturing the usage. The days which are changed are highlighted with the blue bar, and fortnightly windows where the change has been detected are highlighted with the green bar.

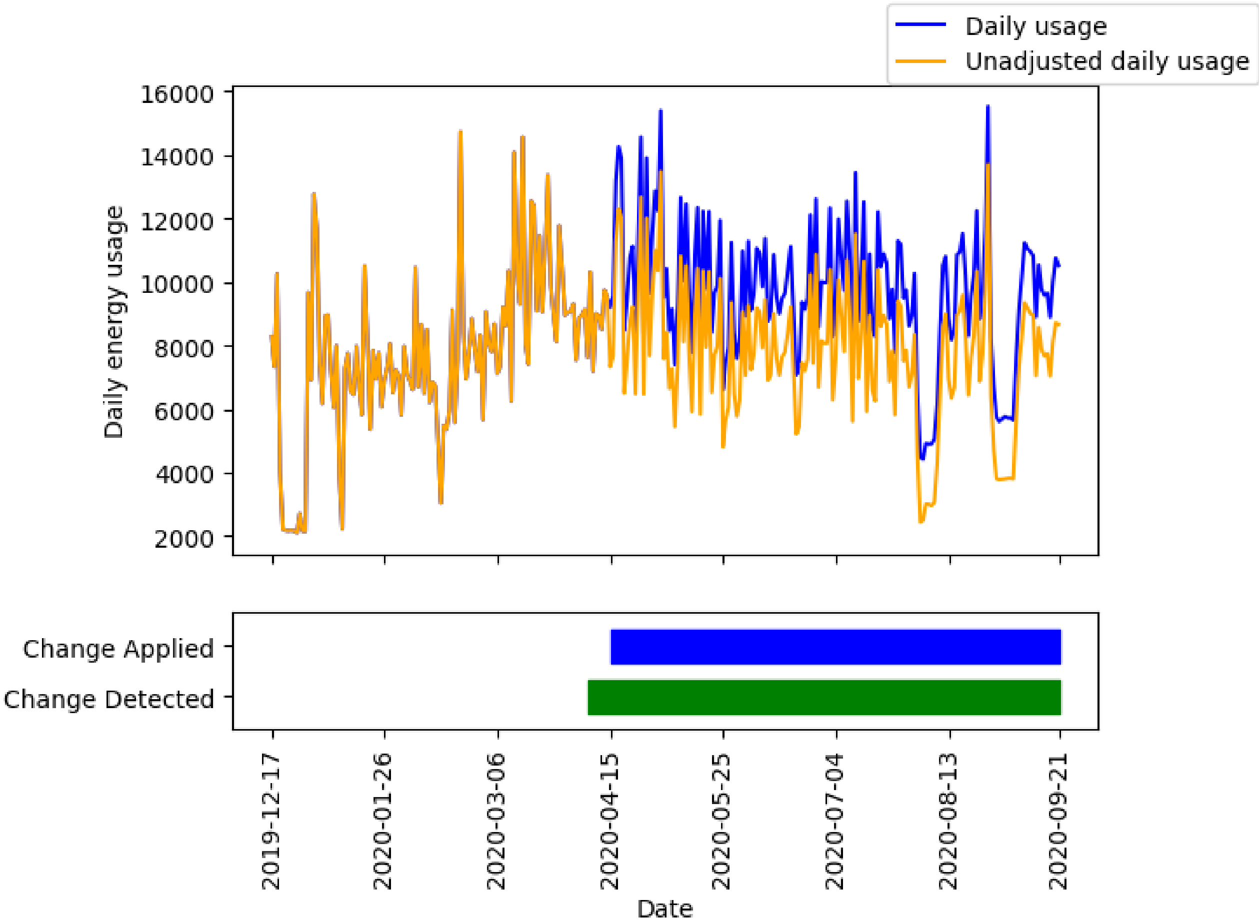

Figure 13: A demonstration of the change detection model on the overnight medical device scenario, showing daily energy usage and the areas where the change is applied and detected.

The figure shows how the model can detect the change in energy usage pattern throughout the period onto which the change has been applied and shows no false positive detection during the period where the change is not applied, with the exception of a few days prior to the change. The exception is because the detection is applied retrospectively on a rolling basis, so the 14 day window starting prior to the change sees sufficient outliers to highlight the full 14 day period as having changed. This is a clear example of how the model can be used across households to show changes in usage pattern, and the results of this detection can be utilised to further investigate the causes of the change in household characteristics through manual intervention or further automated detection, and prompt further action in determining if the household can be classified as vulnerable.

5 Conclusions

This report presents our increased understanding of the characteristics of energy usage patterns observable on smart meter data aggregated into 30-minute windows on a per-household basis. In particular, how usage can be modelled as a series of regularly repeating daily patterns. The initial model concept of detecting outliers to observe a change in behaviour from a modelling of expected energy usage from a series of training days for a household has been further refined and developed to remove noise from large peaks in energy usage and broaden the dependency on daily patterns rather than treating half hours in isolation.

We have also developed a model for how vulnerabilities can be represented as patterns in data usage, on a varying level of severity, and verified that these patterns are within the bounds of expected changes in usage due to changing circumstances.

