Saving Electrical Energy in Commercial Buildings

by

Ryan Case

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Mathematics
in
Computer Science

Waterloo, Ontario, Canada, 2012

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

With the commercial and institutional building sectors using approximately 29% and 34% of all electrical energy consumption in Canada and the United States, respectively, saving electrical energy in commercial and institutional buildings represents an important challenge for both the environment and the energy consumer. Concurrently, a rapid decline in the cost of microprocessing and communication has enabled the profileration of smart meters, which allow a customer to monitor energy usage every hour, 15 minutes or even every minute. Algorithmic analysis of this stream of meter readings would allow *1*) a building operator to predict the potential cost savings from implemented energy savings measures without engaging the services of an expensive energy expert; and *2*) an energy expert to quickly obtain a high-level understanding of a building's operating parameters without a time-consuming and expensive site visit. This thesis develops an algorithm that takes as input a stream of building meter data and outputs a building's *operating parameters*. This output can be used directly by an energy expert to assess a building's performance; it can also be used as input to other algorithms or systems, such as systems that *1*) predict the cost savings from a change in these operating parameters; *2*) benchmark a portfolio of building; *3*) create baseline models for measurement and verification programs; *4*) detect anomalous building behaviour; *5*) provide novel data visualization methods; or *6*) assess the applicability of demand response programs on a given building. To illustrate this, we show how operating parameters can be used to estimate potential energy savings from energy savings measures and predict building energy consumption. We validate our approach on a range of commercial and institutional buildings in Canada and the United States; our dataset consists of 10 buildings across a variety of geographies and industries and comprises over 21 years of meter data. We use K-fold cross-validation and benchmark our work against a leading black-box prediction algorithm; our model offers comparable prediction accuracy while being far less complex.

Acknowledgements

I would like to thank my advisor, Dr. S. Keshav, for his patient support, criticism and advice throughout my graduate studies. His commitment to the success of his students goes far beyond his professional duties; I consider myself very fortunate to have had him as an advisor.

I wish to thank my co-supervisor, Dr. Catherine Rosenberg, for providing useful suggestions, input and feedback throughout my research, and to Dr. Ian Rowlands and Dr. Daniel Lizotte for serving as readers for my thesis.

I owe a great intellectual debt to the people at Pulse Energy; their input and contributions formed the foundation of this thesis. I benefited greatly from conversations with Owen Rogers, Chris Porter, Jason Kok, James Christopherson, Kevin Tate, Brent Bowker, Neil Gentleman, Chuck Clark, Craig Handley and undoubtedly many others to whom I apologize for their omission.

Dedication

To my parents.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

Three hundred and fifty three billion dollars (U.S.) was spent on 3,724 TWh of electricity in the United States in 2009 [1]. In Canada during the same year, approximately 500 TWh of electricity was consumed, producing a total of 85.61 megatonnes of carbon dioxide [2] [3]. With commercial and institutional buildings responsible for 29% of 2009 electricity use in Canada [2] and 34% in the United States [1], saving energy in the commercial and institutional building sectors represents an important challenge for both the environment and the energy consumer.

This challenge coincides with the rapid decline in the cost of microprocessing and communication that has enabled the recent proliferation of advanced ("smart") metering infrastructure: as of 2010, approximately 14% of residential and 11% of commercial customers had smart meter installations in the United States [4]. In Ontario alone, 3.36 million customers out of a total of 4.75 million had smart meters in 2009, with the installed capacity rising to 4.57 million by 2010 [5]. Instead of receiving a consolidated bill from the utility each month, the customer can now learn of his or her usage every hour, every 15 minutes or even every minute. We believe an algorithmic analysis of this stream of meter readings can enable energy savings in commercial and institutional buildings. This thesis contains the development, validation and results of such analysis.

1.1 Definitions

We define the output of a stream of meter readings to be *meter data*, a plot of meter data over time to be a *load profile*, and a load profile over the course of one day to be

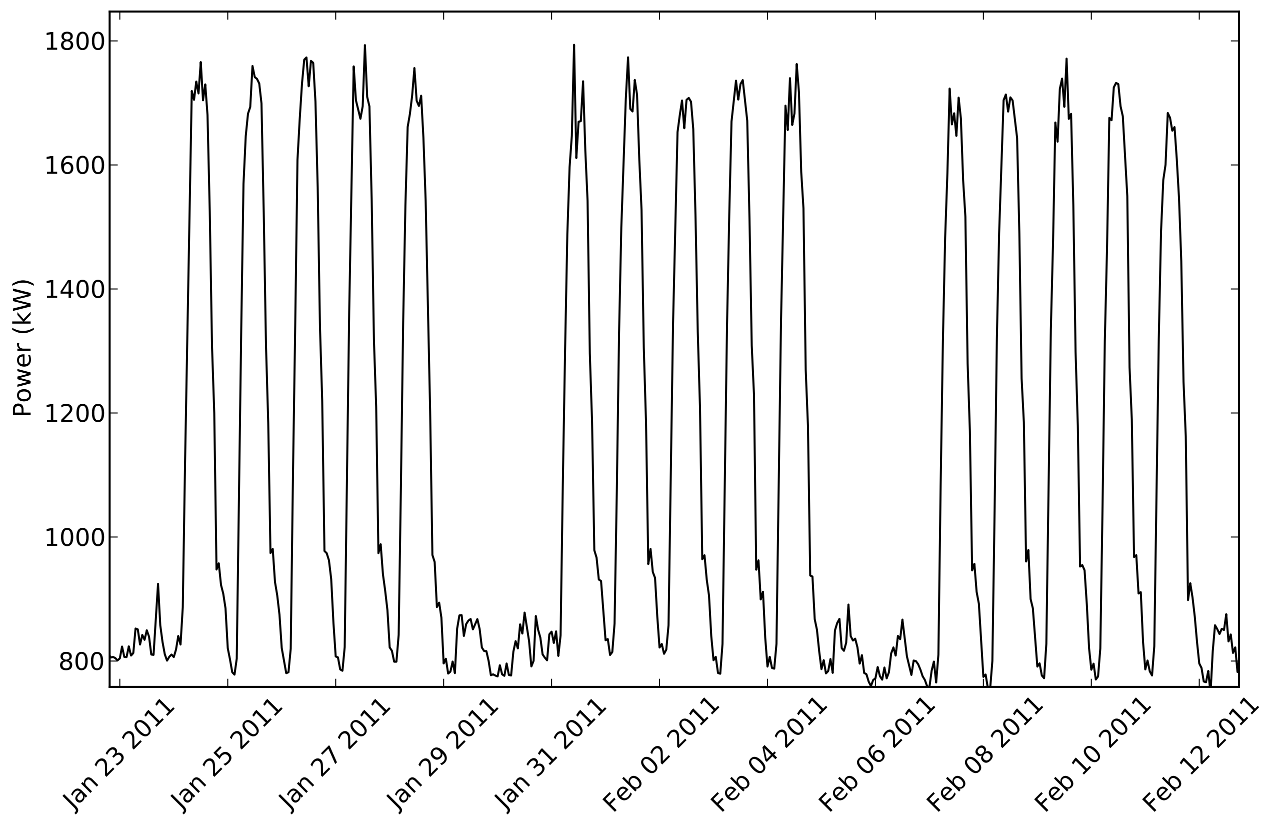

Figure 1.1: Load profile of an office building from our dataset (Office 1).

a *daily load profile*. Consider the load profile in Figure 1.1. This building's daily load profile is relatively consistent during each weekday, and assumes a different shape on the weekend. If this sample is an indication of the building's overall electrical energy consumption pattern, then the building appears to have two *operating modes*, or groups of days which, given the same weather pattern, have similar daily load profiles. Recognizing operating modes is fundamental to understanding the building's energy consumption patterns, and can be used to identify *energy savings measures* (ESMs). ESMs are actionable items which can be undertaken to conserve electrical energy in a building. Categories of ESMs are retrofits (upgrading equipment), operational (changing equipment operation), or behavioural (changing the occupants' behaviour). The process of undertaking ESMs involves several stakeholders, which we discuss in Section 1.3.

1.2 Building Operating Modes

Consider Figure 1.2, which shows different daily load profiles. Are there four modes of operation, or three? Identifying the operating modes of a building is an ill-posed clustering problem. Arguably there are three modes, implying the statuatory holiday is in the same mode as the weekend. This in itself may be an insight: how similar are statuatory holidays to weekends from the perspective of energy consumption? The statuatory holiday in this case is similar to the weekend, but there is a deviation in demand during the morning. Identifying the root cause of this difference could lead to an energy savings opportunity.

1.3 Stakeholders

The process of saving electrical energy in commercial and industrial buildings in the United States and Canada can involve many roles and stakeholders, each of which may be fulfilled by separate, or even multiple, parties. The following is a list of roles that may be filled during the energy saving process:

- Meter operator: installs and maintains the metering equipment.
- Data warehouse: stores and provides access to the meter data.
- Building operator: maintains the building, its equipment, and manages the corresponding budget.

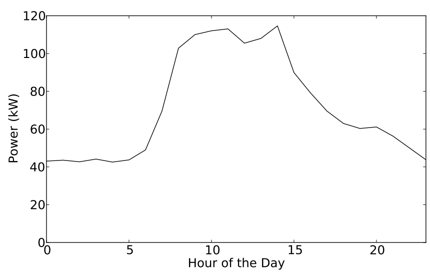

(a) Weekday load profile

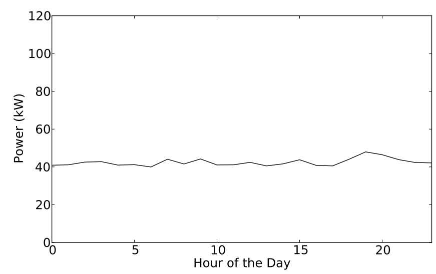

(b) Weekend load profile

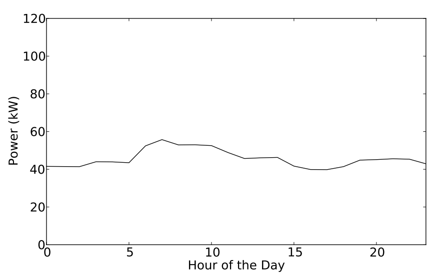

(c) Statuatory holiday load profile

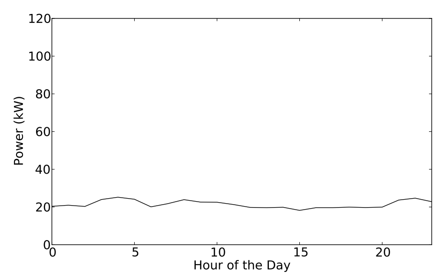

(d) Summer holiday load profile

Figure 1.2: Load profiles and operating modes in a high school from our dataset (School 1).

- Energy expert: uses its domain knowledge, proprietary tools, tools provided by the software provider, and knowledge of the building provided by the building operator to create a list of candidate ESMs or "recommendations" for the building operator.

- Software provider: provides software tools for meter data visualization, analysis and reporting for use by energy experts or building operators.

- ESM implementor: implements the chosen ESM(s).

We assume any omitted roles will have a negligible effect on the discussion. Examples of omitted roles could be the owner of the building, or the party which evaluates the results of the ESM.

We illustrate how these roles might be fulfilled with an example: a case study [6] from B.C. Hydro's "Continuous Optimization" program [7]. This program focuses on re-commissioning and improving efficiency in commercial buildings:

- Meter operator: B.C. Hydro, a Canadian electric utility, upgrades metering equipment at the building and connects the meter to its metering infrastructure; alternatively the building operator installs the meter if a BC Hydro meter is not already in place [8].

- Data warehouse: BC Hydro accesses the data through its metering infrastructure; alternatively, the customer may be required to provide data collection services, eg. through their Internet connection. The data is then stored and used by the energy expert or software provider, and potentially by BC Hydro itself.

- Building operator: Jawl Properties, who also own the building(s)

- Energy expert: SES Consulting

- Software provider: SES Consulting, sub-contracts the software requirements to Pulse Energy

- ESM implementor: SES Consulting oversees the implementation process, which is paid for by the building operator

1.4 Problem Statement

The current practice for identifying energy savings involves two primary parties: *1*) building operators, who have a deep knowledge of the building's operation but lack the resources, such as time or expertise, to investigate energy savings; and *2*) energy experts, who provide expertise and resources as a service, and may lack detailed knowledge of how a particular building is run. We propose to analyze a stream of meter readings that would allow:

1. A building operator to determine the potential cost savings from implementating energy savings measures (ESMs) without engaging the services of an expensive energy expert

2. An energy expert to quickly obtain a high-level understanding of a building's operating parameters without a time-consuming and expensive site visit

Therefore, although the selection of appropriate ESMs for a particular building requires a degree of understanding of its operations that is beyond the scope of this work, we believe that there are still substantial benefits from an algorithmic approach to meter stream data analysis.

Importantly, such an approach must take into account the fact that a building operates in one of several modes, and that its energy consumption varies dramatically with operating mode. Moreover, the approach must model building energy use in terms of human-understandable operating parameters (such as base and peak load as a function of the ambient temperature) so that an energy expert can assess the gains from changes in these operating parameters due to an ESM. These two criteria serve to motivate the work presented in this thesis.

In summary, the goal of this thesis, as illustrated in Figure 1.3 is to develop an algorithm that takes as input a stream of building meter data and that outputs the building's operating parameters. This output can be used directly by an energy expert to rapidly assess the building's performance. It can also be used as input to other systems, for example systems that *1*) predict the cost savings from a change in these operating parameters; *2*) benchmark a portfolio of building; *3*) create baseline models for measurement and verification programs; *4*) detect anomalous building behaviour; *5*) provide novel data visualization methods; or *6*) assess the applicability of demand response programs on a given building.

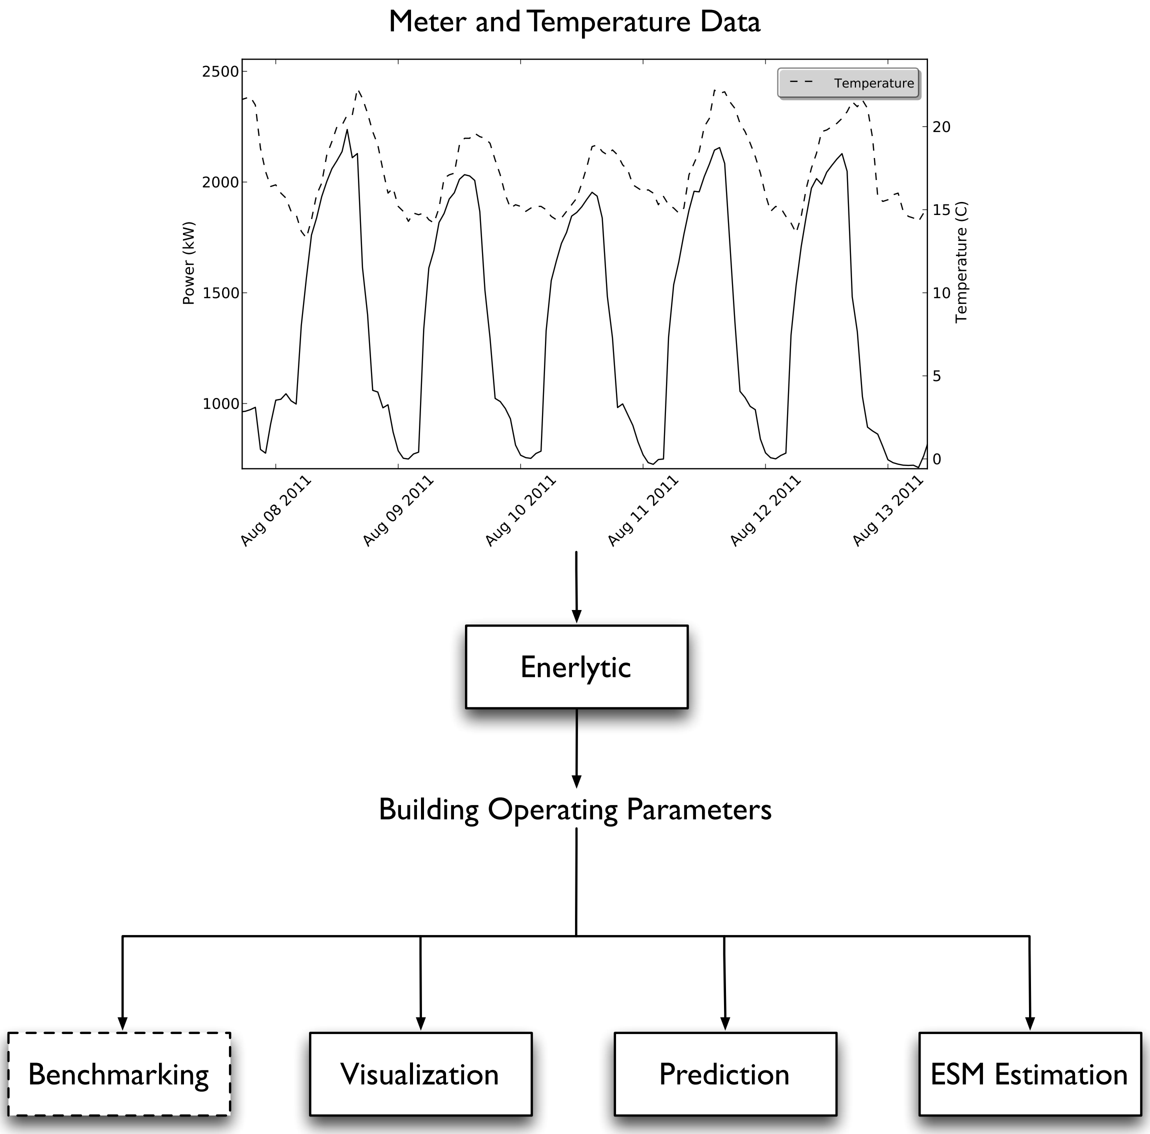

Figure 1.3: Goal of the thesis.

1.5 Assumptions

We assume that:

1. For sake of simplicity, meter readings are taken at an hourly interval. There are no technical limitations preventing this work to being extended to meter readings at other intervals.
2. The outdoor weather can be approximated with hourly dry-bulb temperature measurements from a nearby weather station.
3. If the building parameters are being used for prediction, temperature data can also be predicted for the corresponding period.
4. Predictions need not be made for operating modes that are not in the input data

1.6 Solution Overview

We propose an algorithm, *Enerlytic*, which partitions meter data into operating modes using a process we call *labelling*, and creates a model of a building's electrical energy use during each operating mode as a function of the outdoor air temperature. The model parameters are amenable to human interpretation. We develop a novel algorithm to extract the *peak* and *base* periods of each day, and model these loads as a function of the outdoor temperature using piecewise linear regressions. We validate the output of Enerlytic, the building's *operating parameters*, by using them to predict over 21 years of meter data on 10 different buildings across North America. Our prediction accuracy is comparable to a competitive black-box method [9]. We also show how these operating parameters can be used as input to other algorithms, for example to estimate the potential for energy savings in a building.

1.7 Thesis Organization

In Chapter 2 we provide an overview of the related literature. Chapter 3 describes the development of our solution, and Chapter 4 discusses the results. Chapter 5 provides additional discussion, and Chapter 6 provides concluding remarks, as well as limitations to our solution and areas of future work.

Chapter 2

Related Work

This thesis develops an algorithm that takes as input a stream of building meter data and outputs a model which can provide an energy expert with a high-level understanding of a building. The output also serves as input to other algorithms, for example for estimating the potential cost savings from implementing energy savings measures, or for predicting a building's energy use. We now survey prior work with similar goals.

Advanced tools for building modelling and simulation, such as EnergyPlus [10] [11] are well established in the building design and efficiency community. These tools are useful for advanced analysis of building energy consumption, particularly when no meter data is available. However, their complexity and need for physical building data make them difficult to calibrate [12] and ill-suited to meet our goals.

Creating energy models from meter data and weather parameters—such as temperature and humidity—for predicting energy consumption has also been well studied. This problem was investigated in detail for the 1993 ASHRAE "Great Energy Predictor Shootout" [13], which resulted in many robust methods for predicting energy use. Some entries, such as the winning entry from David MacKay [14], involved systems to determine which inputs were most relevant for accurate prediction, illustrating possible correlations between inputs and output. Many of the top performers, including the winning entry, are neural-network-based models that provide little insight into the building's energy consumption patterns. The black-box nature of these prediction algorithms make them unsuitable for our purpose; however, we use a more recently developed black-box prediction model for benchmarking the prediction accuracy of our work. This model is based on kernel regression [9], and performs comparably to the top prediction models in the ASHRAE shootout based on prediction accuracy.

"Grey-box" models, which contrast with "black-box" models, use the most similar approach to our own. These models use parameters that are related to the physics of the building or its energy consumption patterns. Some, like Lee and Brown [15], do not use meter data as input but instead require other parameters of the physical building, such as the floor plan. Others use well established methodology for fitting piecewise linear regression models to the relationship between outdoor dry-bulb air temperature and average daily energy consumption; for example, the work by Kissock et al [16]. The piecewise linear regression model fits two linear regressions which are continuous at a single change-point—the point where one regression changes to the other, as illustrated in Figure 2.1. The resulting regression is in a shape similar to a hockey stick, with the bend occurring at the change-point. If a building has electric heating or cooling, a building's energy consumption follows a linear or uncorrelated relationship with outdoor temperature until some temperature, known as the "base temperature" is reached, and then this relationship changes, becoming strongly correlated with temperature. The base temperature has a direct relationship with the set point of the building's heating or cooling system, so this can be a useful parameter to identify. Further, the slopes of the regressions indicate the increase in the amount of energy expended per degree increase in the outdoor air temperature. These slopes, known as the heating and cooling gradients, serve to identify physical elements of the building, such as thermal components (the building envelope, and heating and cooling systems).

Sever et al [17] show how grey box models can be used to estimate energy savings from model parameters. However, they make two assumptions that are not compatible with our goals. First, the building is assumed to have a single operating mode. This dramatically limits the insight which the model can provide, since it does not capture this fundamental aspect of the building's energy consumption. Second, and most importantly, their model input is the average daily energy consumption as opposed to hourly data. This prevents analysis of the daily load profile, such as the identification of peak or base periods, and makes it impossible, for example, to estimate the energy savings due to an adjustment in the average peak load or by shortening the peak load period.