By applying the improved model to implementations of these vulnerability usage features we have demonstrated its capability to provide an automated detection of changes in household behaviour with a high degree of confidence. Further, a demonstration of the utility of the model and alerting mechanism for capturing and highlighting changes to an operator has been shown across a range of households and vulnerability scenarios.

6 Recommendations

This project has demonstrated the ability for a conceptual model to characterise usage patterns into indicators that a household may have had changed behaviour. Within the broader context of the vulnerability problem space, this is a key feature in being able to understand the basic smart meter data but is insufficient of itself to provide a full picture of the vulnerability characteristics of any given household. This broader context is where the larger value of the modelling approach becomes apparent.

6.1 Further model development

The change detection model can be further developed to improve its ability to indicate changes have occurred, and some technical areas have been identified as being beyond the scope of this project.

Firstly, the minimum distance detection method could be improved for better characterise shifts or noise in the patterns of the household. In this implementation, the test days are considered outliers if they sit beyond the mean distance from the nearest cluster from the training set; this methodology could be improved to observe distribution shifts in the minimum distance, and by adding further context such as exceptional weather and working patterns such as holiday periods. Parameters such as number of clusters could be adjusted on a per-household basis to improve performance against highly variable households.

The cluster proportion methodology requires significant further development to be applicable in the general case. Many houses show seasonal trends in the proportion of days that are attributed to each cluster, which is not unexpected, much of household energy usage is dictated by the weather and daylight hours, however this means that detecting changes based on proportionality is difficult to remove from general variation. As such, further thought and development using the context of each given period should be used.

6.2 Model application

The next phase of business development around the usage of the model for the end goal of improving vulnerability detection is in further working on its application. In this project we have demonstrated the ability to highlight changes in a households usage to the user, but further intervention is required to determine the nature of the change and develop into household characteristics. There are two areas for further development; providing the model outputs to human users who would be able to apply external context to each individual household, or a process by which the model outputs alongside the external context could be used to create a secondary model. This secondary model would combine the indicated changes and any external context to identify an expectation of types of vulnerability or change.

It is important that any usage of smart meter data, particularly with automated detection methodology, should be performed with an understanding that consumers may be uncomfortable with the data usage and perceived level of oversight. Additional work may be required to ensure that the techniques set out here are used effectively without influencing consumers into reducing smart meter uptake, and to avoid consumer discomfort about the usage of their data.

7 References

[1] “Smart DCC,” [Online]. Available: https://www.smartdcc.co.uk/. [Accessed January 2023].

[2] Z. Hodgins, K. Chippendale and G. Nelson, “VENICE WP2: Modelling Approach & Findings,” FNC 010837/52726R, July 2022.

[3] “SmartMeter Energy Consumption Data in London Households,” 2014. [Online]. Available: https://data.london.gov.uk/dataset/smartmeter-energy-use-data-in-london-households.

[4] “Smart Meters: The Full Story,” [Online]. Available: https://www.smartme.co.uk/meter-data.html#:~:text=Timing%20of%20Data,ready%20for%20viewing%20by%209am.. [Accessed January 2023].

A.1 Scenario Figures

A.1.1 Job loss

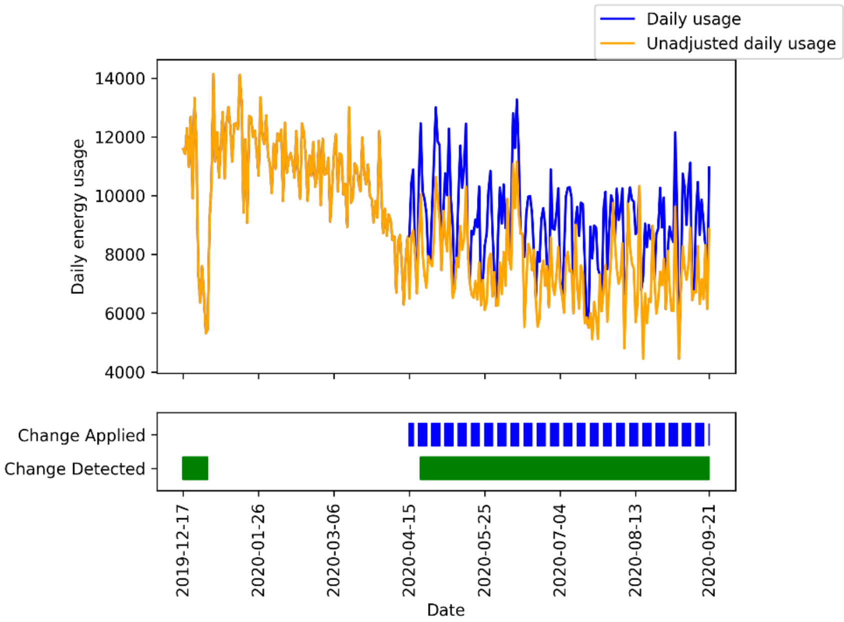

Figure 14: Daily adjusted and unadjusted energy usage for job loss scenario with applied change days shown in blue. The days the model detects as having changes are shown in green.