Other grey-box methods also identify the peak or base periods (or "occupied and unoccupied" periods) but do not derive insight from the building's energy consumption with temperature, limiting the insight the models are able to provide. Mathieu et al [18] predict hourly power using a time-of-week indicator variable and either a linear or local linear regression with outdoor temperature. They demonstrate how their model can be used to estimate the effectiveness of different demand-response strategies through peak-load shifting. However, there are several limitations in their work. The user must manually separate a building's "occupied" mode from its "unoccupied" mode. There is no modelling of the

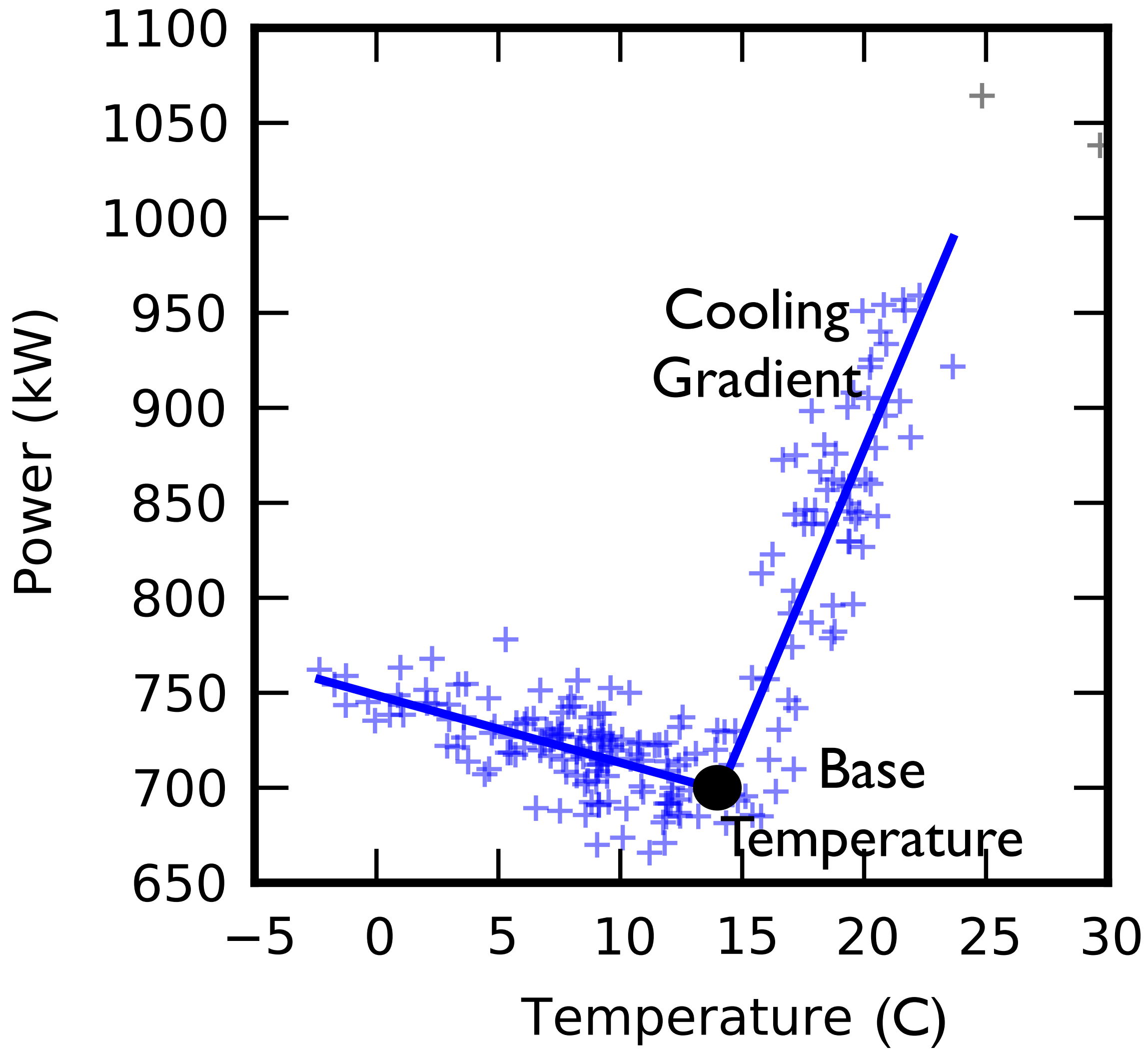

Figure 2.1: Energy signature.

different operating modes of the building, and the temperature modelling cannot be used to estimate energy savings.

Cherkassky et al [19] cluster a building's daily load profile into four periods, then use either a bin-based [20] or non-linear regression with temperature data to predict future consumption. They also separate weekends from weekdays to improve the accuracy of their predictions, thus identifying different operating modes in the two buildings in their test set. However, although the output of their model could potentially be adapted to provide some insight into the building's energy consumption patterns, it is not clear how to manipulate such a model to obtain energy savings estimates or use it as input for other systems.

In summary, although there is no prior work which is directly suitable for our problem statement, our work does extend several prior approaches. We partition the daily load profile similar to Cherkassky et al [19], but instead of using Lloyd's algorithm [21] we develop non-iterative algorithm. We use temperature models like those proposed by Kissock et al [16] for estimating energy savings like Sever et al [17]. The output of our algorithm can be used as an input for estimating energy savings opportunities such as demand response, like Mathieu et al [18]. We now discuss the development of our algorithm, *Enerlytic*, which makes use of these concepts.

Chapter 3

Enerlytic

This chapter describes *Enerlytic*, an algorithm that takes as input hourly meter data and outputs a building's *operating parameters*. These parameters can be used by an energy expert to quickly gain a high-level understanding of a building's electrical energy consumption, and can be used as input to other algorithms and systems.

Enerlytic obtains information about a building's operating modes through a *labeller*—a function that maps a date into an operating mode based on a set of rules. If no labeller is given, Enerlytic uses a default labeller that labels weekdays and weekends; Enerlytic partitions the days in the meter data based on their operating mode. Enerlytic models hourly power consumption as a function of outdoor air temperature using piecewise linear regressions and a pre-processing algorithm we call *summarizing*. Summarizing divides each daily load profile into *peak*, *base*, and *ramping* periods and computes the mean of the peak and base periods, ignoring the ramping period. This acts as a de-noising process, which is illustrated in Figure 3.1.

Enerlytic runs an outlier detection algorithm to verify that the labelling process adequately partitioned the building's operating modes for the subsequent regression fitting. This removes the outliers and alerts the expert that outliers were removed. We use the output of Enerlytic, the building's operating parameters, to create predictions and estimate energy savings in Sections 3.4 and 3.5, respectively. The entire Enerlytic algorithm is displayed in Figures 3.2-3.5.

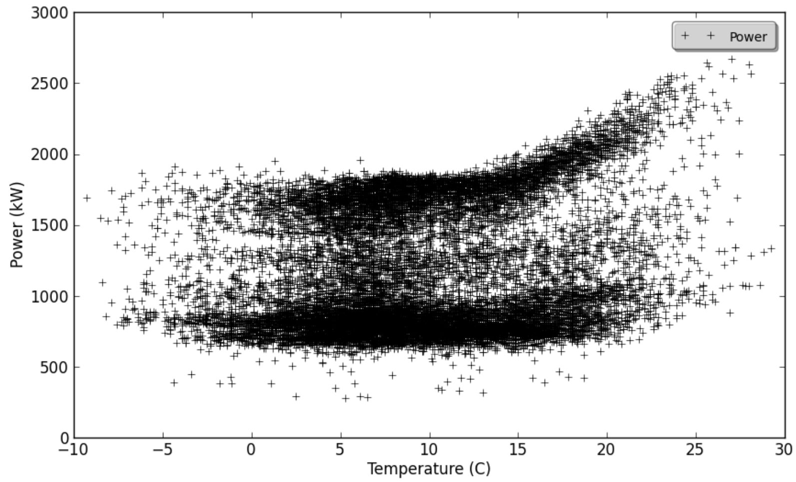

(a) Before summarizing

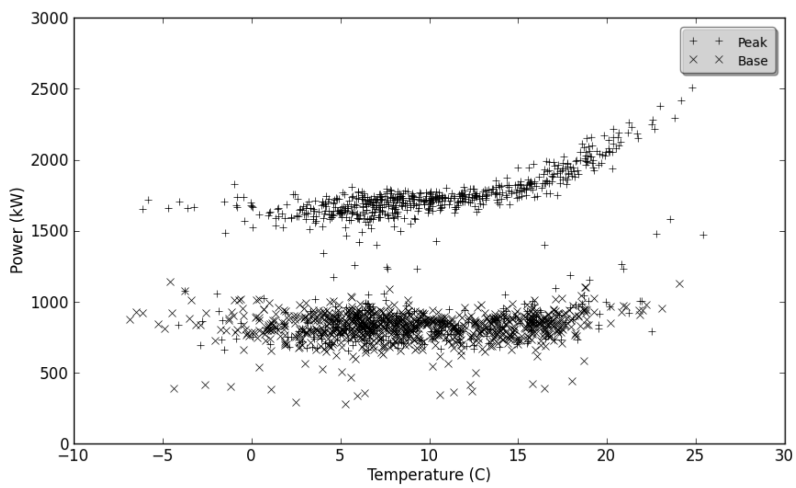

(b) After summarizing

Figure 3.1: Effect of summarizing. Each hourly measurement creates one data point, shown in Figure 3.1(a). Summarizing separates peak and base periods based on the time of day, and averages the measurements during these periods. Figure 3.1(b) shows the effect of this process, which is to de-noise the data for subsequent regression fitting.

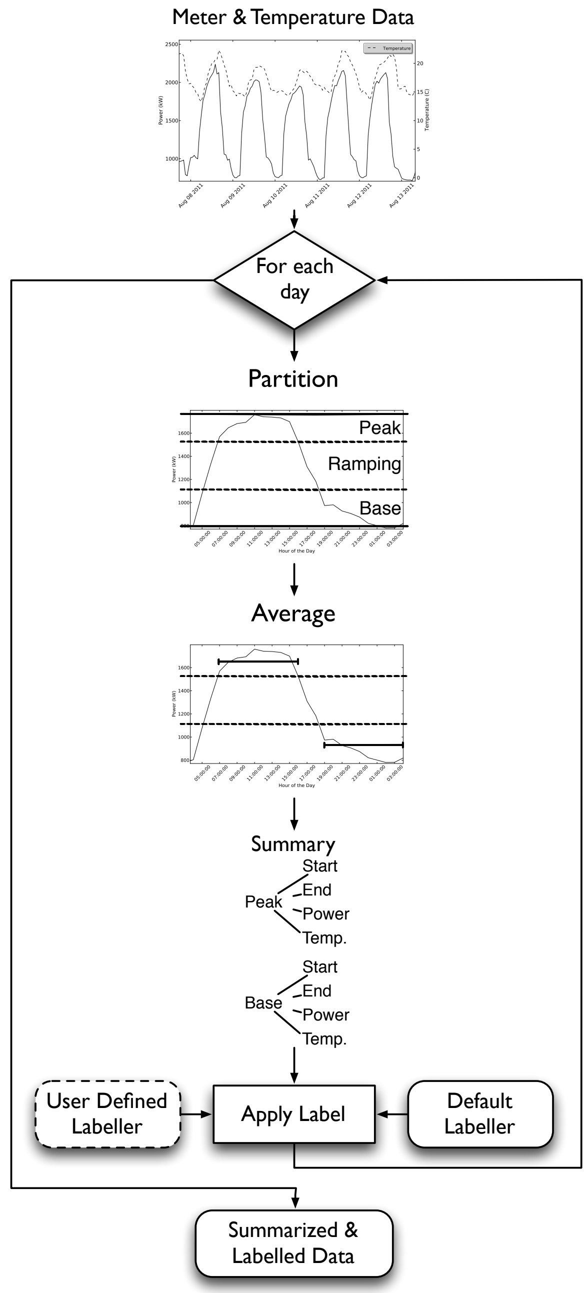

Figure 3.2: Enerlytic algorithm 1: summarizing and labelling.

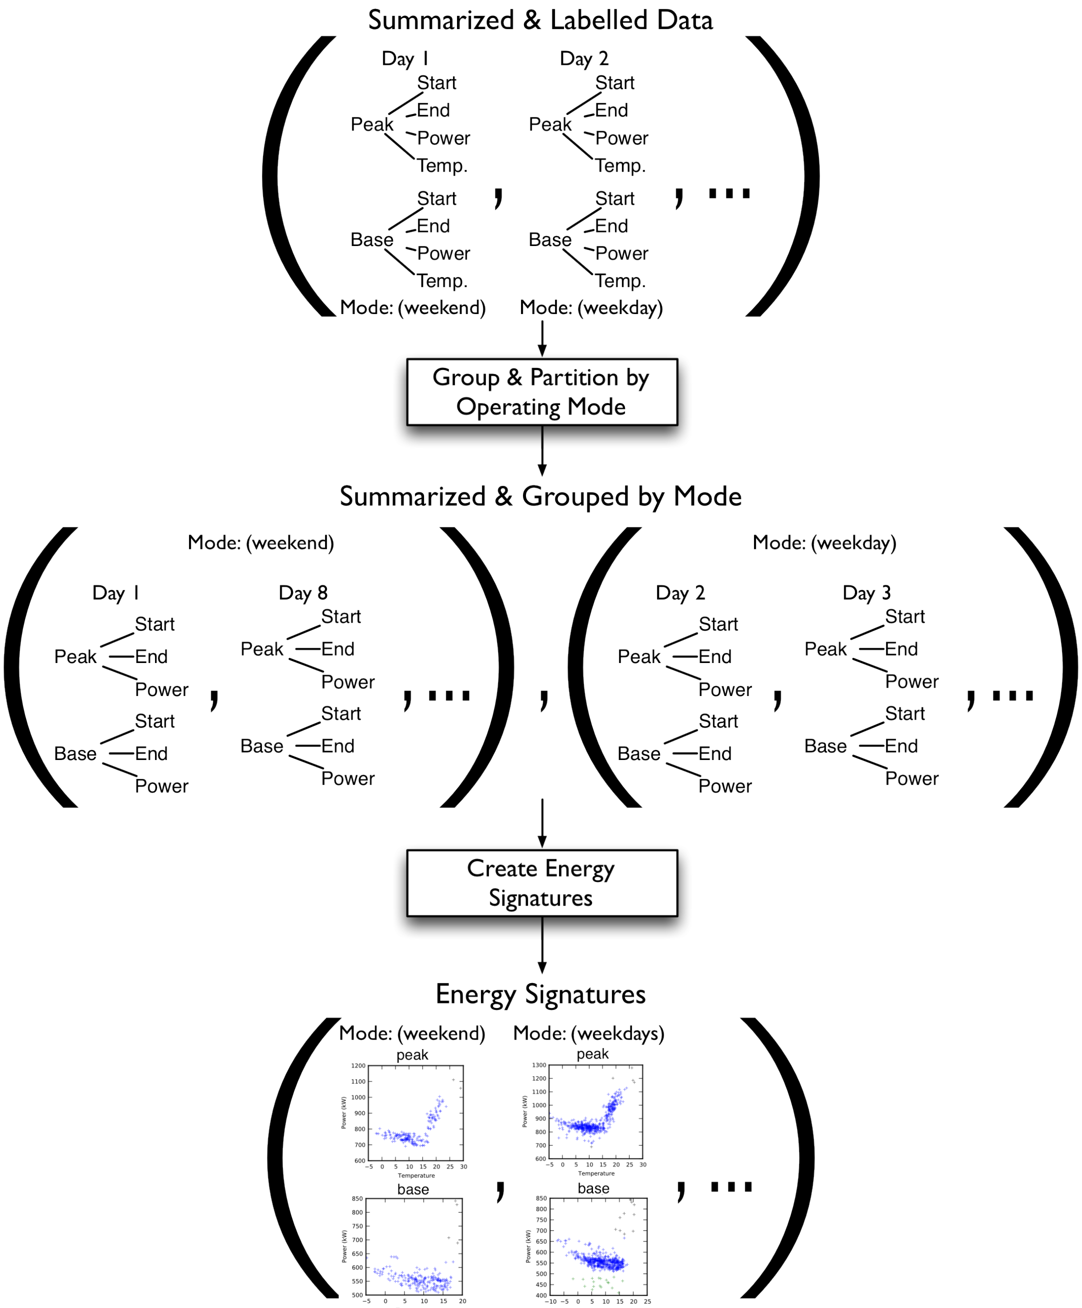

Figure 3.3: Enerlytic algorithm 2: creating energy signatures.

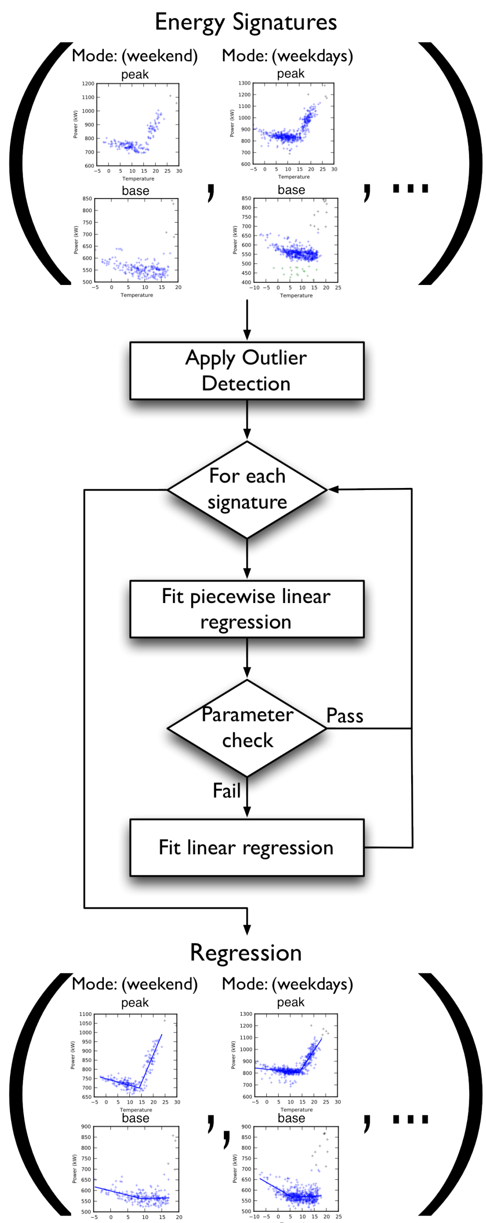

Figure 3.4: Enerlytic algorithm 3: Fitting regressions.

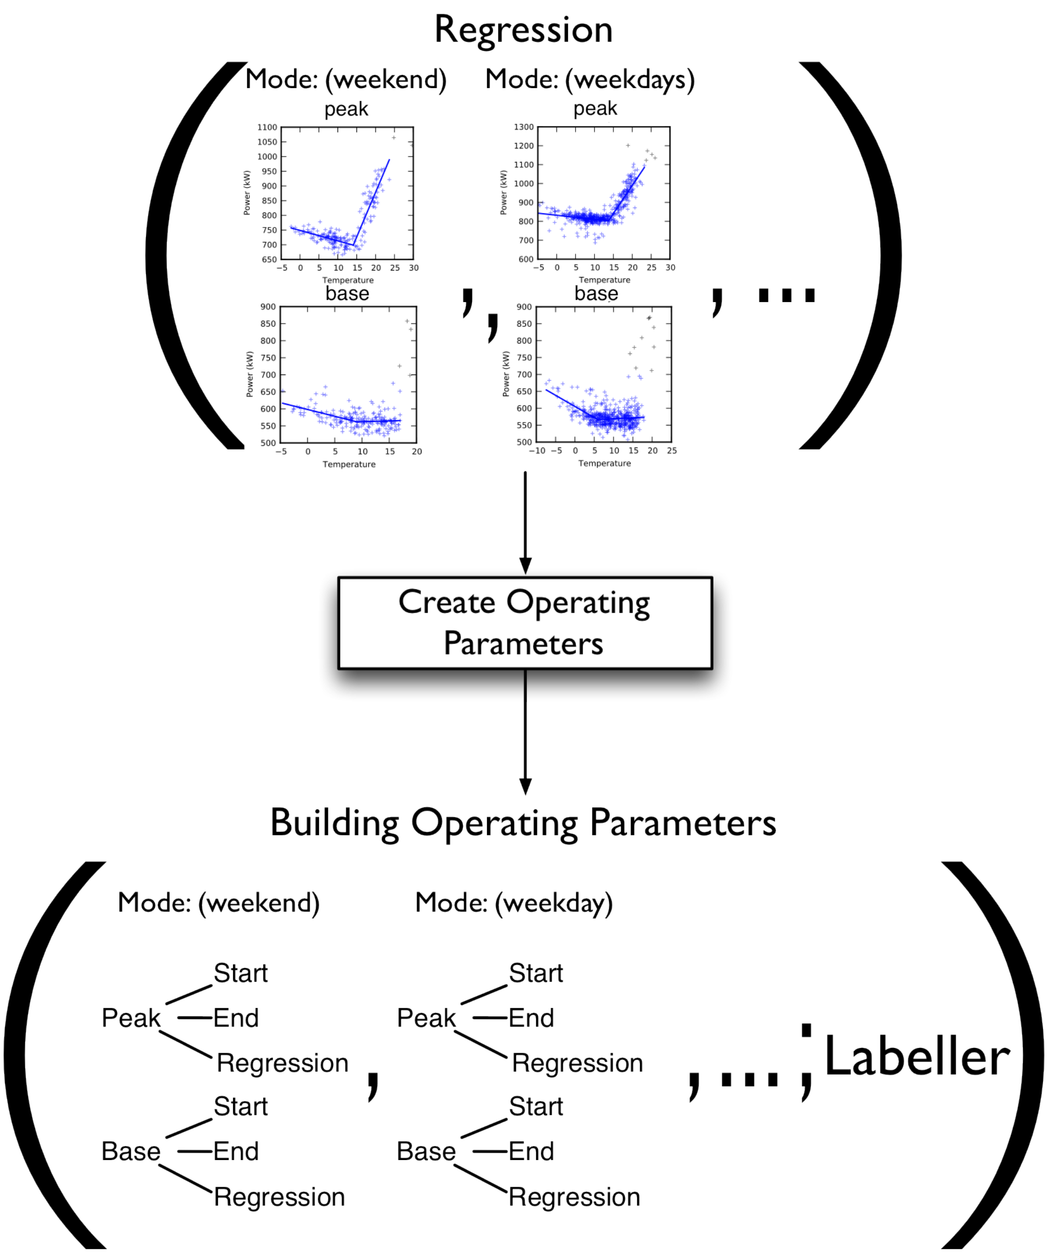

Figure 3.5: Enerlytic algorithm 4: Building operating parameters.

3.1 Summarizing

Summarizing is a pre-processing algorithm that takes hourly meter and temperature data and summarizes each day into *peak*, *base*, and *ramping* periods; these names are inspired by Cherkassy et al [19]. The periods have the following physical interpretation:

- Peak: the building is active
- Base: the building is inactive
- Ramping: the building is changing from active to inactive or vice-versa.

The summarizing algorithm defines the peak, base and ramping periods by dividing each daily load profile into a high and low period, ignoring mid-range values; this corresponds to the peak, base and ramping periods of a building, respectively. We refer to the range of power values in each period as a *bin*. For simplicity, we assume equal bin widths for each of the periods. The output of the algorithm is a list of daily *summaries*; these summaries contain the start and end time and the average power and temperature for both the base and peak periods. A building may oscillate between two periods (bins). We adjust for this by defining the peak period to start the first time a building enters the peak period and the last time it leaves the peak period; similarly for the base period. In the event that there is an oscillation between a peak and a base period, peak periods take precedence over base periods.

3.2 Labelling

Enerlytic obtains information about the building's operating modes through auxiliary *labellers*. Labellers assign each input date to a label corresponding to a building's operating mode. Enerlytic applies a labeller to each date in the training input to obtain the operating mode for that date. A labeller can be created *a priori* by a user and input to Enerlytic; there is also a default labeller that merely labels weekends and weekdays. Enerlytic then partitions the input, grouping the dates by operating mode. For each operating mode, Enerlytic converts the set of daily summaries into *energy signatures*—lists of average temperature and average power—for both the peak and base periods. Enerlytic also averages the start and end times in the peak and base periods; these average start and end times indicate when the peak or base periods typically start and end when a building is in a particular operating mode. The creation of the energy signatures can be seen in Figure 3.4.

Table 3.1: Regression parameter definitions.

Parameter	Definition
β	Y-intercept
T_b	Base temperature
m_h	Heating gradient
m_c	Cooling gradient
m_n	Neutral gradient

3.3 Regression

If the data is not partitioned appropriately, Enerlytic may fit regressions incorrectly due to Simpson's Paradox [22]. If two separate populations are combined because of inadequate partitioning, for example, due to inadequate labelling, the apparent correlation between temperature and power may be distorted. By testing for multiple populations before fitting the regressions, Enerlytic aims to provide a single meaningful regression model and exclude a set of outliers, instead of outputting a potentially meaningless regression model. We implement a method to test for multiple populations using Density-Based Spatial Clustering with Applications with Noise (DBSCAN) [23], a spatial clustering technique which does not require the number of clusters *a priori*. DBSCAN:

1. Finds regions where a given number of points occur within a given density;
2. Uses these regions as the seeds for the clusters;
3. Grows these clusters by adding points which are within the given density; and
4. Labels any points which are not clustered as outliers.

If DBSCAN identifies multiple clusters, Enerlytic chooses the cluster with the greatest number of points, and labels any other clusters as outliers. If any points are labelled as outliers, Enerlytic discards these data and provides a notification. Following outlier detection, Enerlytic fits a piecewise linear regression model to each energy signature. Figures 3.6(a) and 3.6(b) show examples of buildings with electric cooling and heating piecewise linear regression parameters, respectively. The regression parameters are tabulated in Table 3.1, and the piecewise linear regressions are defined in Equations 3.1 and 3.2 for heating and cooling, respectively.

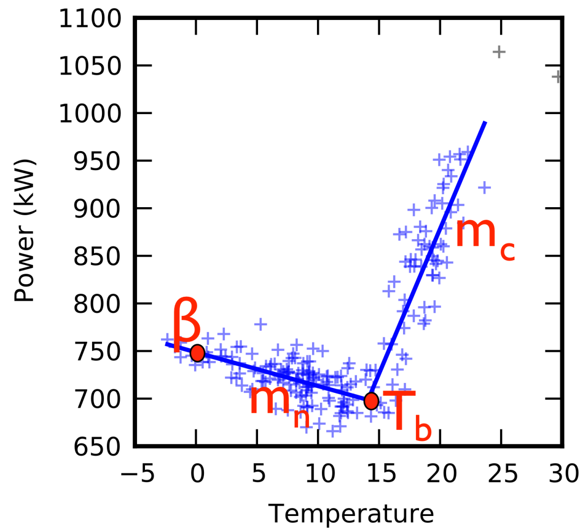

(a) Electric cooling

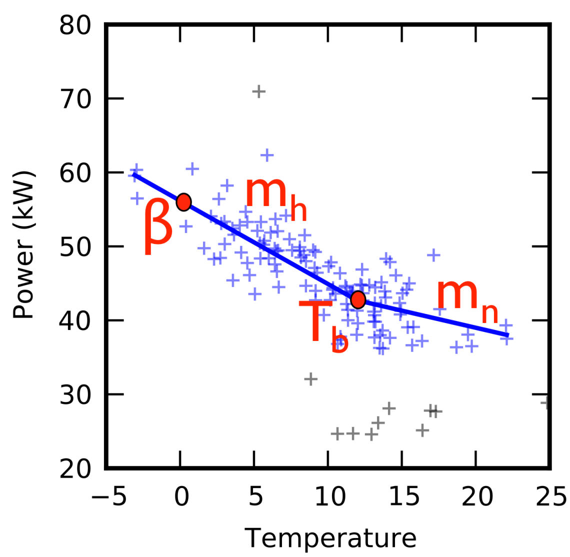

(b) Electric heating

Figure 3.6: Piecewise linear regression parameters. Figure 3.6(a) is taken from a hospital in our dataset (Hospital 1), and Figure 3.6(b) is from an Office building (Office 5).

Table 3.2: Acceptable piecewise parameter ranges

Heating	Cooling
$0 < T_b < 15$	$10 < T_b < 22$
$-0.15 < \frac{m_h}{\mu} < -0.006$	$-0.007 < \frac{m_n}{\mu} < -0.007$
$-0.1 < \frac{m_n}{\mu} < 0.006$	$0.007 < \frac{m_c}{\mu} < 0.06$

$$y_i = \begin{cases} \beta + m_h x_i + e_i & \text{if } x_i \leq T_b \\ \beta + (m_n - m_h)x_i + e_i & \text{if } x_i > T_b \end{cases} \tag{3.1}$$

$$y_i = \begin{cases} \beta + m_n x_i + e_i & \text{if } x_i \leq T_b \\ \beta + (m_c - m_n)x_i + e_i & \text{if } x_i > T_b \end{cases} \tag{3.2}$$

where y_i is the ith observation, x_i is the corresponding independent variable and e_i is the residual or error, assumed to be independent and additive with zero mean and constant variance. This formulation is similar to Toms et al [24]. Enerlytic choses regression parameters that minimize the squared prediction error on the training set using the *SiZer* package [25] in R.

Enerlytic performs a parameter verification process, rejecting piecewise linear models which are unlikely to correspond to electric heating or cooling. For example, if the base temperature of a piecewise model is 0°C, it is unlikely this parameter corresponds to a set point. We reject the model if the parameters fall outside of the ranges tabulated in Table 3.2. μ represents the average power value of the energy signature. If the parameter check fails and the piecewise model is rejected, Enerlytic fits a linear model.

After fitting the regression models, Enerlytic creates the building operating parameters. This is a set of operating modes, each containing a peak and a base with average start and end times, and a regression model.

3.4 Prediction

The building operating parameters can be used for predicting a building's hourly power usage; these predictions can in turn be used for ESM estimation, discussed in Section

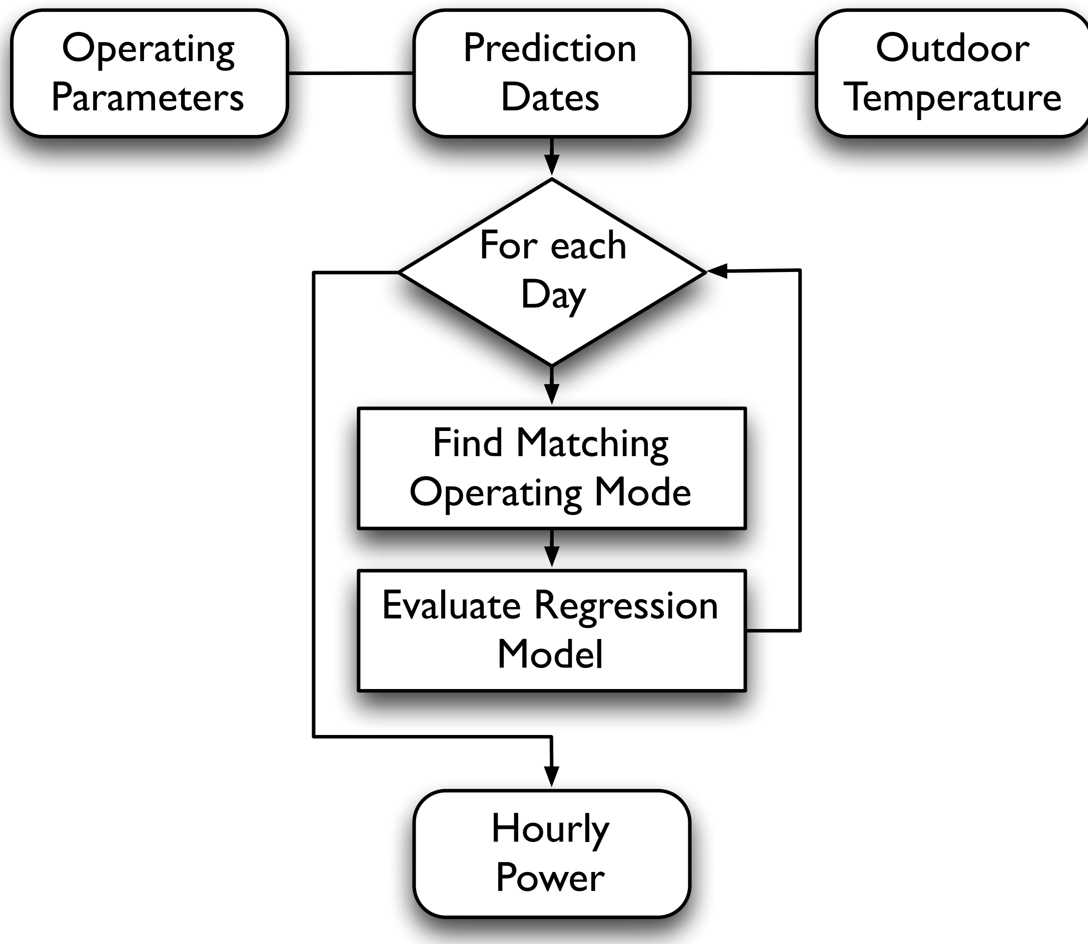

Figure 3.7: Prediction overview.

3.5, for validating the building operating parameters themselves, for forecasting electrical power demand, or creating baseline models to evaluate the effectiveness of an ESM. An overview of the prediction process can be seen in Figure 3.7.

The prediction mechanism takes as input:

1. The building operating parameters
2. A list of dates for which we would like to generate predictions, referred to as the *prediction dates*
3. The hourly temperatures for the prediction dates

The prediction algorithm generates a prediction for each prediction date, and later concatenates the results. It uses the labeller to determine the operating mode of each prediction date, and retrieves the operating parameters corresponding to that operating mode. The peak and base regression models are then evaluated using the temperature input. Once the prediction algorithm has predicted the peak and base power value for each day, it re-creates each daily load profile by filling the peak and base periods with the predicted peak and base values, respectively, and linearly interpolating during the ramping periods.

3.5 ESM Estimation

The building operating parameters can be used to estimate energy savings; this is done by changing the building operating parameters. The energy savings estimates do not correspond to savings from a particular action or ESM. Instead, the savings correspond to a category of possible ESMs, which we call *ESM scenarios*. This provides a building operator or an energy expert with the ability to estimate the potential energy savings due to a change in energy consumption patterns. Relating the changes in the building operating parameters to the building's equipment and operation requires a deep knowledge of how a specific building is run. The building operator or energy expert must generate a list of building-specific actions or ESMs which will accompany the ESM scenario. Estimates are obtained by substracting the predicted usage obtained using learned model parameters from the predicted usage under the modified model parameters. We demonstrate five ESM scenarios:

1. Peak or base average power reduction

2. Peak or base period reduction
3. Change of base temperature
4. Change of cooling or heating gradient
5. Change of operating mode

Figures 3.8 and 3.9 illustrate ESM scenarios 1-4, assuming a fixed cost of electricity of \$0.05 per kWh. The first two ESM scenarios result from changing the parameters of the daily load profile. For each predicted daily profile, we reduce the average peak by 10%, or decreasing the length of the peak period by one hour. An overview of the ESM estimation process can be seen in Figure 3.10.

Scenarios three and four arise from changing the parameters of the regression models. Here, we adjust any piecewise linear regressions according to the ESM scenario: for example increasing, the cooling base temperature by 1°C, or by reducing the cooling gradient by 10%. The final what-if scenario arises from creating a new labeller and changing the predicted operating mode of the building. For example, assuming we have data to support the ESM scenario, we could estimate the difference in energy consumption if: *a*) a retrofit had not taken place, for example to verify energy savings; *b*) an anomalous event had not occurred, such as equipment malfunctioning or being left on; or *c*) a low-power period, such as a holiday period, was extended.

3.6 Concluding Remarks

In this chapter we developed an algorithm, Enerlytic, which takes as input hourly power consumption, temperature data, and an optional labeller containing information about the building's operating modes, and creates a set of building operating parameters. We demonstrated how a building's operating parameters can be used for prediction and to estimate energy savings. We believe there are many ways to extend this model, and we continue this discussion in Chapter 5, after presenting the model validation in Chapter 4.

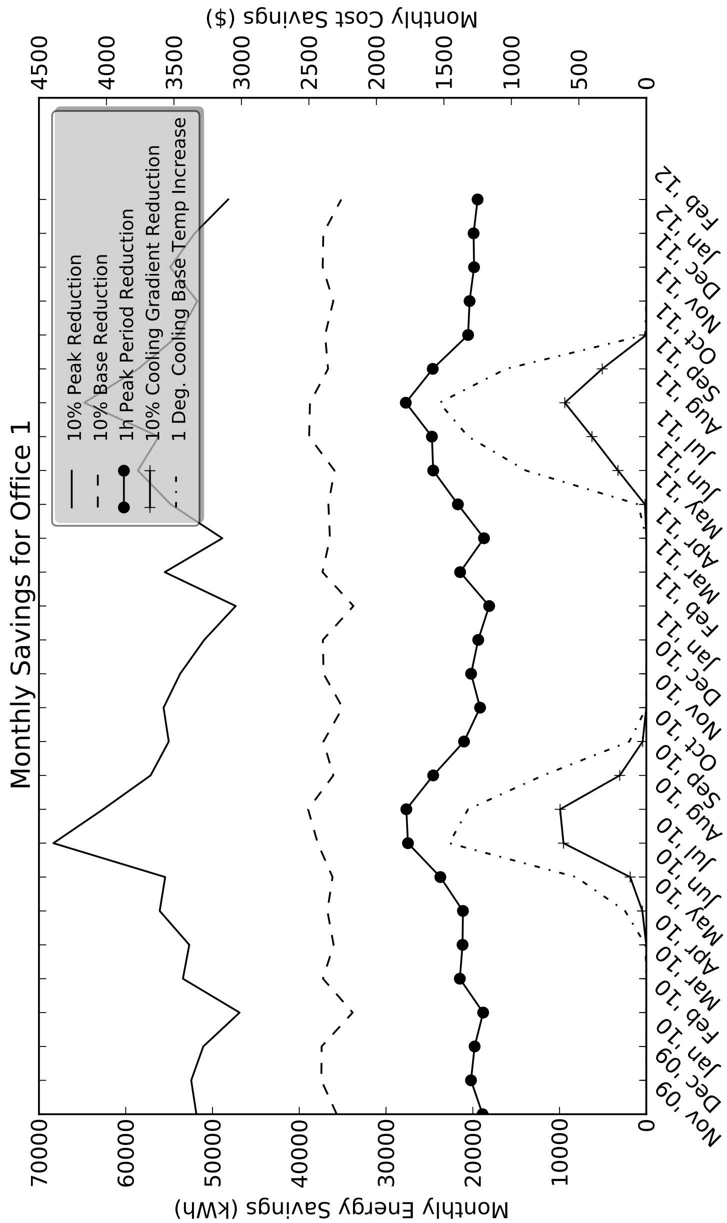

Figure 3.8: ESM estimation for an office building.

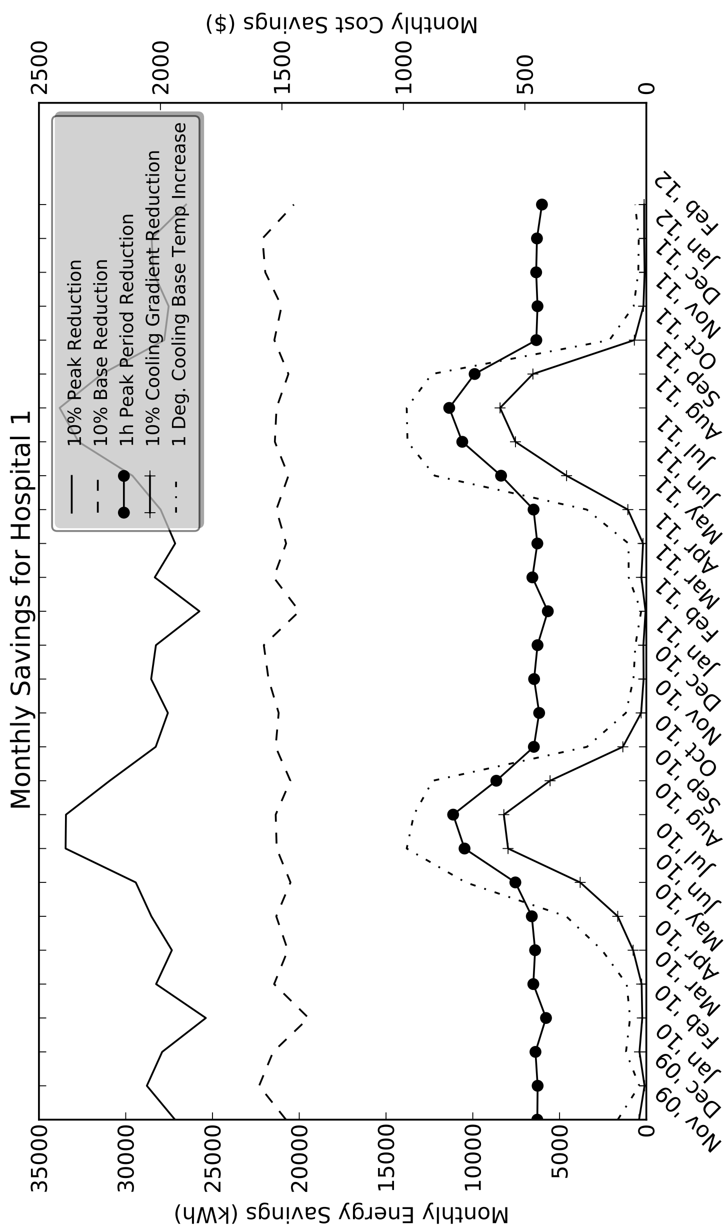

Figure 3.9: ESM estimation for a hospital.