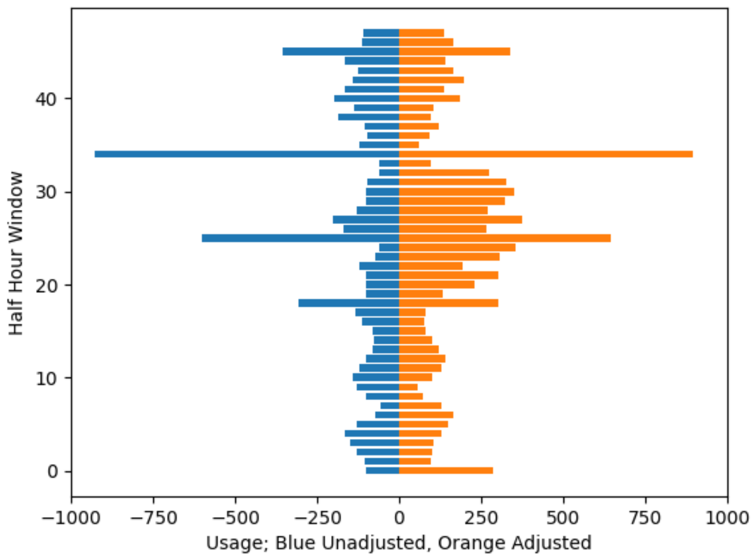

Figure 15: Energy usage before and after job loss usage feature applied for an example day.

A.1.2 In-home carer

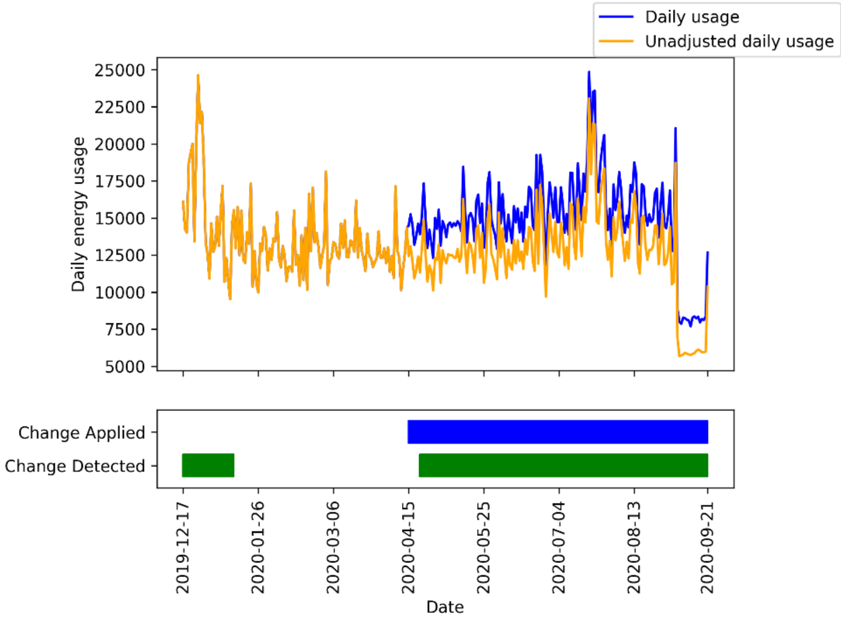

Figure 16: Daily adjusted and unadjusted energy usage for in-home carer scenario with applied change days shown in blue. The days the model detects as having changes are shown in green.

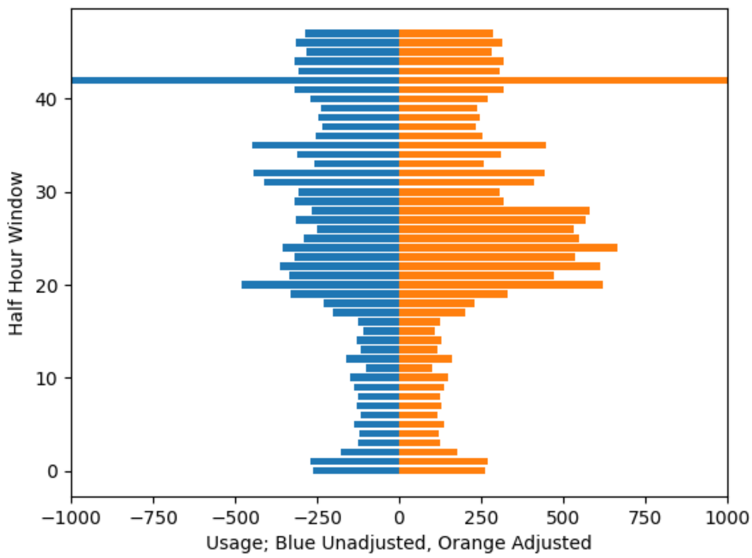

Figure 17: Energy usage before and after in-house carer usage feature applied for an example day.