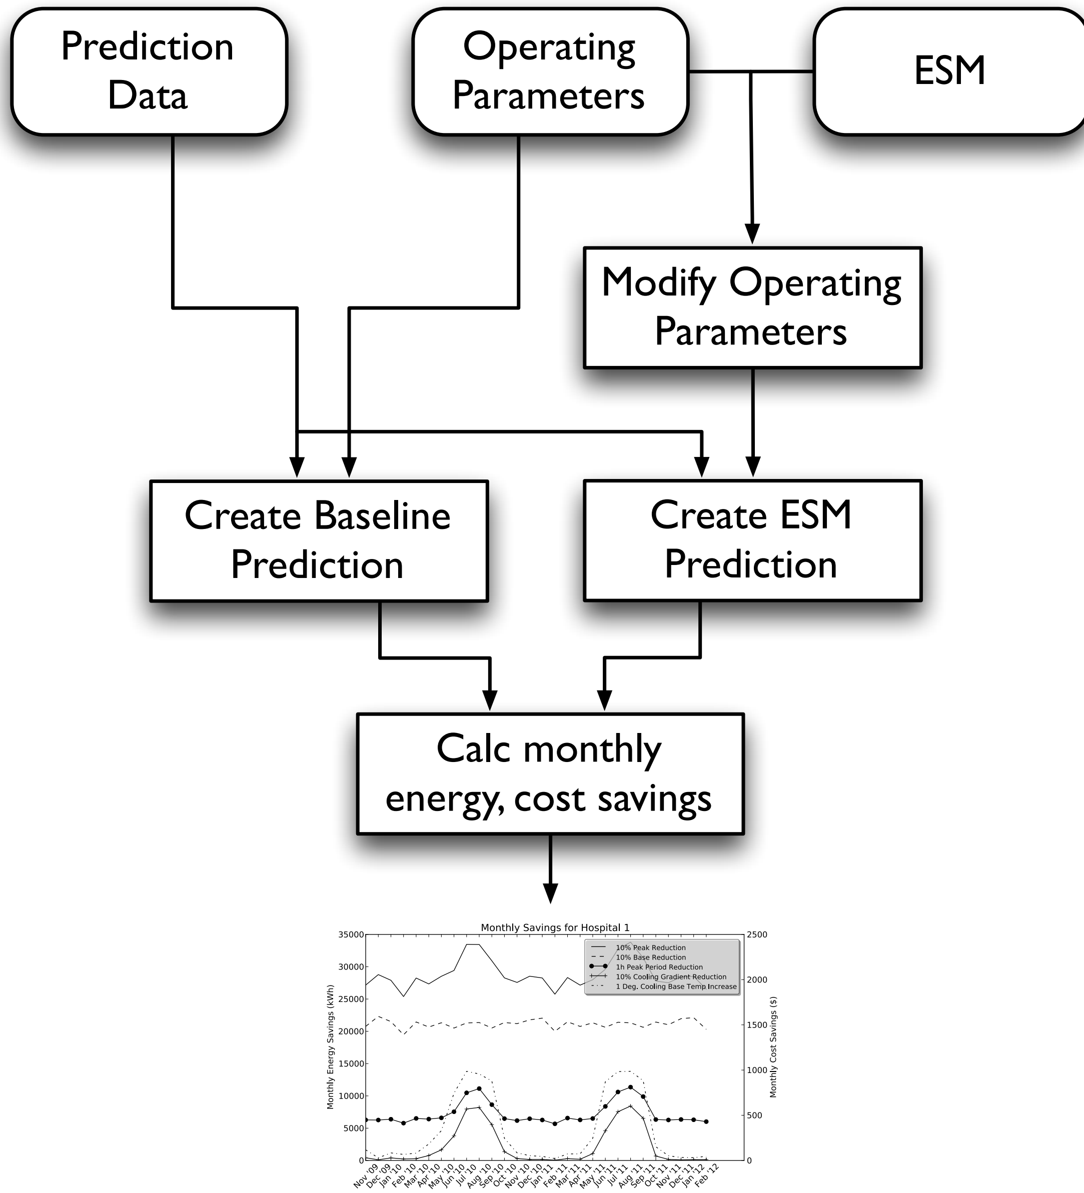

Figure 3.10: ESM estimation overview.

Chapter 4

Evaluation

We evaluate the building operating parameters by testing the prediction algorithm on a dataset provided by Pulse Energy, a leading Energy Information System (EIS) provider [26]. The dataset contains:

- Two schools, two hospitals, one grocery store and five office buildings
- Buildings in Western and Eastern Canada, and North-Western and South-Eastern United States.
- An average of over two years of data per building for a total of over 21 years of meter data
- Hourly average power consumption and accompanying hourly outdoor temperature data

Table A.1 contains a more detailed description of the dataset. We use buildings that span a variety of commercial areas and geographies and use several years of data for each building to evaluate Enerlytic under a range of seasons, climatic zones, and electrical demand patterns.

We measure the accuracy of our predictions using K=4 fold cross validation on each building. The partitions for each fold are chosen at random. The default labeller is used to label and partition the data for each building into weekdays and weekends. We measure prediction accuracy using the coefficient-of-variation of the root-mean-squared-error (CV-RMSE) and mean-bias-error (MBE); these are standard metrics in the building energy prediction literature [14] [9] [13]. We benchmark our prediction algorithm against a state-of-the-art prediction algorithm that uses kernel regression [9].

4.1 Results

Figures 4.1 and 4.2 show, respectively, a summary of the CV-RMSE and MBE performance of our method ("Enerlytic") and the benchmark ("Kernel Regression"). The 90% confidence intervals represent the contributions from each of the four folds.

Enerlytic is less accurate at prediction than the benchmark; this is to be expected due to the relative simplicity and restrictive parameterization of the building operating parameters, and because the primary goal of the building operating parameters is to gain insight and be used as input to other systems, not solely to create an accurate prediction algorithm. We believe the prediction error is comparable enough to current practice to enable the building operating parameters to be used for prediction, and also as input to other systems such as ESM estimation and benchmarking.

The results shown in Figures 4.1 and 4.2 fall into two categories: buildings where the default labeller identifies the building's operating modes, and buildings where the default labeller misses significant operating modes of a building. When a building primarily has weekend and weekday operating modes (Hospitals 1, 2; Offices 1, 2, 4; Store), the prediction performance of their building operating parameters is quite similar to the benchmark. An example of this can be seen in the detailed results of Office 1, shown in Figures 4.3-4.6. Figure 4.3 shows the actual and predicted demand during the prediction period, Figure 4.4 shows the prediction error (residual) and the data used to create (train) the building operating parameters. Figures 4.5 and 4.6 show the energy signatures for the weekday and weekend operating modes, respectively. Statutory holidays are a primary source of error, as seen in Figure 4.4(b) where holidays appear as spikes in the residual, and in Figure 4.5(a), where holidays are classified as outliers. Creating a labeller that labels statuatory holidays may further improve prediction performance.

The remaining buildings (Schools 1, 2; Offices 3, 5) have significant operating modes that were not identified by the default labeller. The schools had summer holiday periods that caused poor regression fits due to Simpson's paradox [22]; an example of this can be seen in Figures 4.7-4.10. Outlier detection was not effective. An improved outlier detection strategy may improve prediction performance, but we suggest future work investigate default labellers that create labels based on the underlying statistics of the data. Figures 4.11-4.14 show an example of the prediction performance for Office 5. Figure 4.13(a) shows there are shifts in the consumption patterns of the building's electric heating system; there are small translations between the "training" and "actual" data, causing a systemic bias in prediction error. These changes in the building's electric heating system can be interpreted as a change in operating mode. This could be captured by labelling the building

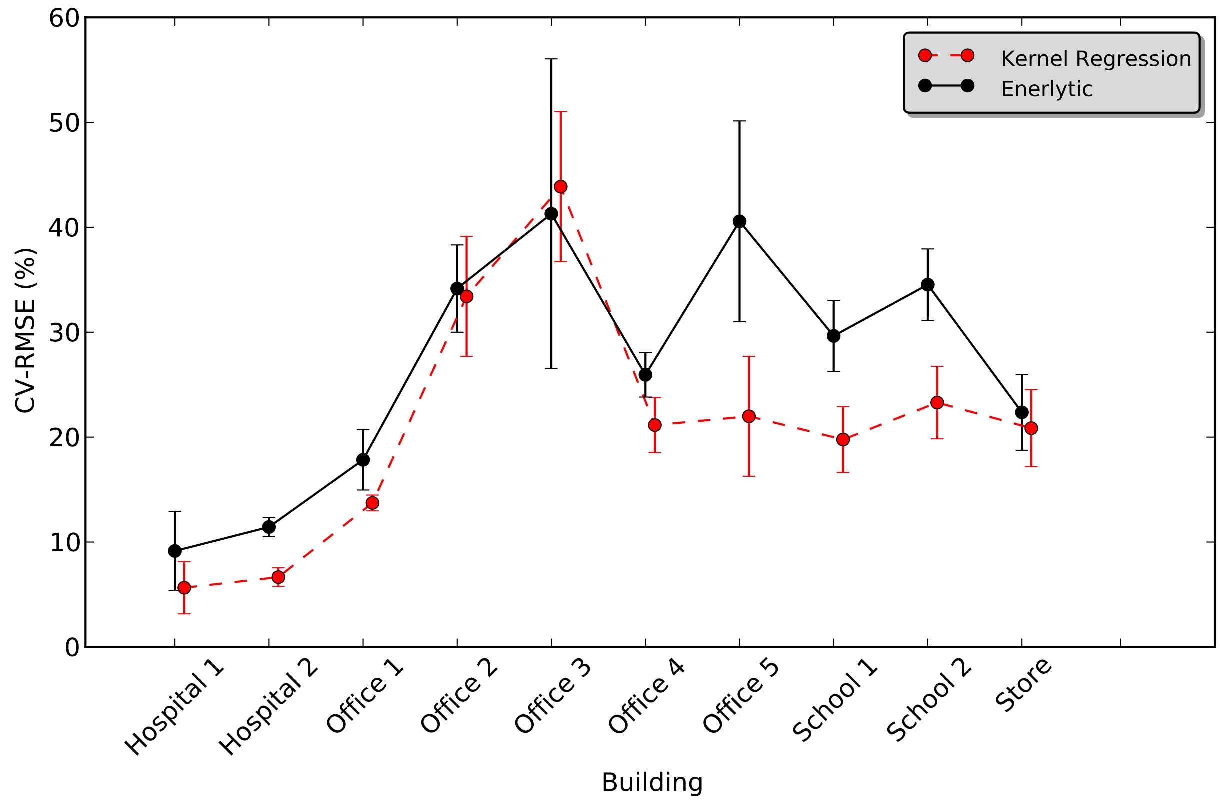

Figure 4.1: CV-RMSE results for all buildings.

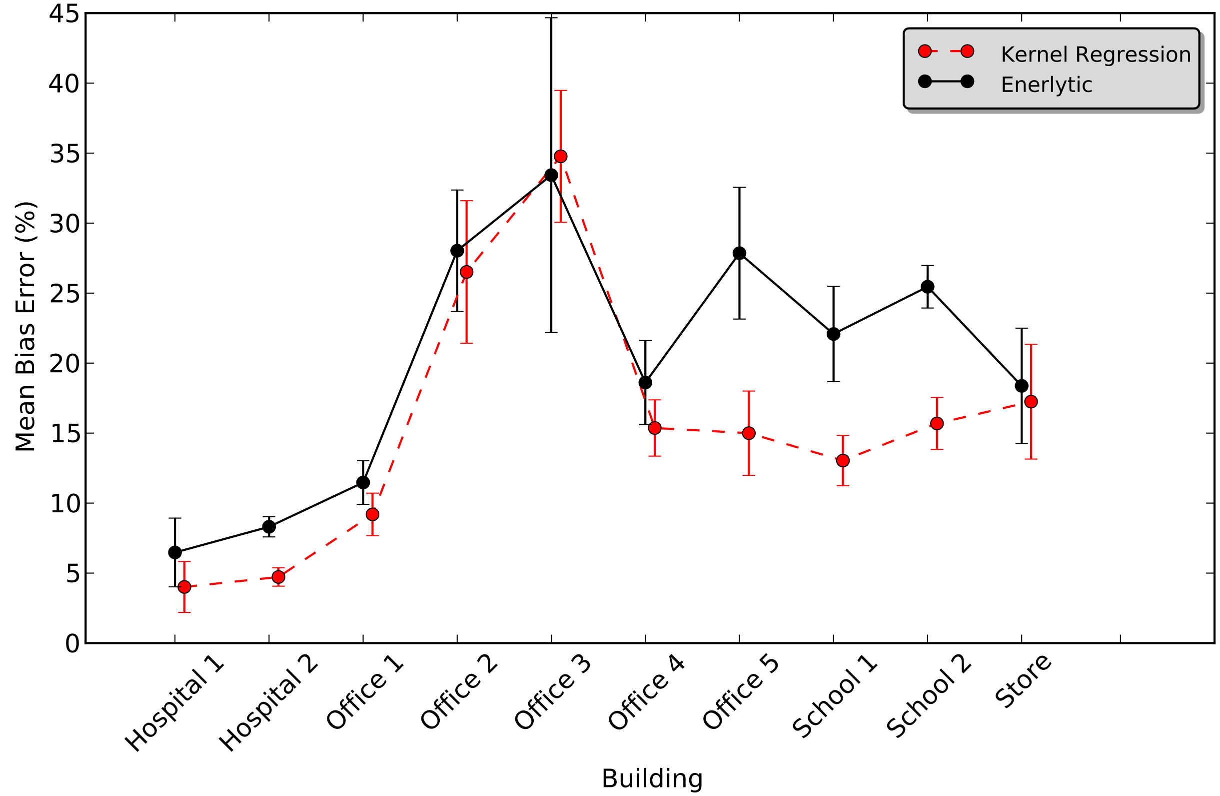

Figure 4.2: Mean Bias Error results for all buildings.

as being in different operating modes during these two periods. This result also shows the building operating parameters generalize poorly compared to the benchmark. Another example of poor generalization (overfitting) can been seen in Figure 4.14(a), where the building operating parameters do not anticipate the use of electric heating on weekends. Future work may improve the model's ability to generalize by considering the underlying distribution of the data: if the variance is sufficiently high, a stochastic process could be used for prediction instead of linear regressions.

Office 3 had gradual changes in the load profile throughout its measurement period; this can be seen by comparing Figure 4.15(a), which shows the data in the first fold, to Figure 4.16(a), which shows the data in the final three folds. Figures 4.15-4.18 show detailed prediction performance for the first fold of Office 3. Figure 4.15(a) shows Office 3 initially had relatively high electric energy consumption; Figure 4.16(a) indicates this consumption level fell and remained relatively low during the rest of the measurement period. When the data from Figure 4.16(a) was used to predict the period shown in Figure 4.15(a), the prediction error was high since the training data was not representative of the demand during the prediction period—the building was in a different operating mode. When predicting folds 2, 3, and 4, the high consumption period fell in the training period instead of the prediction period; the outlier detection mechanism typically labelled the high consumption days as outliers. Figures 4.19-4.22 show the prediction performance of the second fold of Office 3. As seen in Figures 4.21 and 4.22, data in the high consumption period were labelled as outliers. The transient period of high consumption, combined with the outlier detection, led to relatively low prediction error for folds 2, 3, and 4, while the error from fold 1 remained high. This caused a large variation in prediction performance for this building, despite the mean prediction performance being comparable to the benchmark. Like Office 5, prediction performance may be improved by partitioning the dataset into different operating modes using a more advanced labelling scheme, or by using a regression model that is less prone to overfitting.

In summary, Enerlytic predicts well when the operating modes in the training data and prediction period have been identified. The non-linear kernel regression generalizes more effectively than the piecewise linear models, particularly when operating modes are not identified by the labeller. Future work could address these limitations by creating more advanced labellers and by investigating regression methods that are able to both provide insight and also avoid overfitting.

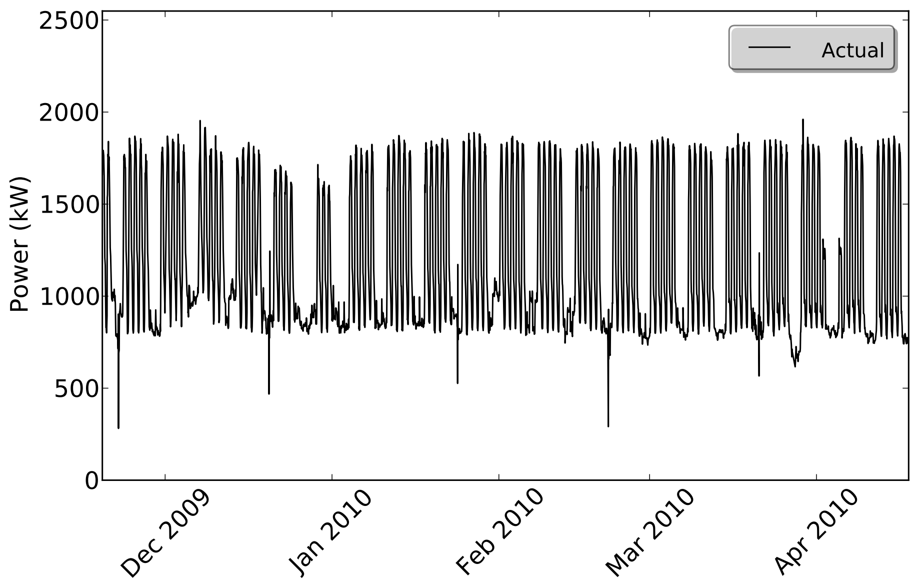

(a) Office 1: Actual Demand

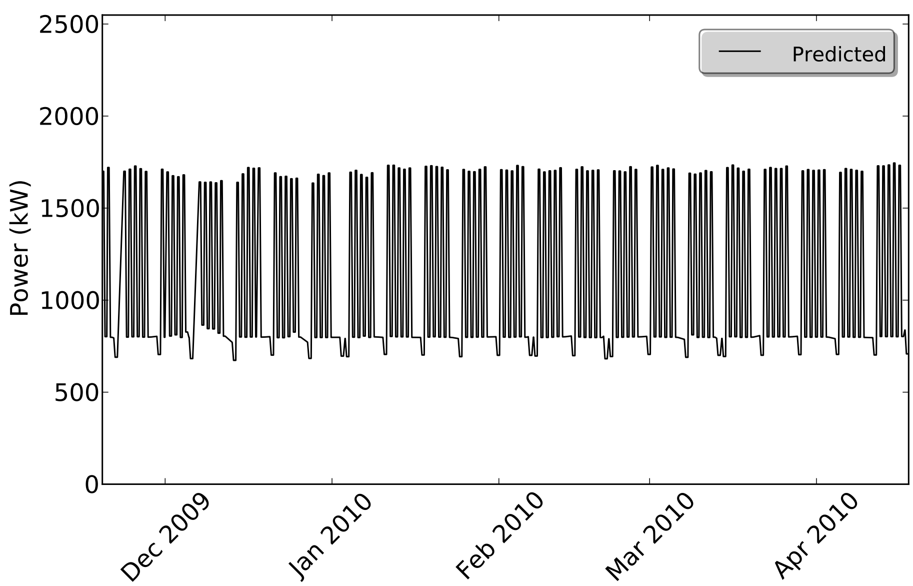

(b) Office 1: Predicted Demand

Figure 4.3: Office 1 actual and predicted demand during predicted period (fold 1).

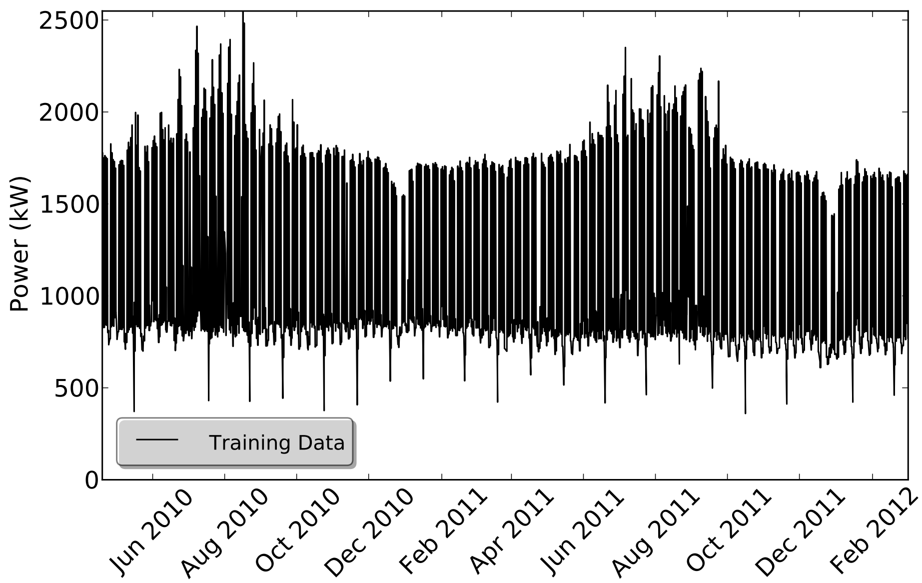

(a) Office 1: Training Demand

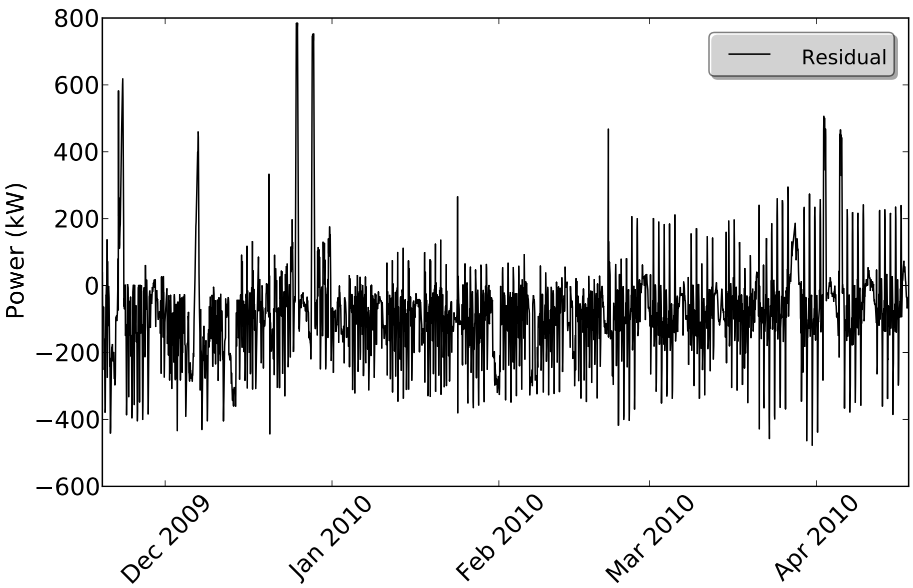

(b) Office 1: Residual Demand

Figure 4.4: Office 1 training and residual demand during predicted period (fold 1).