A.1.3 Dementia

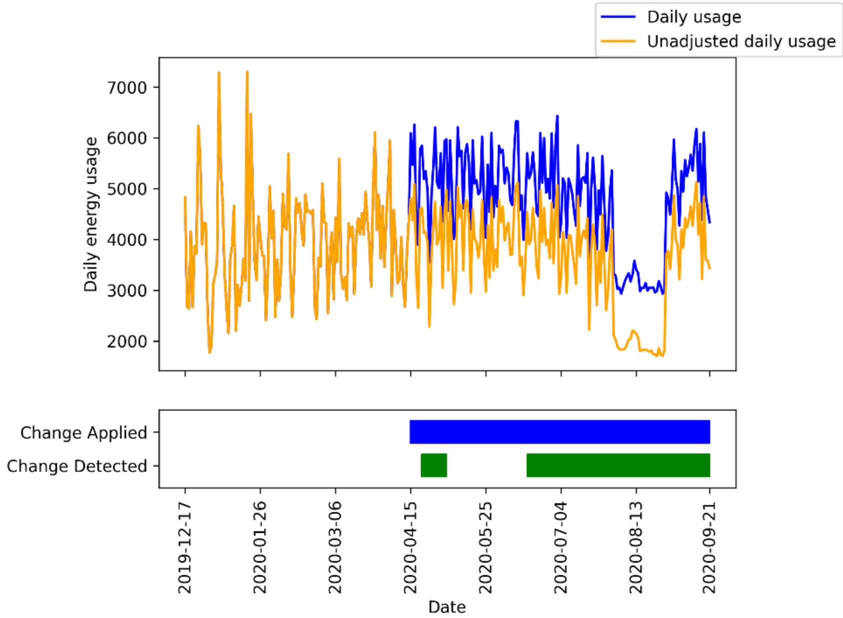

Figure 18: Daily adjusted and unadjusted energy usage for dementia scenario with applied change days shown in blue. The days the model detects as having changes are shown in green.

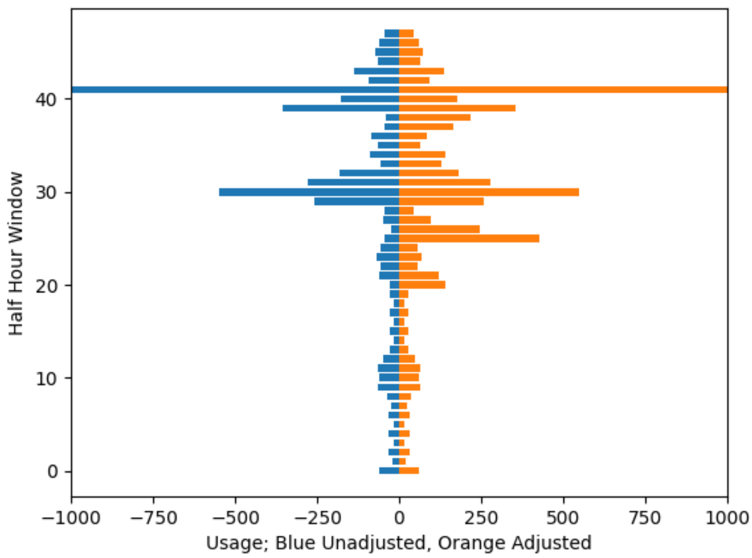

Figure 19: Energy usage before and after dementia usage feature applied for an example day.

A.1.4 Struggle to pay

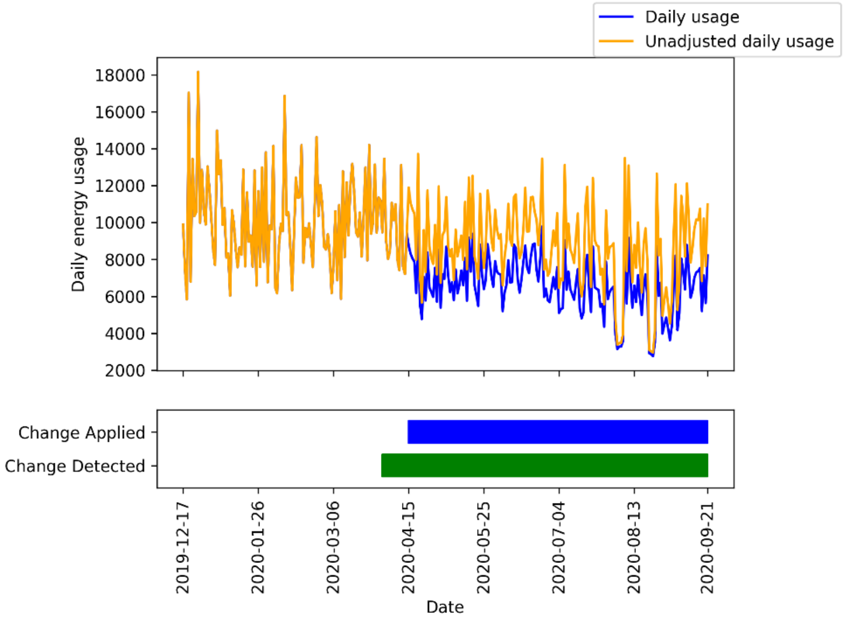

Figure 20: Daily adjusted and unadjusted energy usage for struggle to pay scenario with applied change days shown in blue. The days the model detects as having changes are shown in green.

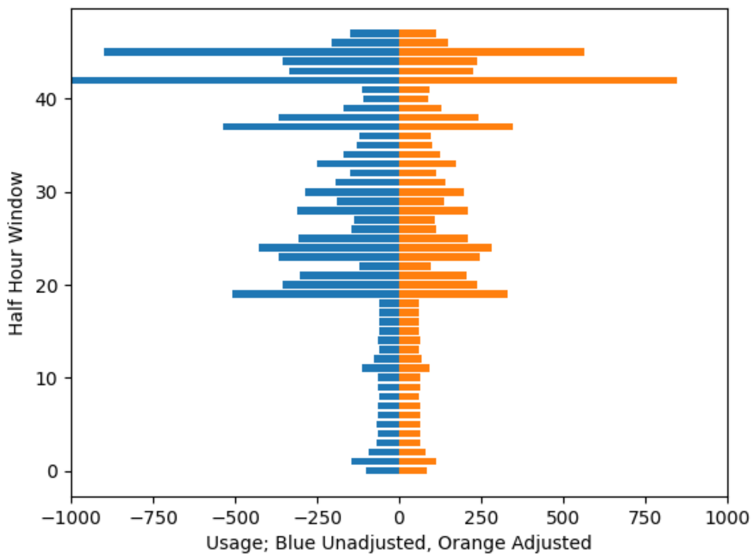

Figure 21: Energy usage before and after struggle to pay usage feature applied for an example day.

A.1.5 New child

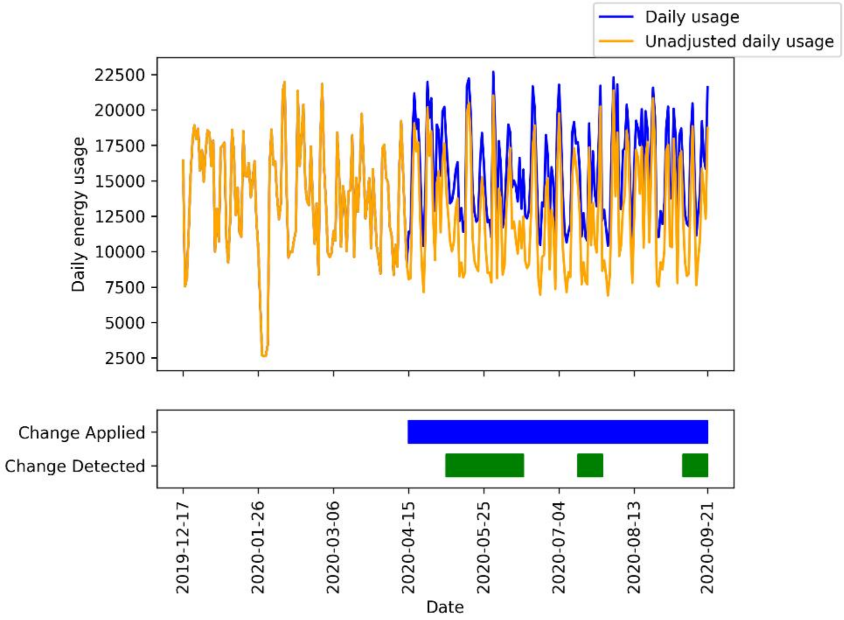

Figure 22: Daily adjusted and unadjusted energy usage for new child scenario with applied change days shown in blue. The days the model detects as having changes are shown in green.

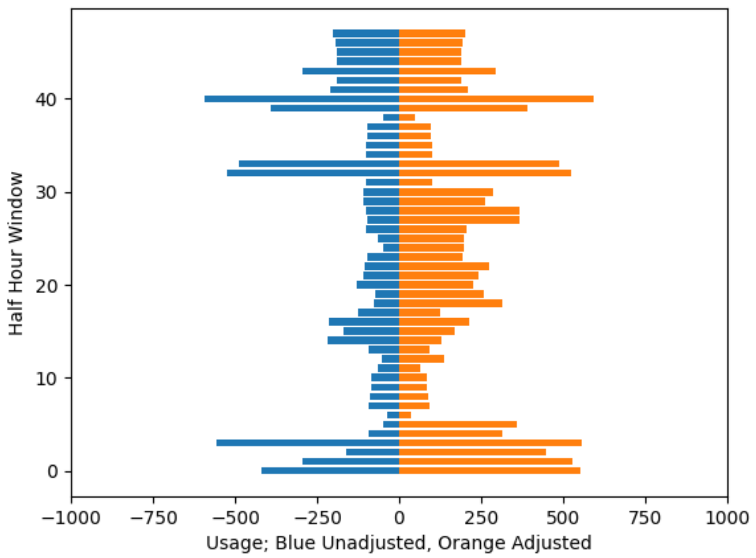

Figure 23: Energy usage before and after new child usage feature applied for an example day.

A.1.6 Overnight medical device

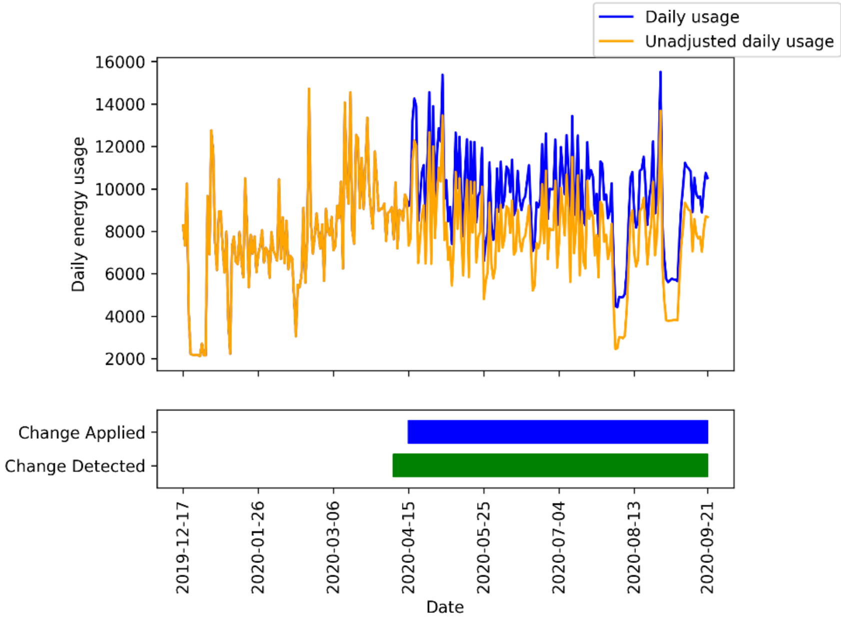

Figure 24: Daily adjusted and unadjusted energy usage for overnight medical device scenario with applied change days shown in blue. The days the model detects as having changes are shown in green.