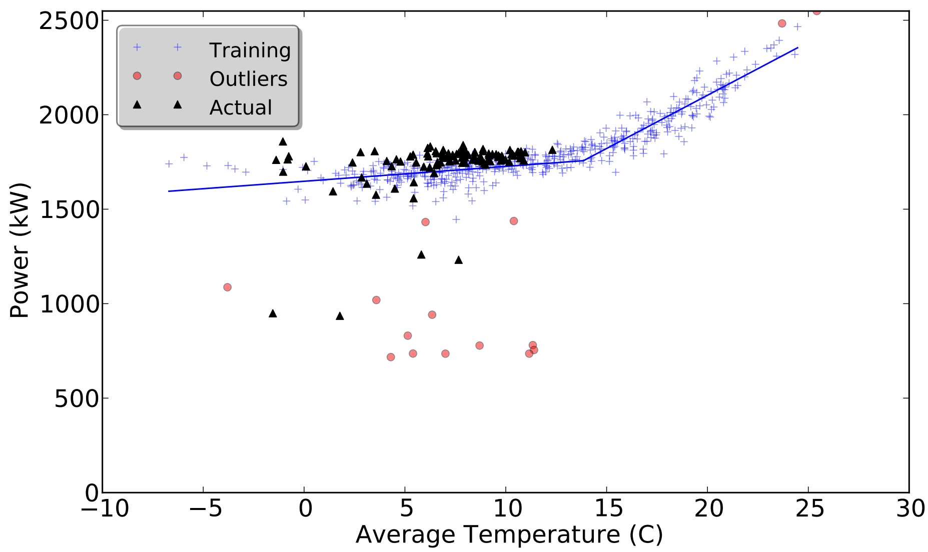

(a) Office 1: Weekday Peak Energy Signature

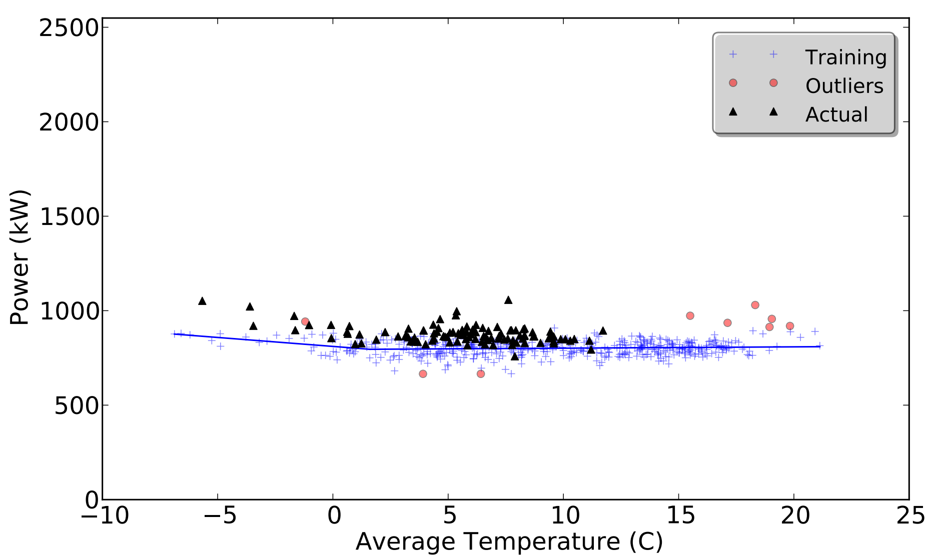

(b) Office 1: Weekday Base Energy Signature

Figure 4.5: Office 1 weekday energy signatures (fold 1).

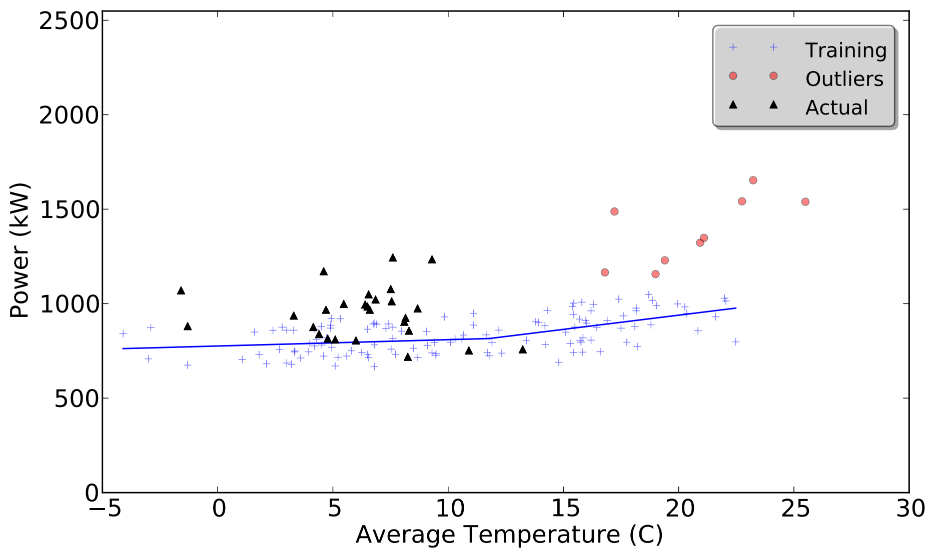

(a) Office 1: Weekend Peak Energy Signature

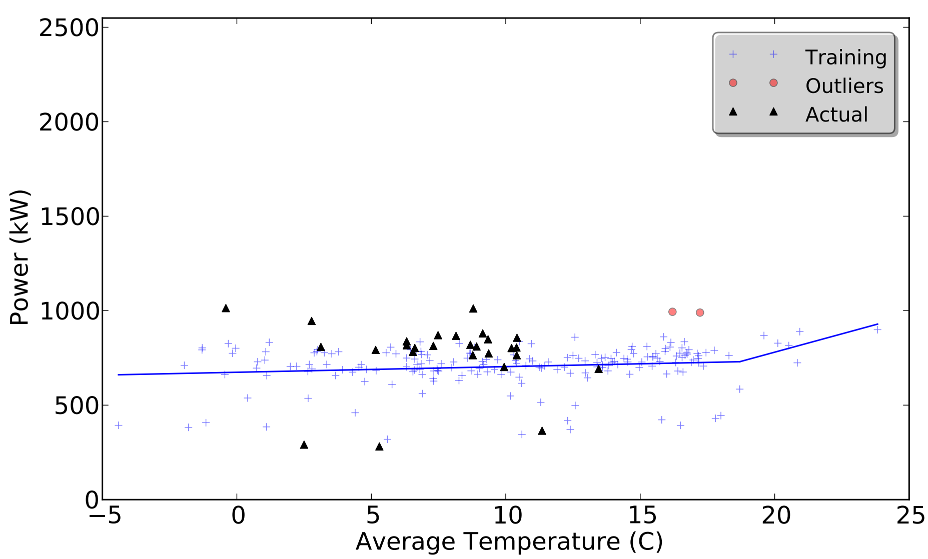

(b) Office 1: Weekend Base Energy Signature

Figure 4.6: Office 1 weekend energy signatures (fold 1).

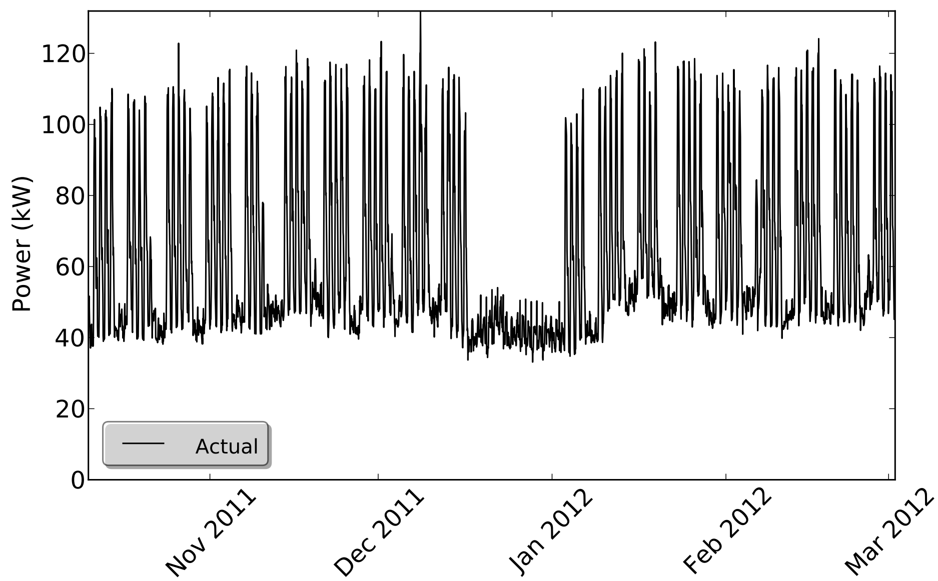

(a) School 1: Actual Demand

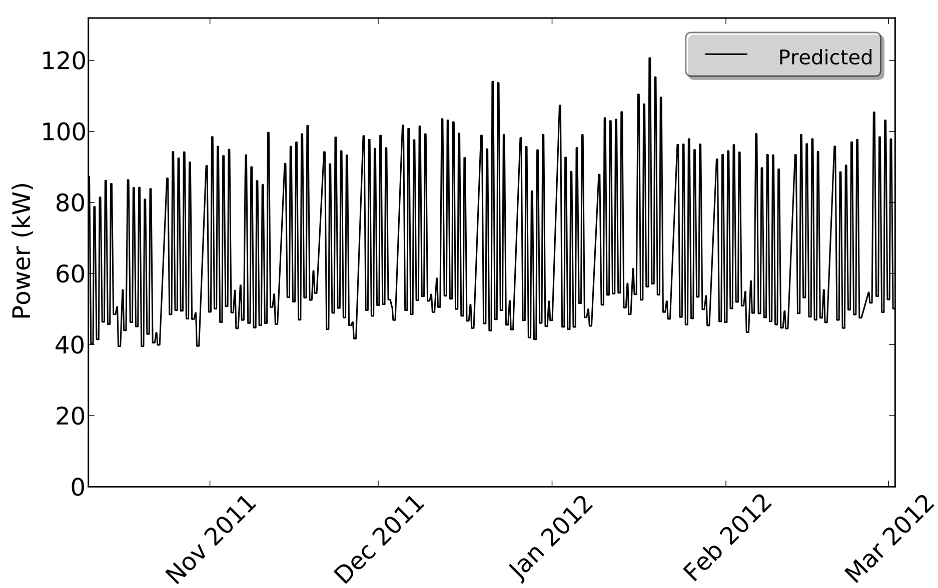

(b) School 1: Predicted Demand

Figure 4.7: School 1 actual and predicted demand during predicted period (fold 4).

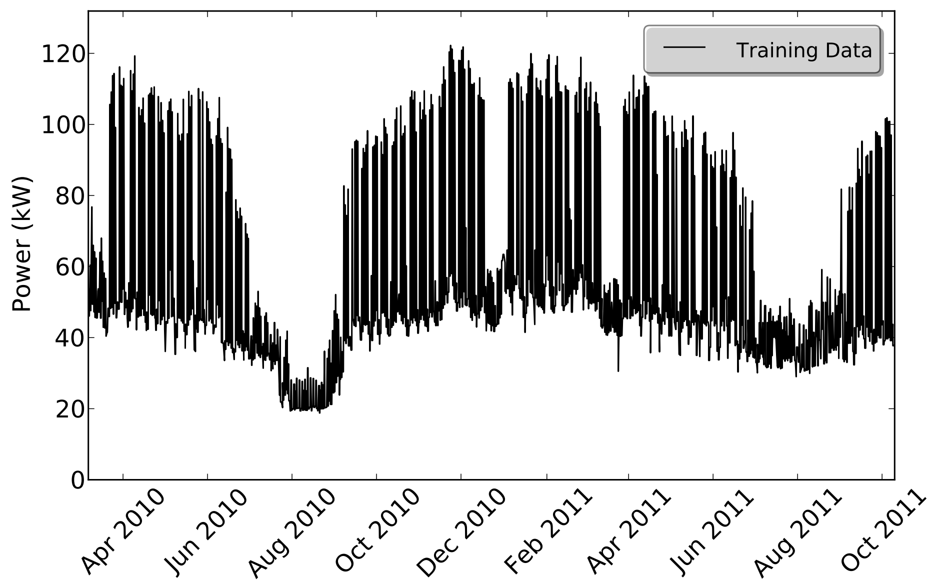

(a) School 1: Training Demand

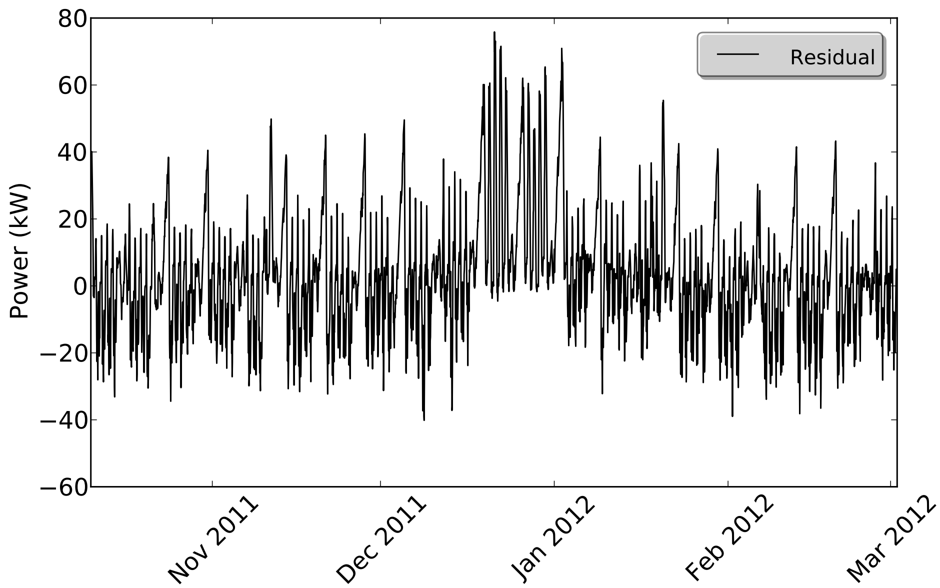

(b) School 1: Residual Demand

Figure 4.8: School 1 training and residual demand during predicted period (fold 4).

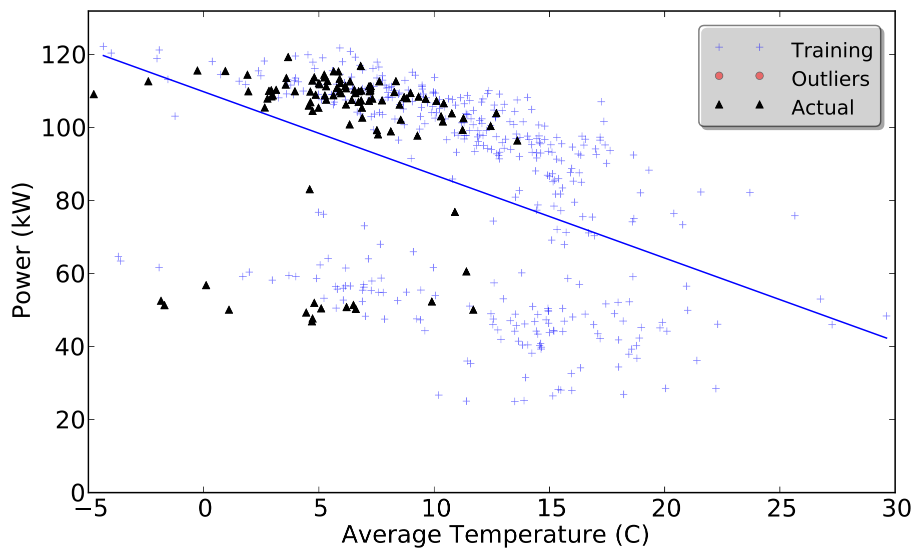

(a) School 1: Weekday Peak Energy Signature

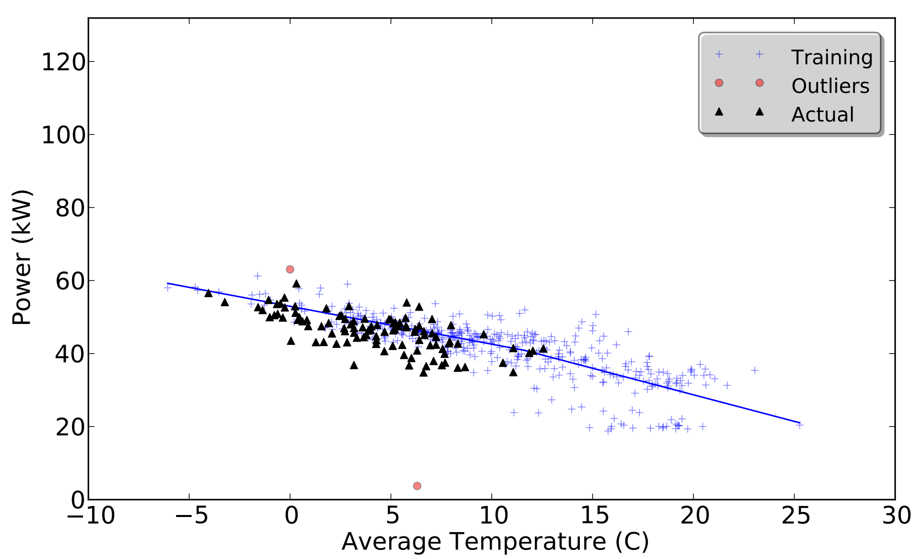

(b) School 1: Weekday Base Energy Signature

Figure 4.9: School 1 weekday energy signatures (fold 4).

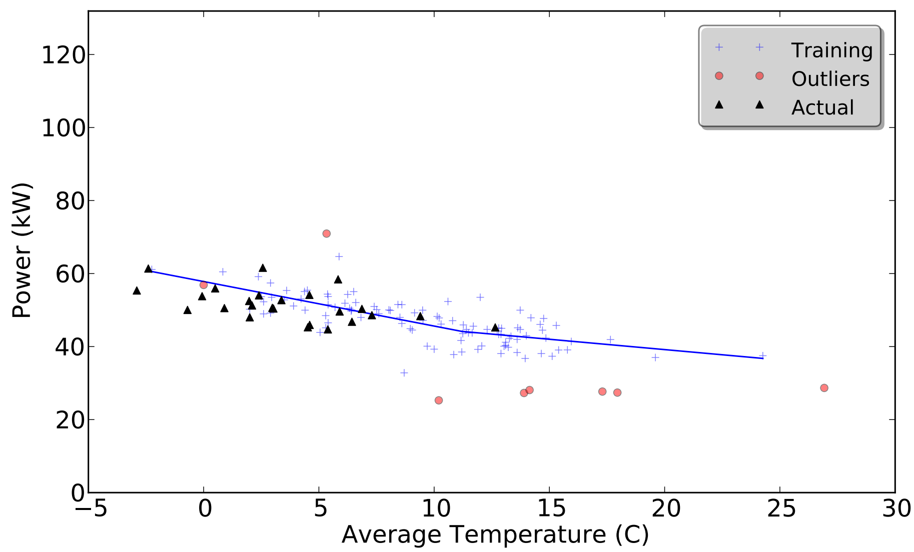

(a) School 1: Weekend Peak Energy Signature

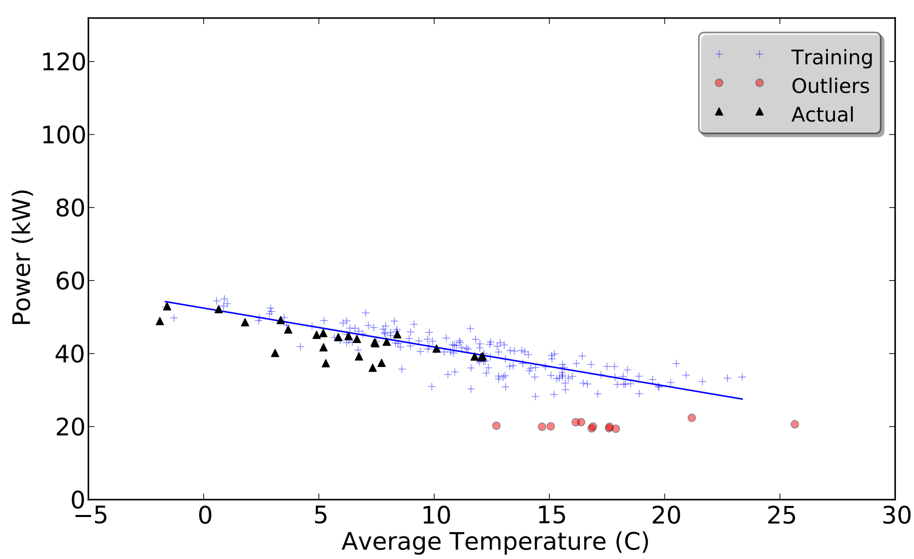

(b) School 1: Weekend Base Energy Signature

Figure 4.10: School 1 weekend energy signatures (fold 4).

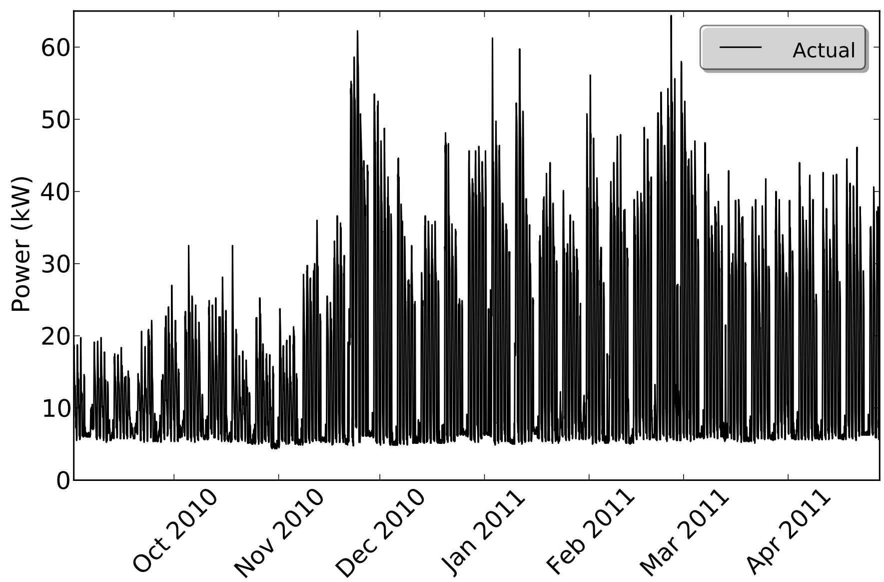

(a) Office 5: Actual Demand

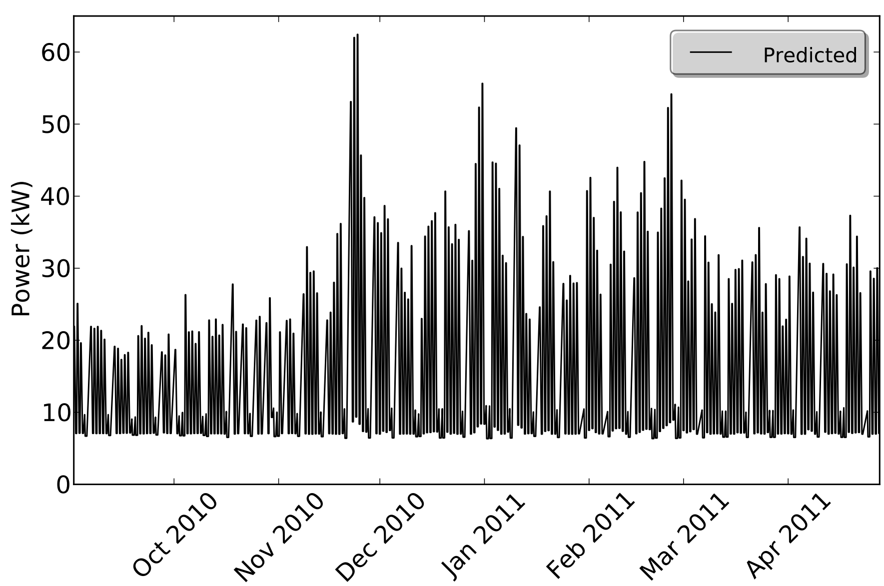

(b) Office 5: Predicted Demand

Figure 4.11: Office 5 actual and predicted demand during predicted period (fold 1).