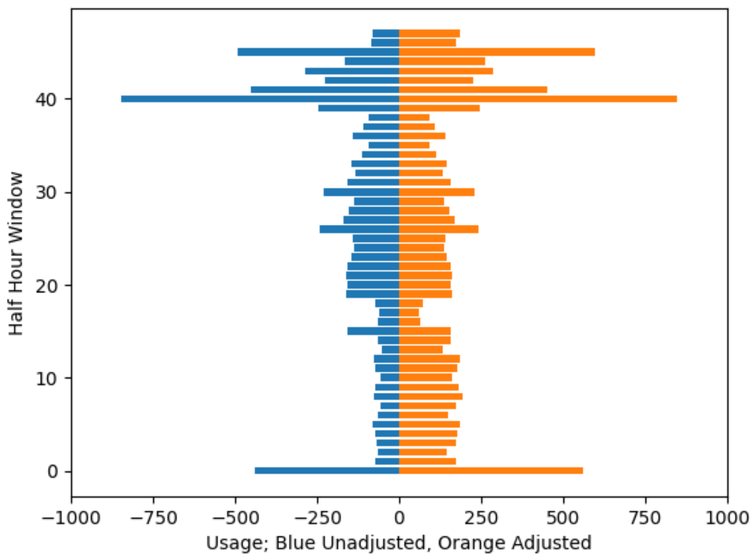

Figure 25: Energy usage before and after overnight medical device usage feature applied for an example day.

A.1.7 Cannot pay – weekly pay check

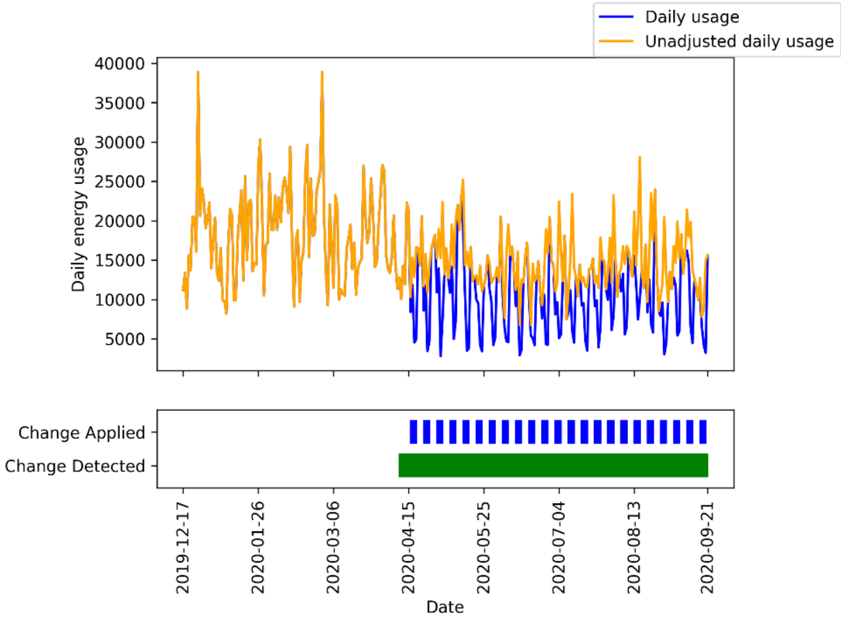

Figure 26: Daily adjusted and unadjusted energy usage for cannot pay - weekly pay check scenario with applied change days shown in blue. The days the model detects as having changes are shown in green.

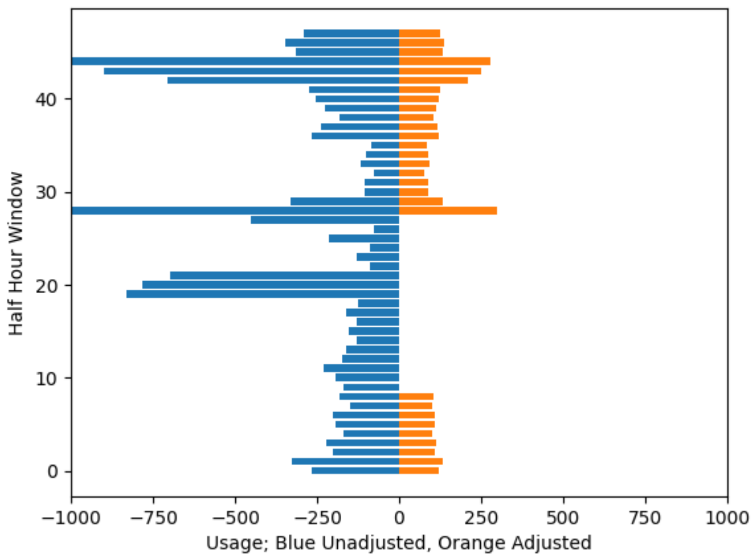

Figure 27: Energy usage before and after cannot pay – weekly pay check usage feature applied for an example day.