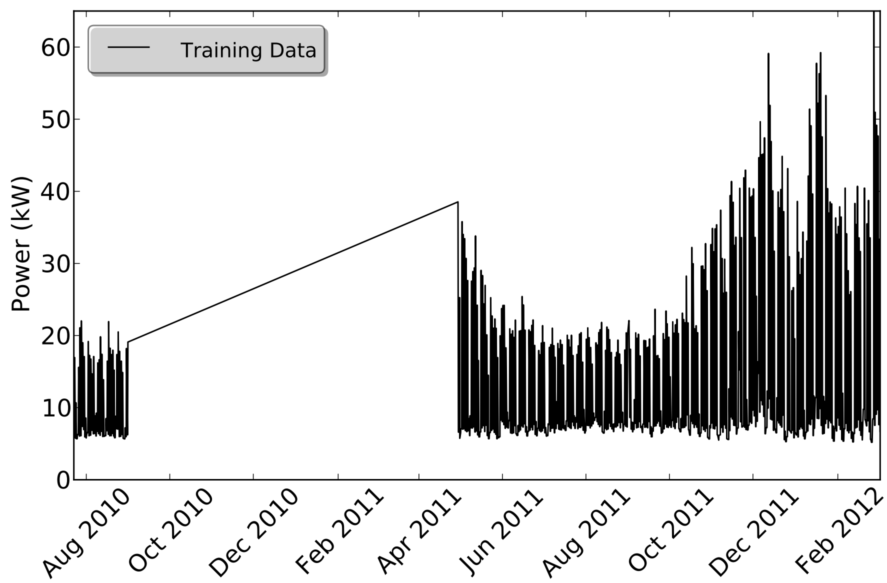

(a) Office 5: Training Demand

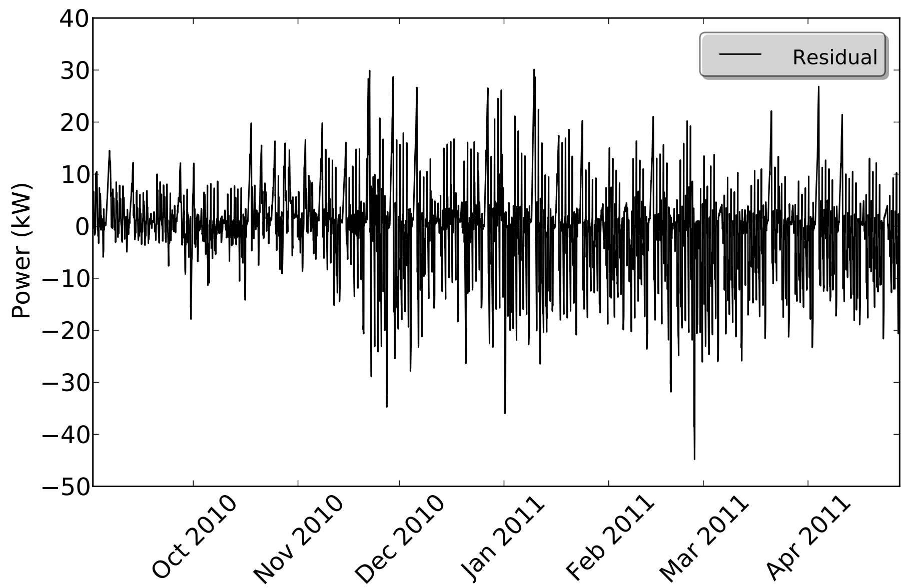

(b) Office 5: Residual Demand

Figure 4.12: Office 5 training and residual demand during predicted period (fold 1).

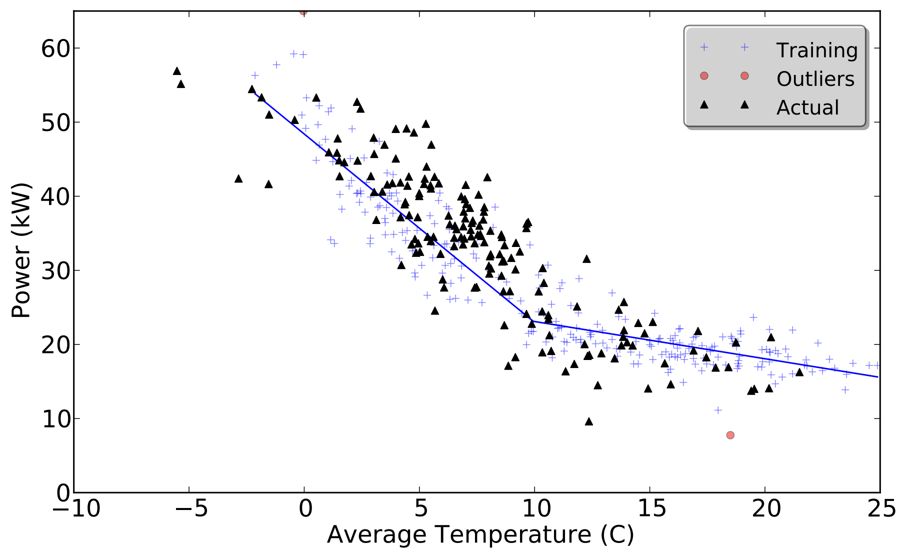

(a) Office 5: Weekday Peak Energy Signature

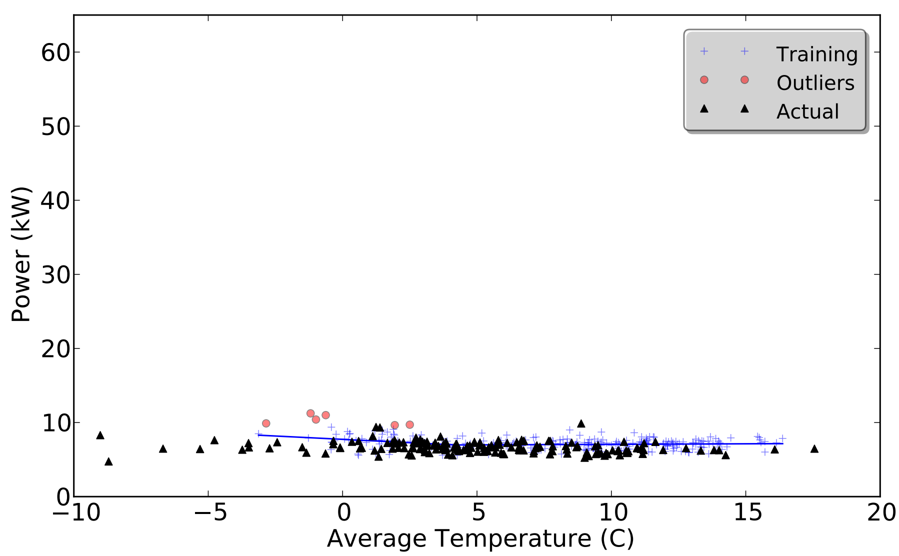

(b) Office 5: Weekday Base Energy Signature

Figure 4.13: Office 5 weekday energy signatures (fold 1).

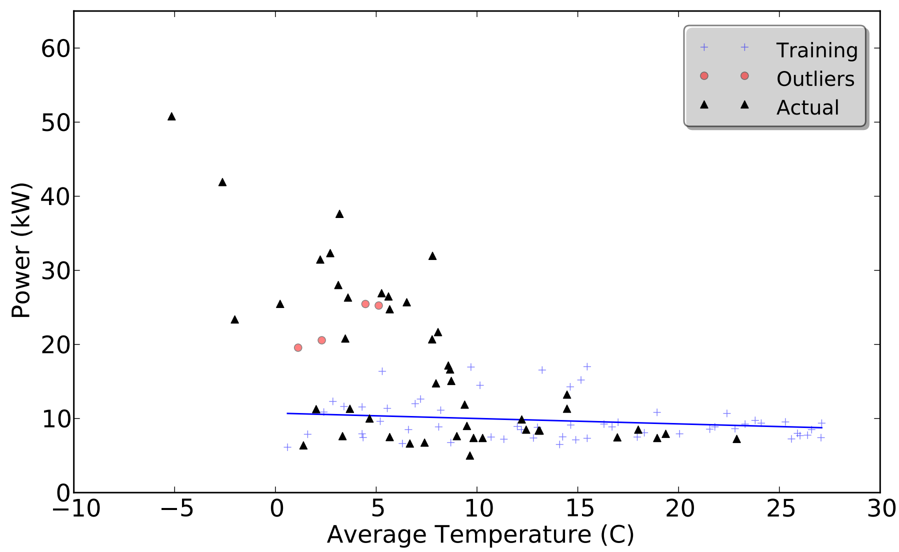

(a) Office 5: Weekend Peak Energy Signature

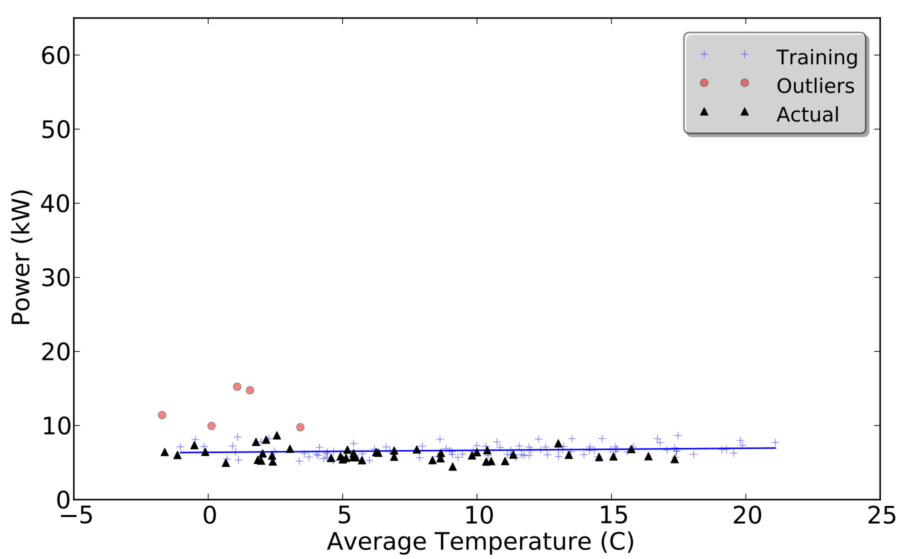

(b) Office 5: Weekend Base Energy Signature

Figure 4.14: Office 5 weekend energy signatures (fold 1).

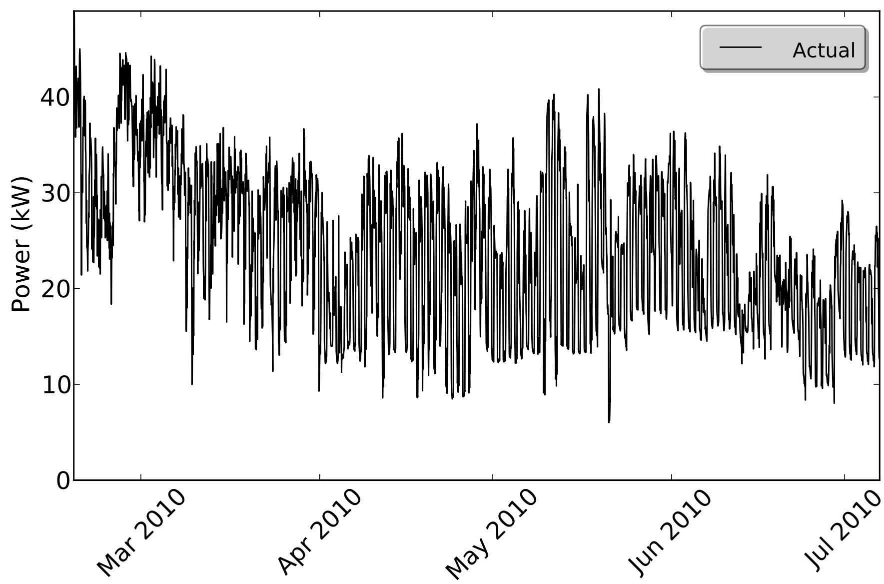

(a) Office 3: Actual Demand

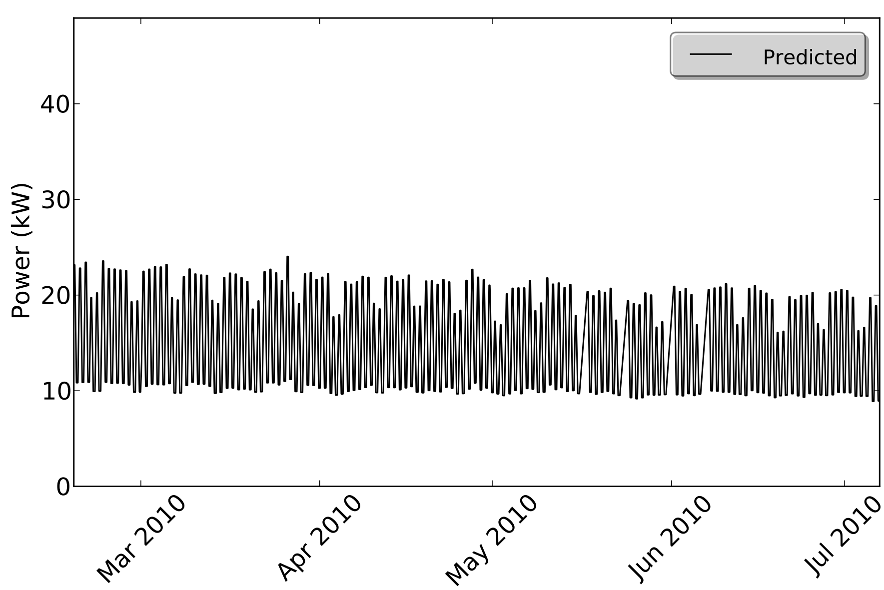

(b) Office 3: Predicted Demand

Figure 4.15: Office 3 actual and predicted demand during predicted period (fold 1).

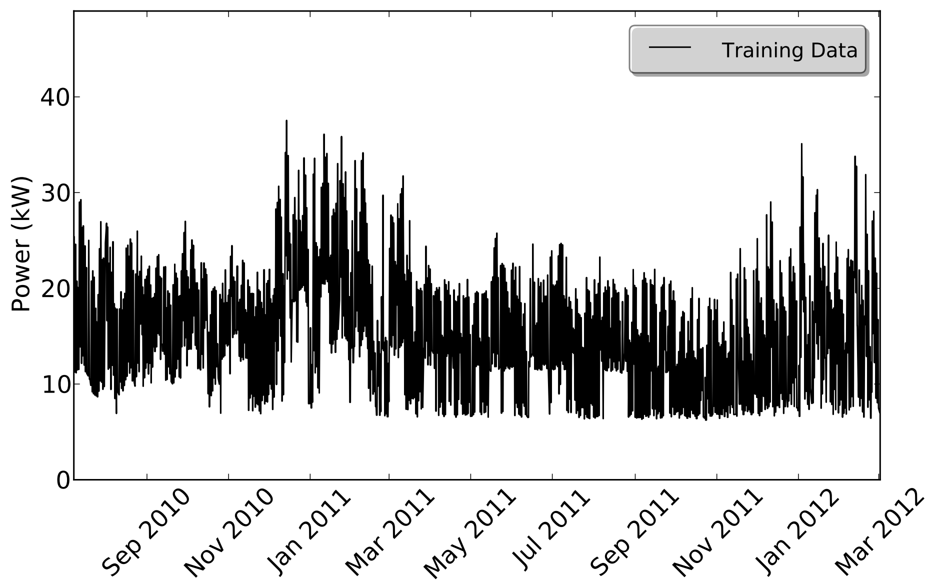

(a) Office 3: Training Demand

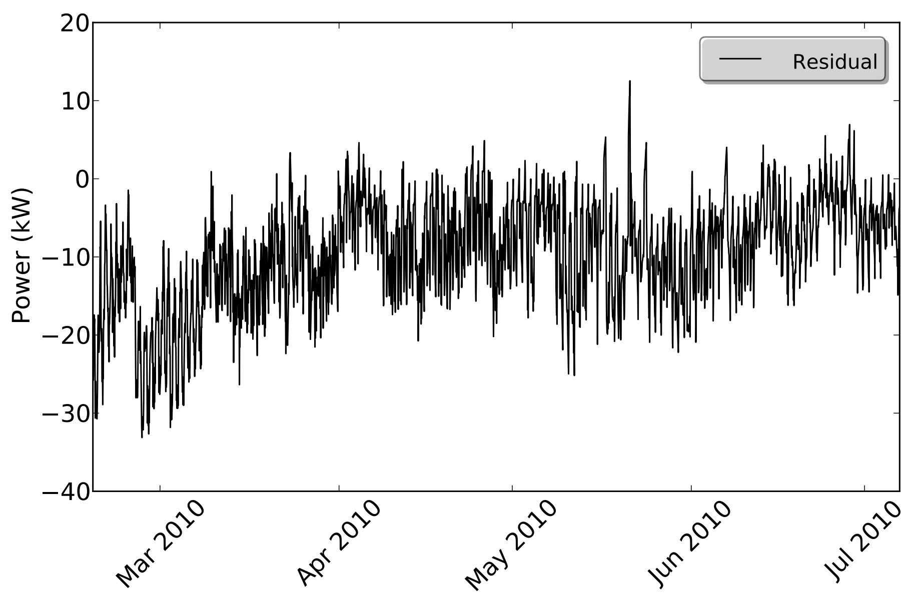

(b) Office 3: Residual Demand

Figure 4.16: Office 3 training and residual demand during predicted period (fold 1).

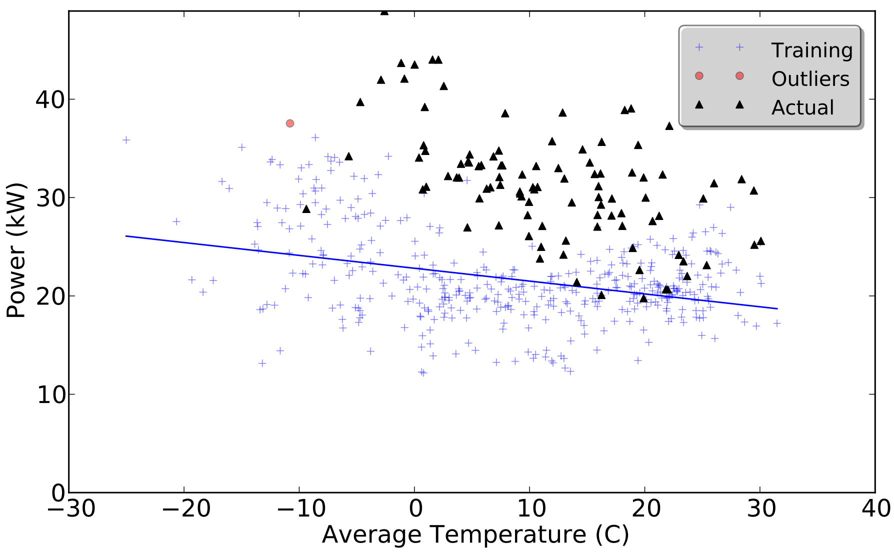

(a) Office 3: Weekday Peak Energy Signature

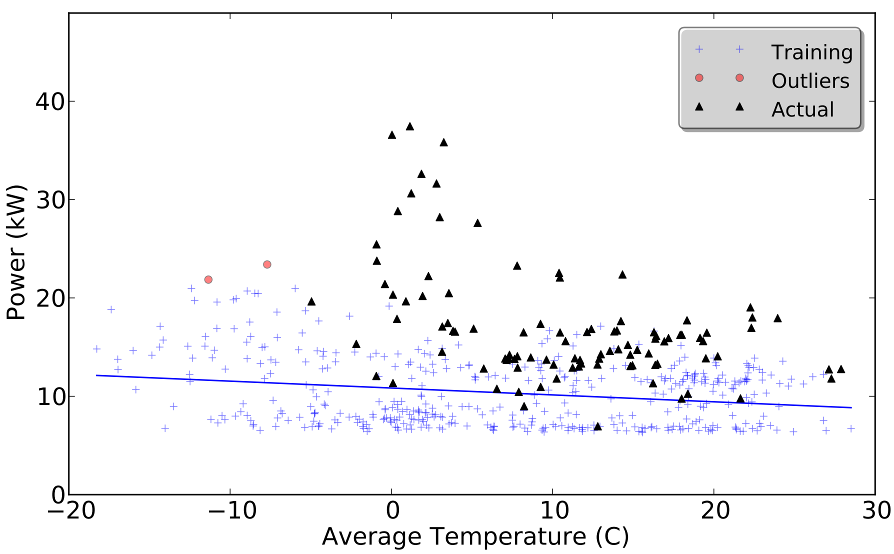

(b) Office 3: Weekday Base Energy Signature

Figure 4.17: Office 3 weekday energy signatures (fold 1).

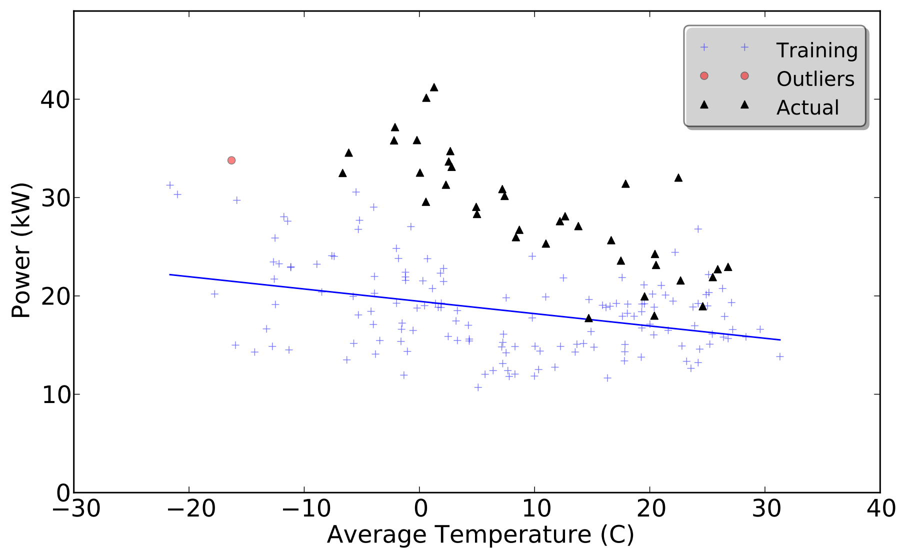

(a) Office 3: Weekend Peak Energy Signature

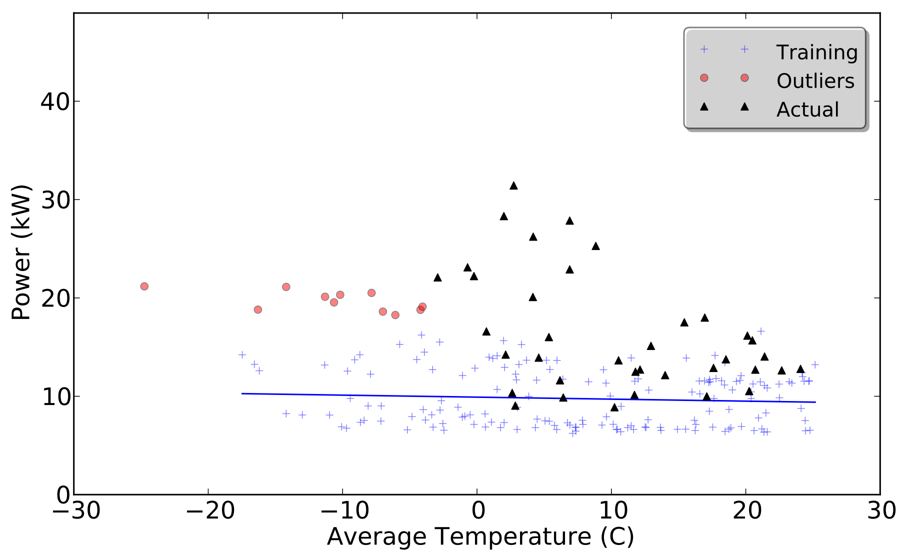

(b) Office 3: Weekend Base Energy Signature

Figure 4.18: Office 3 weekend energy signatures (fold 1).

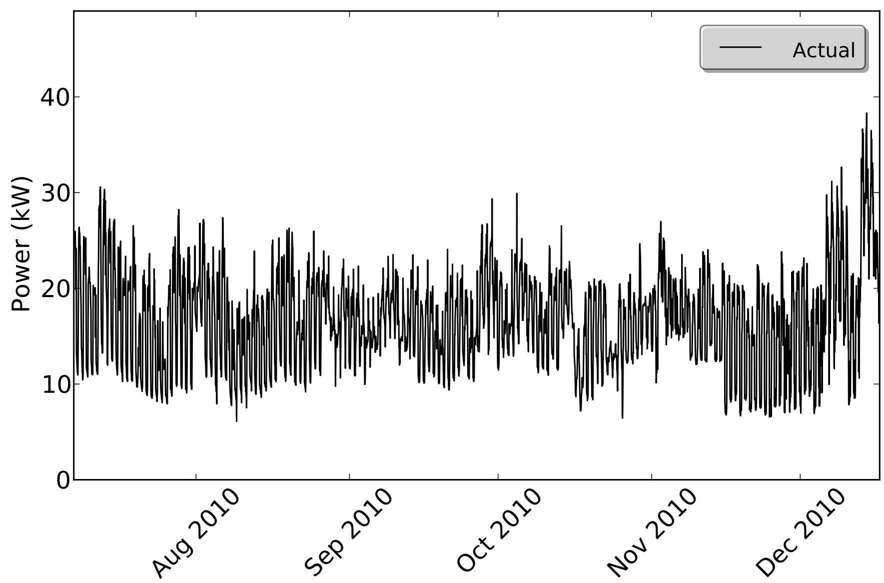

(a) Office 3: Actual Demand

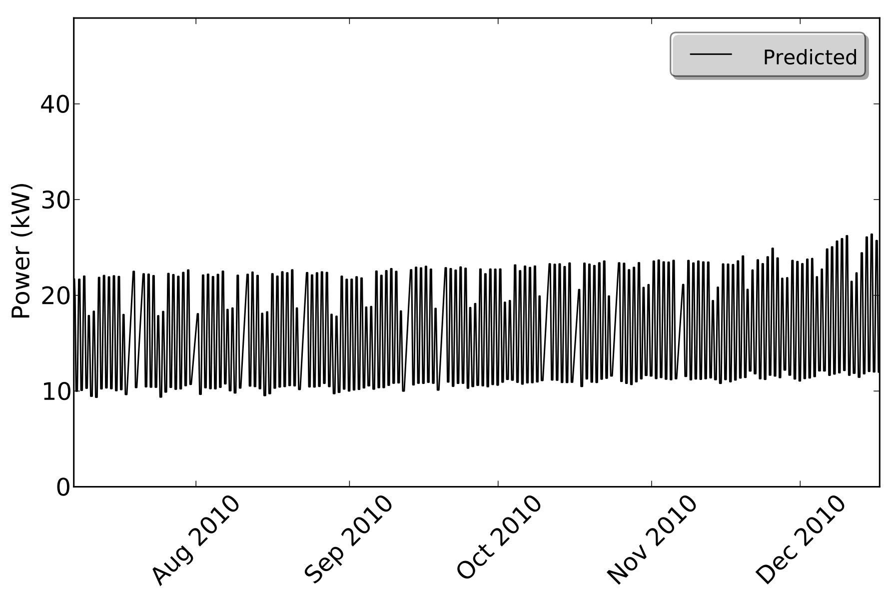

(b) Office 3: Predicted Demand

Figure 4.19: Office 3 actual and predicted demand during predicted period (fold 2).

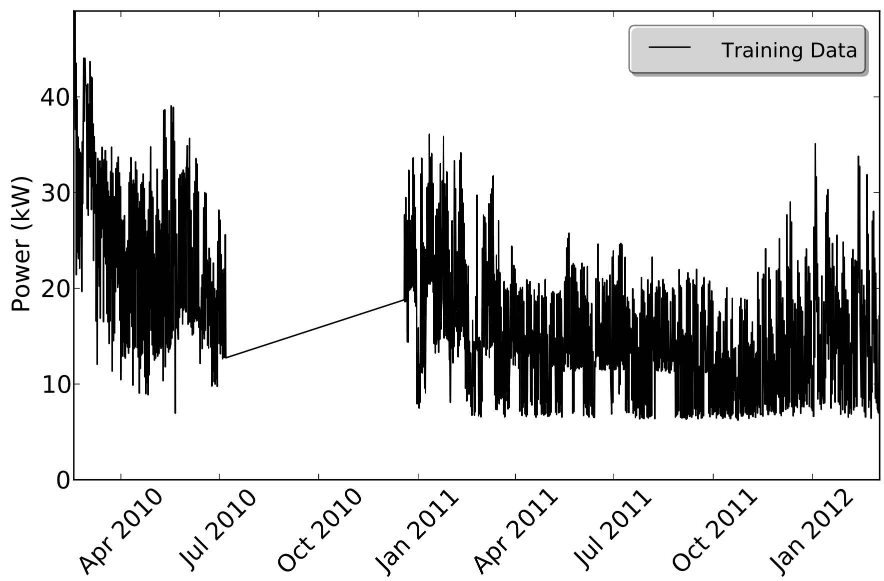

(a) Office 3: Training Demand

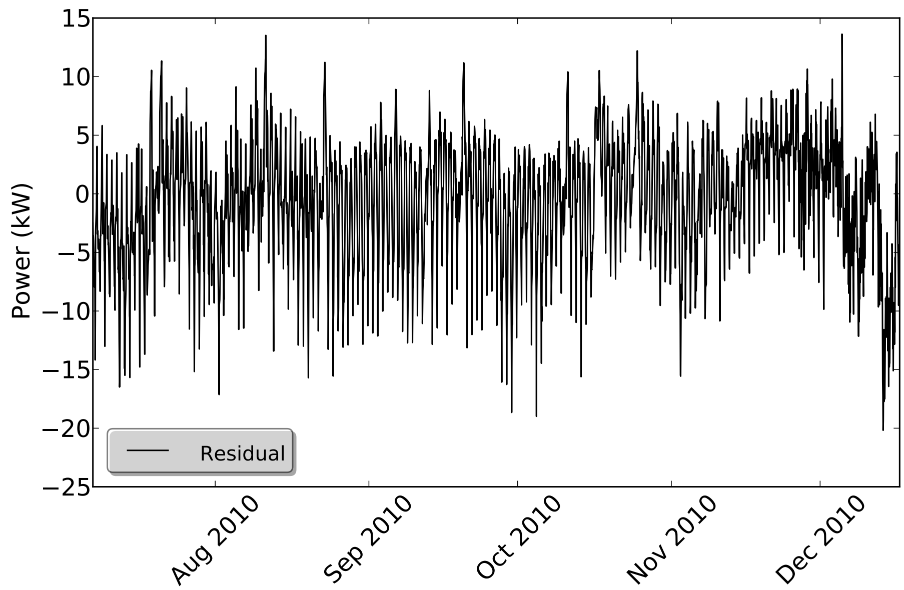

(b) Office 3: Residual Demand

Figure 4.20: Office 3 training and residual demand during predicted period (fold 2).

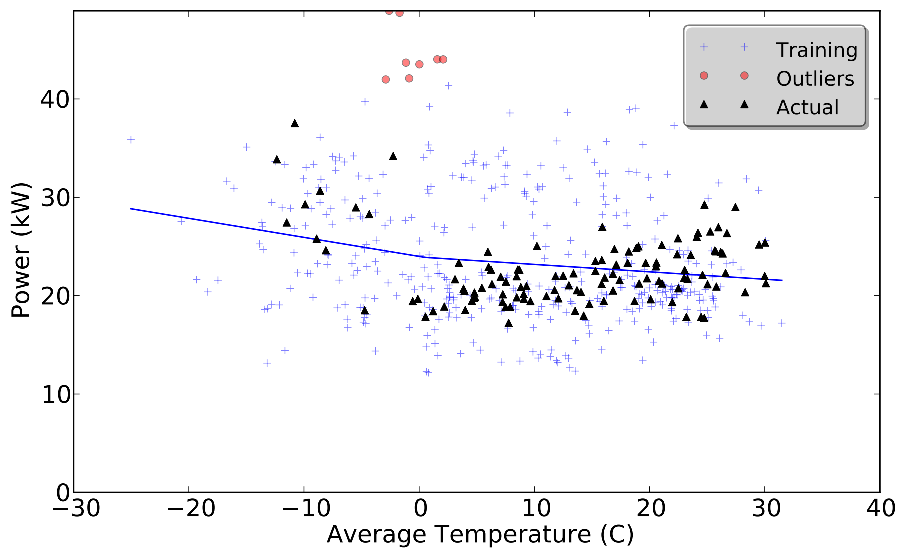

(a) Office 3: Weekday Peak Energy Signature

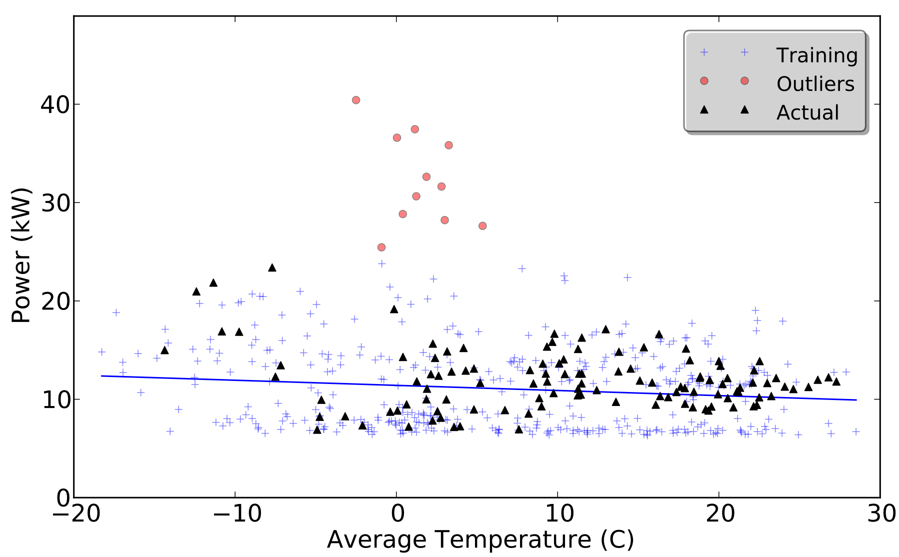

(b) Office 3: Weekday Base Energy Signature

Figure 4.21: Office 3 weekday energy signatures (fold 2).

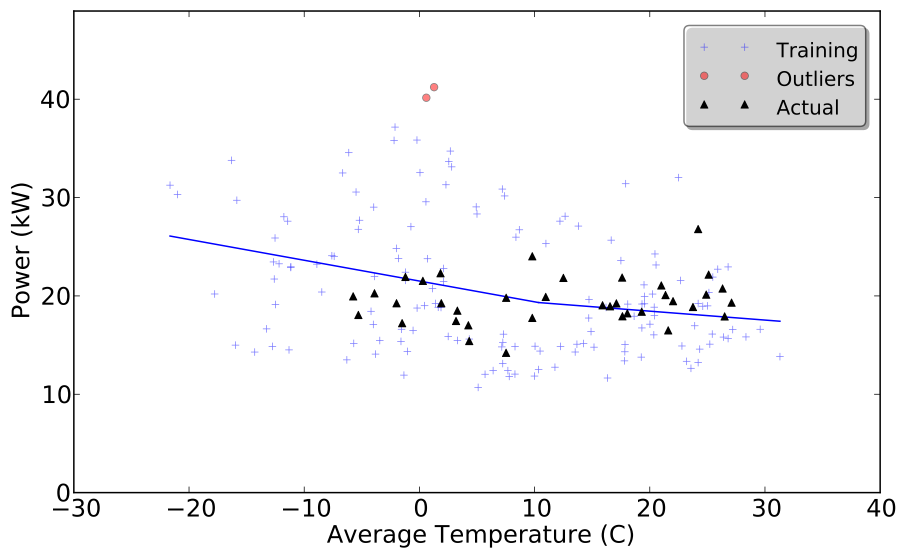

(a) Office 3: Weekend Peak Energy Signature

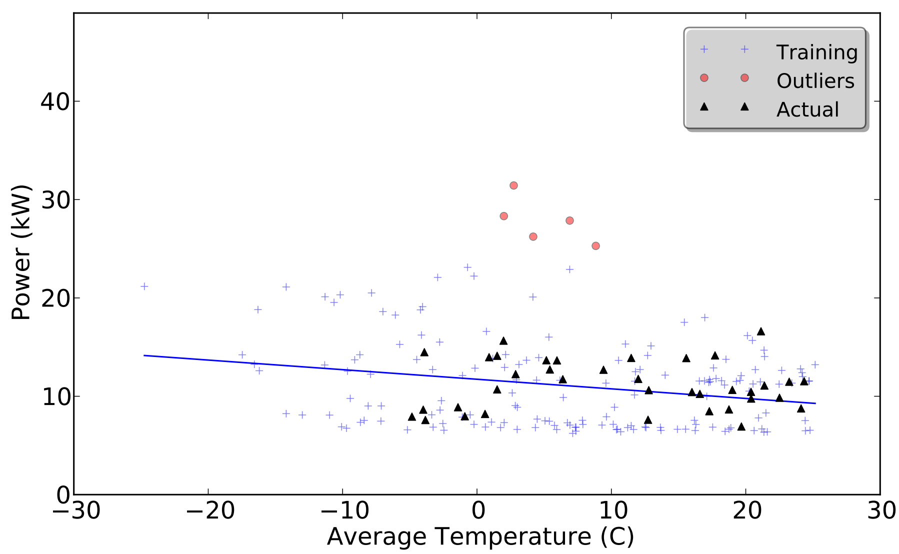

(b) Office 3: Weekend Base Energy Signature

Figure 4.22: Office 3 weekend energy signatures (fold 2).

Chapter 5

Implementation

The software for this thesis, written in the Python programming language [27], has been documented, packaged and distributed to our industry partner. Developing this thesis in a popular, open-source, object-oriented programming language means that this work can be more easily deployed and integrated into a production system. To demonstrate this, we used the Numpy [28], Scipy [29], and SK-Learn [30] scientific libraries for Python to develop the core numeric functionality of this thesis; these libraries have fast and efficient C data structures and functions. We extended this core functionality by implementing the data store with PostgreSQL [31], and developing a web-based user interface (UI) using the Django web framework [32]. This architecture is illustrated in Figure 5.1.

The web-based UI was used to generate many of the plots in this thesis. In addition to the web-based UI, we created an Application Programming Interface (API) web service that allowed us to programmatically generate plots using Matplotlib [33] and Highcharts [34]; this provided a useful platform for rapid prototyping and experimentation. Screenshots of the web-based user interface can be seen in Figures 5.2-5.4. It is our hope that this implementation will assist in the adoption of the algorithms described in this thesis.

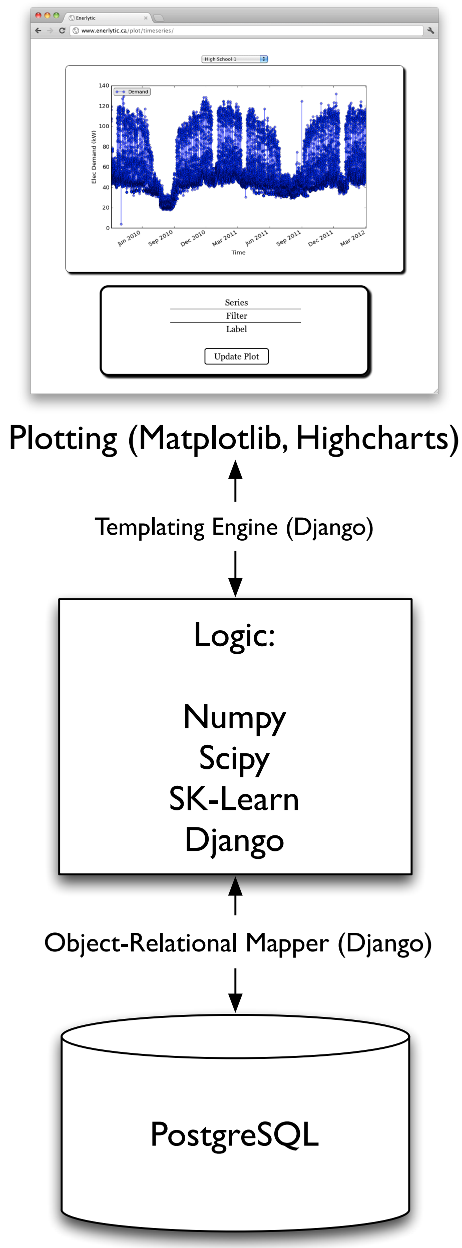

Figure 5.1: Implementation architecture.