A.1.8 Cannot pay – monthly pay check

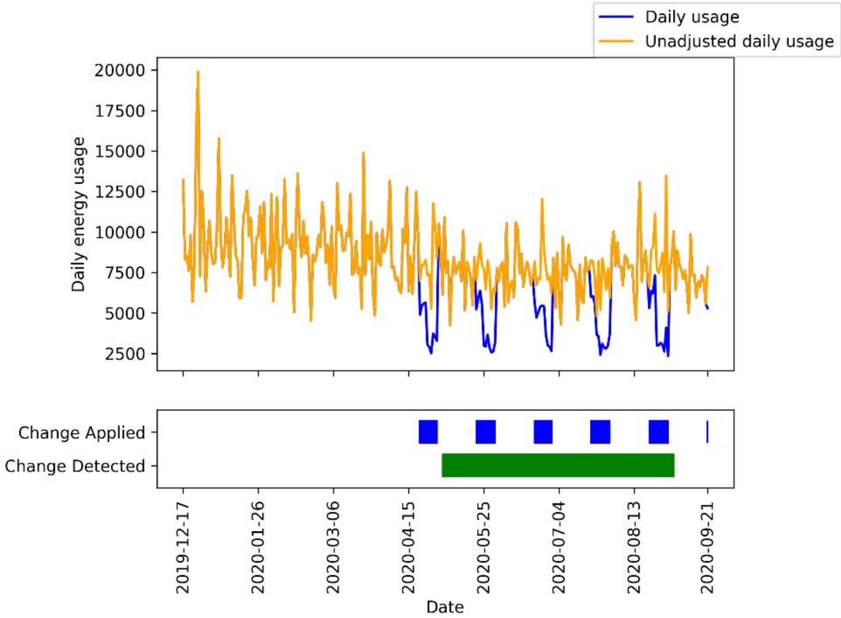

Figure 28: Daily adjusted and unadjusted energy usage for cannot pay - monthly pay check scenario with applied change days shown in blue. The days the model detects as having changes are shown in green.

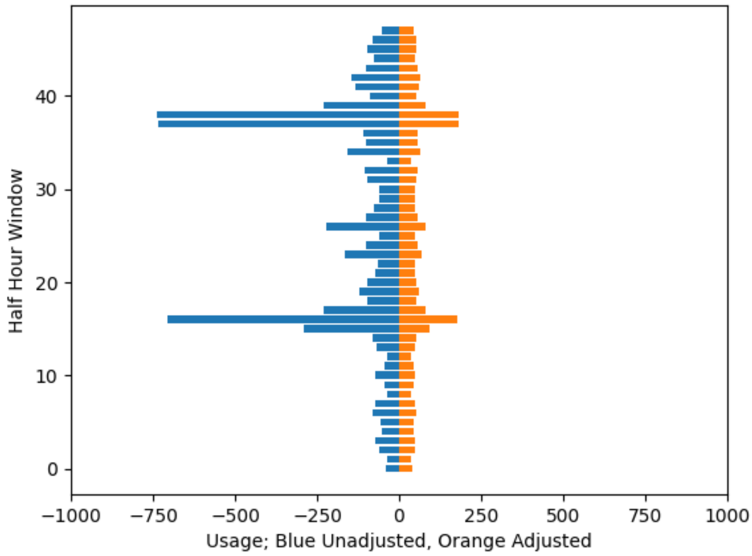

Figure 29: Energy usage before and after cannot pay – monthly pay check usage feature applied for an example day.

A.1.9 Post-hospital

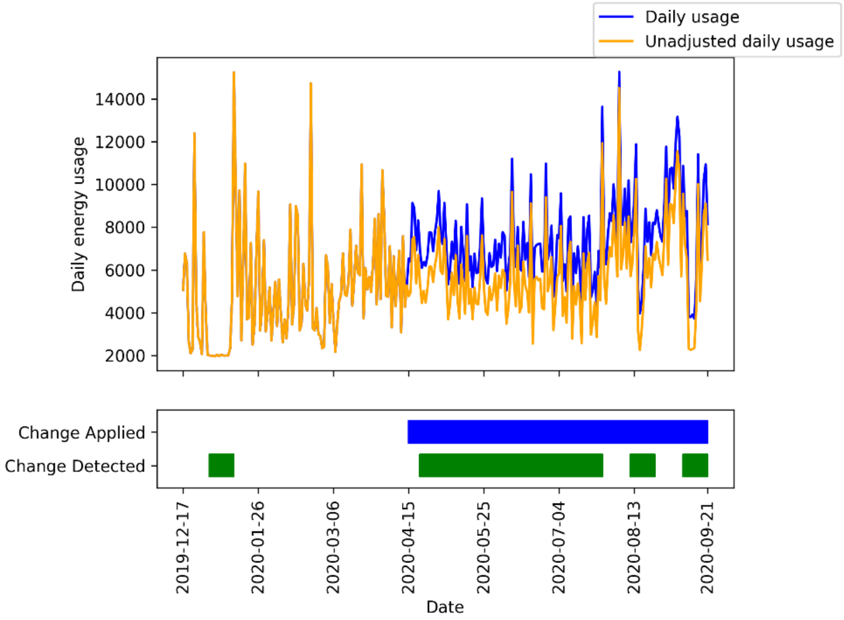

Figure 30: Daily adjusted and unadjusted energy usage for post-hospital scenario with applied change days shown in blue. The days the model detects as having changes are shown in green.

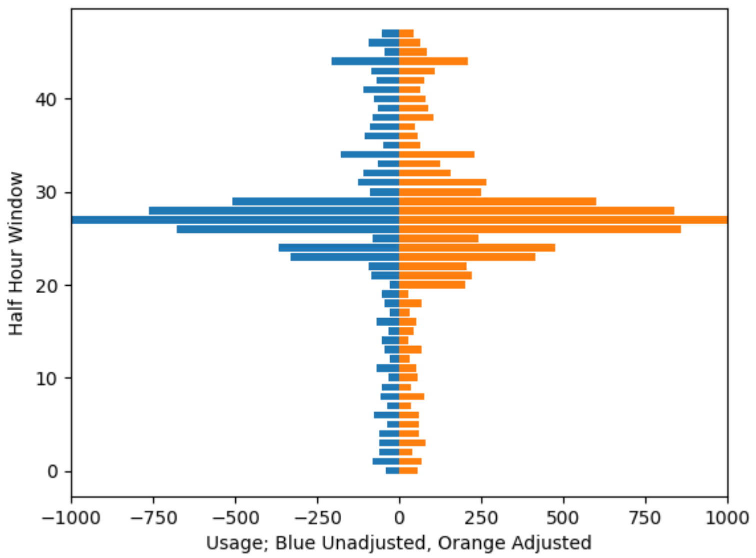

Figure 31: Energy usage before and after post-hospital usage feature applied for an example day.

A.1.10 Reduction in storage heater usage

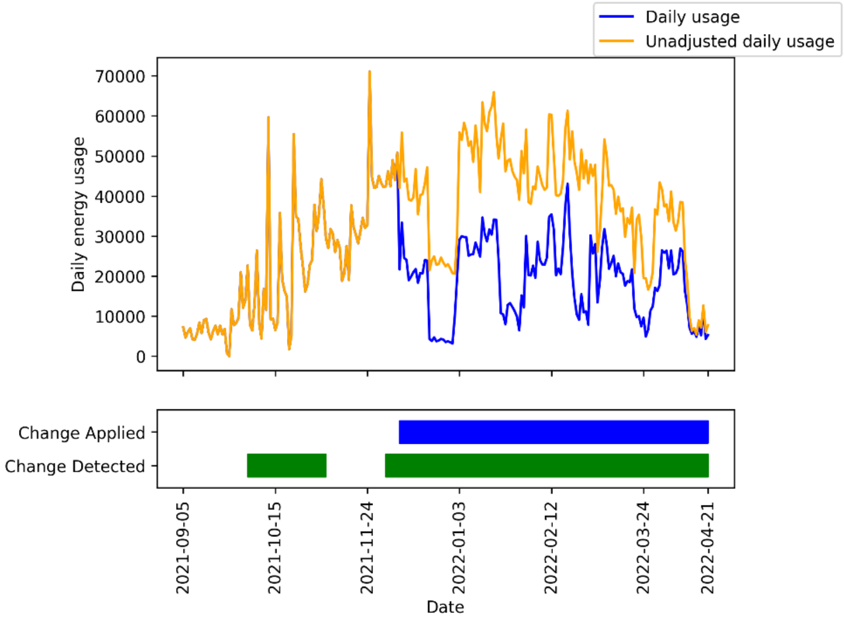

Figure 32: Daily adjusted and unadjusted energy usage for reduction in storage heater usage scenario with applied change days shown in blue. The days the model detects as having changes are shown in green.

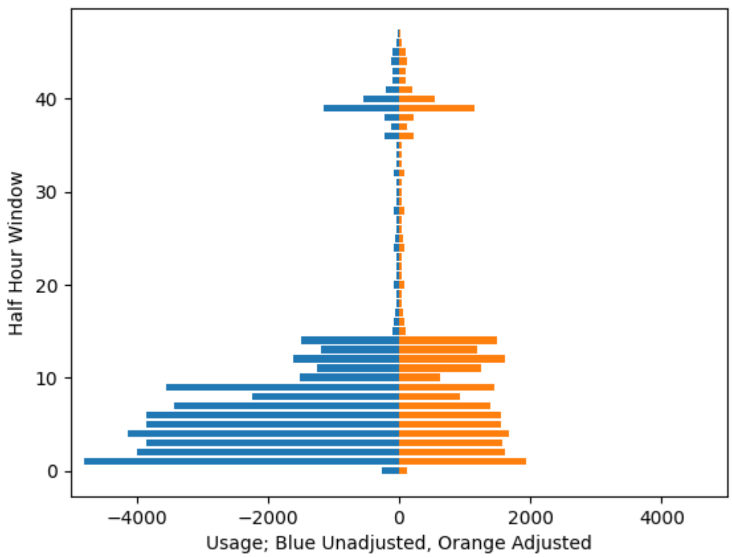

Figure 33: Energy usage before and after reduction in storage heater usage feature applied for an example day.