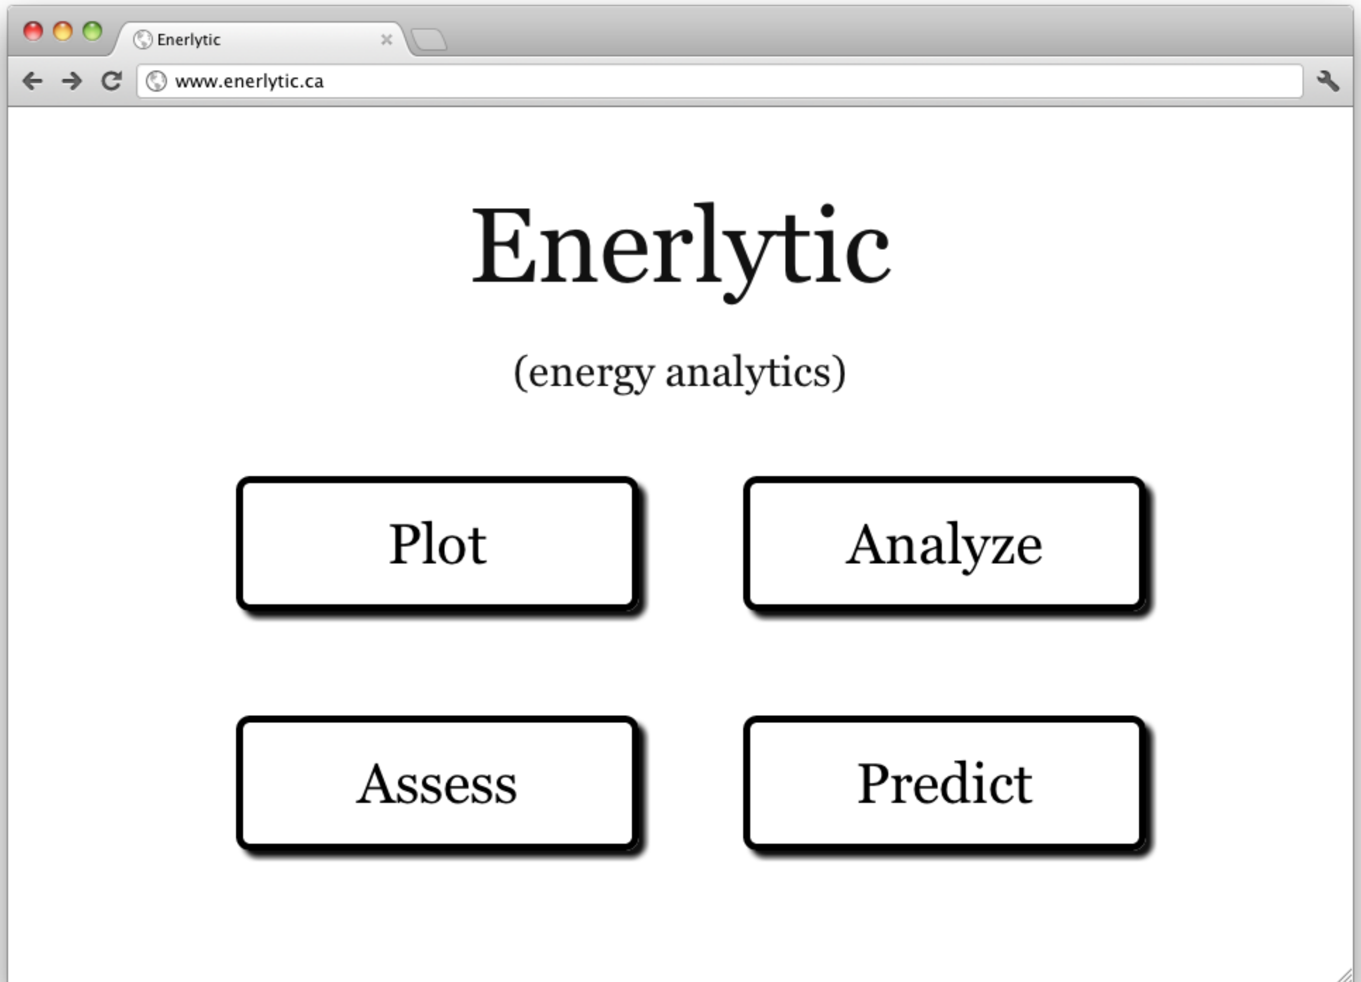

Figure 5.2: Web-based user interface: Home page.

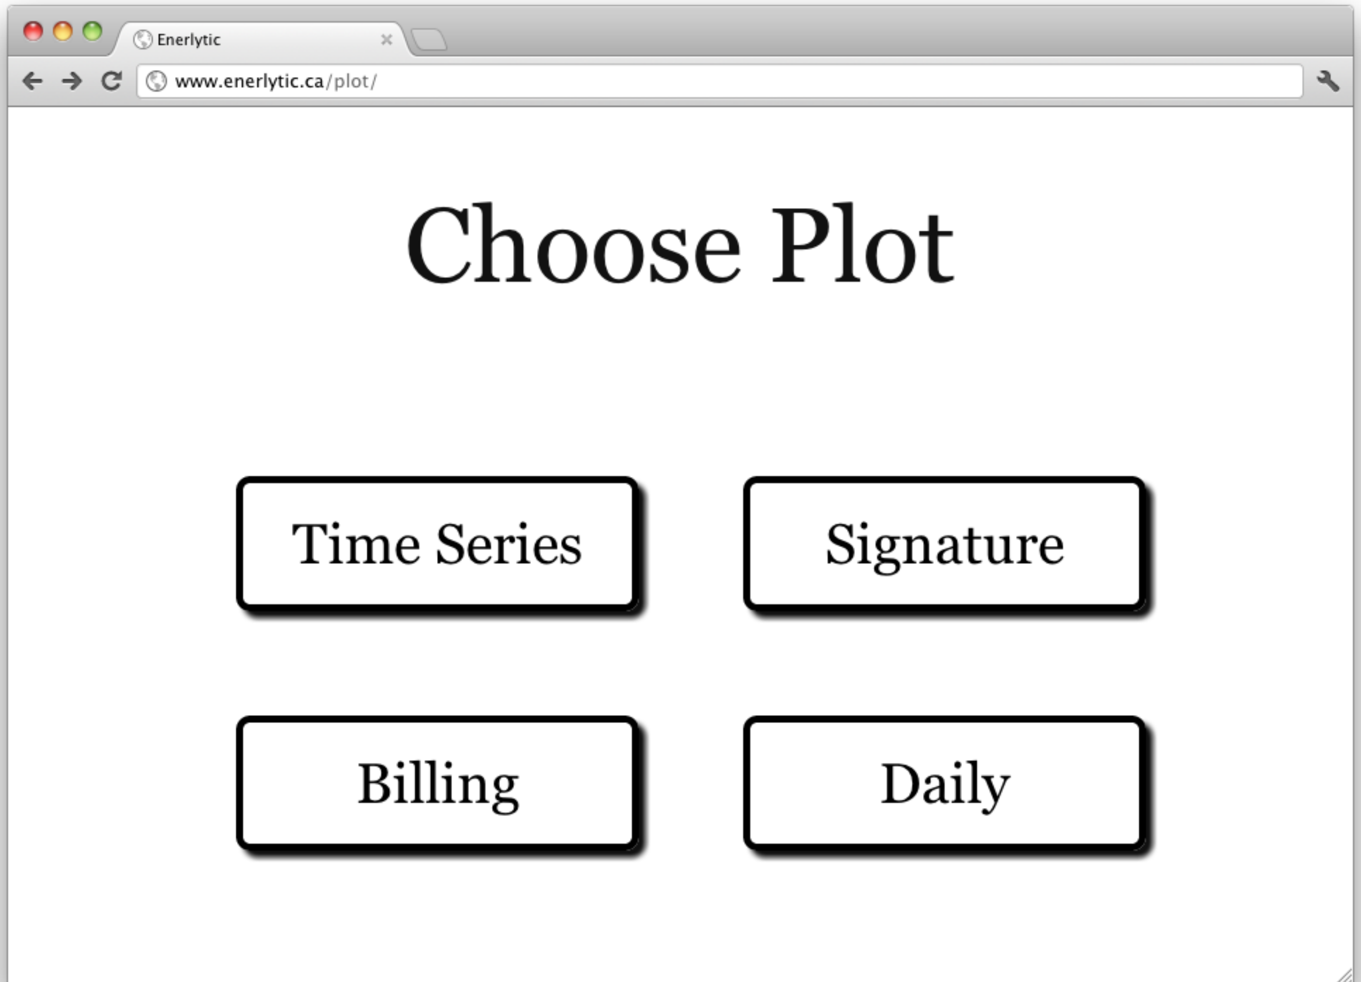

Figure 5.3: Web-based user interface: Plot selector.

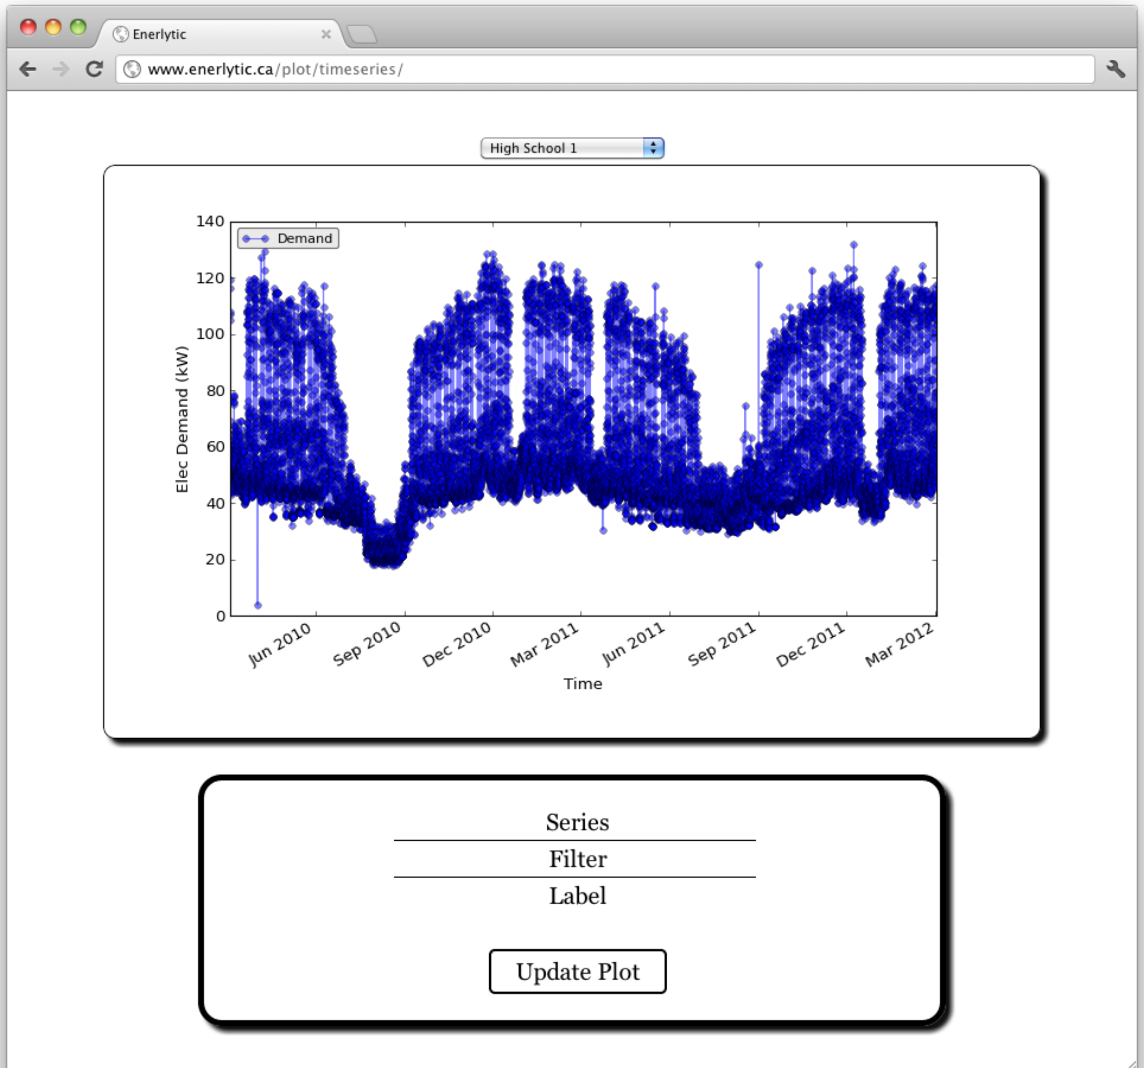

Figure 5.4: Web-based user interface: Plot viewer.

Chapter 6

Conclusion

This thesis describes the development, implementation, and evaluation of Enerlytic, an algorithm that takes as input electrical meter data from a commercial building and outputs its operating parameters. These operating parameters can be used by an energy expert to rapidly assess the building's energy consumption patterns without a time consuming and expensive site visit. Moreover, we have shown how these operating parameters can be used as input to other systems for prediction and estimating possible energy savings. A building operator can use Enerlytic to determine potential cost savings from implementing energy savings measures without engaging the services of an expensive energy expert.

Prior work includes the use of piecewise linear models for estimating energy savings in buildings [16] [17], and partitioning the daily load profile for prediction [19] or demand response [18]. However, this work either uses daily averages of energy useage, or does not develop a model which can be manipulated to obtain energy savings estimates. Our approach takes into account the various operating modes of a building through labelling. It models the peak and base load as a function of outdoor air temperature using piecewise linear regressions. The algorithmic creation of the building operating parameters enables future work, such as benchmarking a portfolio of buildings, demand response, and novel visualization methods, as discussed below.

We tested our approach on a dataset containing 10 buildings and 21 years of meter data; we show that the building operating parameters offer comparable performance to a black-box prediction method used in the industry [9]. This thesis is the product of a year long collaboration with Pulse Energy. The formulation of this problem is a direct result of working with the energy experts at Pulse Energy and from conversations with commercial building operators. The software for this thesis has been documented, packaged, and distributed to an industry partner. We hope that this thesis will allow energy experts and building operators to easily realize financial and environmental benefit from the algorithmic analysis of commercial meter data.

6.1 Limitations

There are several limitations to this thesis.

1. The prediction accuracy of the building operating parameters is comparable to the benchmark model in the majority of buildings in our dataset; in some buildings, however, the prediction accuracy is significantly lower. We attribute this to poor regression fits. To address this, we suggest an iterative regression mechanism that fits a set of regression models and considers the tradeoff between model complexity versus an error metric, such as the CV-RMSE prediction error. Models in this set could include non-linear or bin-based [20] models.

2. We use only one benchmark for the accuracy of Enerlytic's predictions; this benchmark is not the most accurate prediction method available [9]. Future work could consider evaluating the prediction accuracy of the building operating parameters against other prediction models such as those described in references [14] [18] [19].

3. Our work considers only electrical energy; it is possible to apply Enerlytic to other fuel sources, such as natural gas.

4. We have not verified that energy savings measures estimates lead to energy savings in buildings. Future work could verify this with help from industry partners.

5. We introduce a mechanism for labelling the operating modes of a building, but test the prediction accuracy using only the default labeller. We hope that future work will create case studies evaluating the effect of labelling on both prediction accuracy and providing insight into energy consumption patterns.

6. Enerlytic uses a piecewise regression with a single change-point, so that it is applicable to buildings with either electric heating or cooling or neither. Future work can investigate buildings with both electric heating and cooling.

7. We have evaluated our work using only hourly meter data; we encourage future work to investigate the effects of different data resolutions.

6.2 Future Work

In addition to addressing the limitations of this thesis, we believe there are many opportunities for extending our work. We would like to quantify the computational cost of Enerlytic compared with other models such as Brown et al [9]. We suspect Enerlytic, due to its relative simplicity, is considerably faster. If true, we believe this is a significant opportunity. If predictions can be generated on-the-fly, they do not need to be cached or persisted; this greatly simplifies the implementation, maintenance and scalability of any production system that uses Enerlytic. On-the-fly predictions may also lead to improved user experience by allowing a user to interact with the building's operating parameters in real time.

We believe the building operating parameters can be used for many other applications. Several that are not explored in this thesis include: novel methods for visualization, such as visualizing the average peak and base start times for an operating mode; evaluating demand-response opportunities using peak-load shifting; and benchmarking a portfolio of buildings. Future work should formalize the building operating parameters by specifying an application programming interface (API).

References

[1] United States Energy Information Administration. Electric Power Annual 2011, Table ES1. http://www.eia.gov/electricity/annual/pdf/tablees1.pdf. Accessed: 29/06/2012.

[2] Natural Resources Canada. Total End-Use Sector-Energy Use Analysis. http://oee.nrcan.gc.ca/corporate/statistics/neud/dpa/tablesanalysis2/aaa_ca_1_e_4.cfm. Accessed: 29/06/2012.

[3] Natural Resources Canada. Total End-Use Sector-GHG Emissions. http://oee.nrcan.gc.ca/corporate/statistics/neud/dpa/tablesanalysis2/aaa_00_1_e_4.cfm. Accessed: 29/06/2012.

[4] United States Energy Information Administration. Frequently Asked Questions: How many smart meters are installed in the U.S. and who has them? http://www.eia.gov/tools/faqs/faq.cfm?id=108&t=3. Accessed: 29/06/2012.

[5] Ontario Independent Electricity System Operator. Smart Grid Forum Report, May 2011. http://www.ieso.ca/imoweb/pubs/smart_grid/Smart_Grid_Forum-Report-May_2011.pdf. Accessed: 29/06/2012.

[6] B.C. Hydro. Continuous Optimization Success Stories: Jawl Properties. http://www.bchydro.com/etc/medialib/internet/documents/power_smart/success_stories/A11_530_success_story_Jawl.Par.0001.File.A11-530-success-story-Jawl.pdf. Accessed: 29/06/2012.

[7] B.C. Hydro. Continuous Optimization Program. http://www.bchydro.com/powersmart/business/commercial/continuous_optimization.html. Accessed: 29/06/2012.

[8] B.C. Hydro. Continuous Optimization Program Services Funding for Commercial Buildings Agreement. http://www.bchydro.com/etc/medialib/internet/documents/power_smart/commercial/c_op_agreement_schedule.Par.0001.File.C-OP-Agreement-Schedule.pdf. Accessed: 29/06/2012.

[9] Matthew Brown, Chris Barrington-Leigh, and Zosia Brown. Kernel regression for real-time building energy analysis. *Journal of Building Performance Simulation*, 5(4):1–14, August 2011.

[10] DB Crawley, LK Lawrie, and CO Pedersen. Energy plus: energy simulation program. *ASHRAE Journal*, 4(42):49–56, 2000.

[11] D.B. Crawley, L.K. Lawrie, C.O. Pedersen, F.C. Winkelmann, M.J. Witte, R.K. Strand, R.J. Liesen, W.F. Buhl, Y.J. Huang, R.H. Henninger, et al. EnergyPlus: An Update. In *Proceedings of SimBuild, Building Sustainabilty and Performance Through Simulation*. Boulder, August 2004.

[12] JS Haberl and TE Bou-Saada. Procedures for calibrating hourly simulation models to measured building energy and environmental data. *Journal of Solar Energy Engineering*, 120:193–204, 1998.

[13] J.F. Kreider and J.S. Haberl. Predicting hourly building energy use: the great energy predictor shootout- overview and discussion of results. *ASHRAE transactions*, 100:1104–1118, 1994.

[14] DJC MacKay. Bayesian nonlinear modeling for the prediction competition. *ASHRAE Transactions*, 100:1053–1062, 1994.

[15] Kyoung-ho Lee and James E Braun. Development and Application of an Inverse Building Model for Demand Response in Small Commercial Buildings. In *Proceedings of SimBuild*, pages 1–12, 2004.

[16] J K Kissock, T A Reddy, and D E Claridge. Ambient-Temperature Regression Analysis for Estimating Retrofit Savings in Commercial Buildings. *Journal of Solar Energy Engineering*, 120(3):168, 1998.

[17] Franc Sever, JK Kissock, and Dan Brown. Estimating Industrial Building Energy Savings using Inverse Simulation. *ASHRAE Transactions*, 117(1):348–352, 2011.

[18] Johanna L Mathieu, Phillip N Price, Sila Kiliccote, and Mary Ann Piette. Quantifying Changes in Building Electricity Use, With Application to Demand Response. *IEEE Transactions on Smart Grid*, 2(3):507–518, 2011.

[19] Vladimir Cherkassky, Sohini Roy Chowdhury, Volker Landenberger, Saurabh Tewari, and Paul Bursch. Prediction of Electric Power Consumption for Commercial Buildings. In *2011 International Joint Conference on Neural Networks (IJCNN 2011 - San Jose)*, pages 666–672. IEEE, August 2011.

[20] M.R. Brambley, S. Katipamula, and P. ONeill. Diagnostics for monitoring-based commissioning. In *Proceedings of the 17th National Conference on Building Commissioning*, 2009.

[21] S. Lloyd. Least squares quantization in pcm. *Information Theory, IEEE Transactions on*, 28(2):129–137, 1982.

[22] E.H. Simpson. The interpretation of interaction in contingency tables. *Journal of the Royal Statistical Society. Series B (Methodological)*, 13(2):238–241, 1951.

[23] M. Ester, H.P. Kriegel, J. Sander, and X. Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of the 2nd International Conference on Knowledge Discovery and Data mining*, volume 1996, pages 226–231. AAAI Press, 1996.

[24] J.D. Toms and M.L. Lesperance. Piecewise regression: a tool for identifying ecological thresholds. *Ecology*, 84(8):2034–2041, 2003.

[25] D.L. Sonderegger, H. Wang, W.H. Clements, and B.R. Noon. Using sizer to detect thresholds in ecological data. *Frontiers in Ecology and the Environment*, 7(4):190–195, 2008.

[26] Pulse Energy. Energy Management Software for Utilities, Commercial and Institutional Buildings. http://www.pulseenergy.com. Accessed: 29/06/2012.

[27] Van Rossum, G. and others. Python Programming Language. http://www.python.org. Accessed: 29/06/2012.

[28] Oliphant, T.E. and others. Numpy. http://www.numpy.org. Accessed: 29/06/2012.

[29] Eric Jones, Travis Oliphant, Pearu Peterson, et al. SciPy: Open source scientific tools for Python. http://www.scipy.org. Accessed: 29/06/2012.

[30] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine Learning in Python . *Journal of Machine Learning Research*, 12:2825–2830, 2011.

[31] The PostgreSQL Global Development Group. PostgresSQL: The world's most advanced open source database. http://www.postgresql.org. Accessed: 29/06/2012.

[32] Django Software Foundation. django. http://www.djangoproject.com. Accessed: 29/06/2012.

[33] John D. Hunter. Matplotlib: A 2d graphics environment. *Computing In Science & Engineering*, 9(3):90–95, May-Jun 2007.

[34] Highsoft Solutions AS. Highcharts. http://www.highcharts.com. Accessed: 29/06/2012.

APPENDIX

Appendix A

Description of the Dataset

The algorithm presented in this thesis is evaluated on a dataset that spans a variety of commercial areas and geographies. Table A.1 tabulates the details of this dataset.

Table A.1: Description of the dataset

Building	Location	Duration
Hospital 1	Western Canada	2 years, 29 days
Hospital 2	Western Canada	2 years, 104 days
Office 1	Western Canada	2 years, 104 days
Office 2	Eastern Canada	3 years, 60 days
Office 3	South-Eastern United States	2 years, 14 days
Office 4	South-Eastern United States	1 year, 227 days
Office 5	North-Western United States	1 year, 222 days
School 1	Western Canada	1 years, 363 days
School 2	Western Canada	2 years, 104 days
Store	Western Canada	2 years, 102 days