Single-Pixel Thermopile Sensors

for people counting applications

Erik Hagenaars

Master Thesis

TU Delft

Single-Pixel Thermopile Sensors

for people counting applications

by

Erik Hagenaars

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on February 28, 2020 at 12:00 PM.

Student number:	4272404	
Project duration:	April 1, 2019 – February 1, 2020	
Thesis committee:	Prof. dr. ir. G. J. T. Leus,	TU Delft, supervisor
	Dr. R. T. Rajan,	TU Delft
	Dr. A. Pandharipande,	Signify

This thesis is confidential and cannot be made public until February 28, 2022.

An electronic version of this thesis is available at http://repository.tudelft.nl/.

TU Delft

Abstract

People counting data in offices is used in many applications like HVAC system control and space management to increase comfort, decrease energy consumption and optimise space utilisation. In contrast to past approaches using imaging modalities that tend to be either expensive or intrusive, we consider single-pixel thermopile sensors for people counting. These sensors may already be deployed as part of a smart lighting system to provide temperature data for HVAC controls.

Firstly, a statistical sensor model for thermopile temperature measurements is proposed. The proposed people counting method enhances the CUSUM RLS algorithm to estimate temperature change caused by people entering or leaving. We estimate mean temperature changes upon detection of an occupancy event, and then estimate the people count using a maximum likelihood on the estimated temperature change. Finally, PIR vacancy data is merged with the people count estimation to increase accuracy. We obtain an average counting error of 0.11 and 0.19 for 90% of the instants respectively when considering 15 minute windows for simulated and experimental datasets.

A second aspect of the thesis considers the problem of commissioning plan detection. We leverage the two-sided CUSUM signals to address this problem. The two-sided CUSUM scores for a pair of sensors are used to calculate similarity measures; these features are used in a Random Forest Classifier to detect commissioning changes of the sensor pair. Using simulated data with the thermopile signal model, we show that the proposed method achieves a true positive rate (determining the correct layout) of 90.2% and false positive rate of 1.3%.

Preface

Before you lies the master thesis ‘Single-Pixel Thermopile Sensor for people counting applications‘. A thesis where we create a statistical model for the sensors in an office setting and tackle the problems of people counting and commissioning change detection. The sensors and its applications had not yet been researched widely. It has been written as a master thesis for Electrical Engineering at the Technological University Delft. The research was published in February 2020.

The research was conducted at Signify in Eindhoven, the world leader in lighting for professionals, consumers and lighting for the Internet of Things. Signify provided me with the research question and facilitated the requirements for this research such as hardware, data and guidance.

There are many people I would like to acknowledge before the start of the thesis. Without these people, I would have never been able to finish this thesis and stand were I stand today.

This thesis would not be possible without the guidance of my Daily Supervisor, Ashish Pandharipande. Ashish has taught me many valuable lessons in the field on research. I’d also like to thank my professor, Geert Leus.

Many thanks to everyone involved. From coworkers to co-students who evaluated this thesis and helped during brainstorms.

I hope you enjoy your reading.

Erik Hagenaars
Eindhoven, February 2020

Contents

1 Introduction

1.1. Problem description

1.1.1. Space management

Effective space management can have a huge impact on the performance and total costs for businesses and universities. Research has found that in some cases, up to 20% of the work spaces are effectively used at universities[1]. In closed IT work spaces, the average vacancy is up to 65%[2]. It is therefor not a surprise that the costs of space utilisation and real estate can be significantly cut. ‘Some companies use real estate as a lever to improve their competitive position‘[3]. Cost savings on real estate can be cut up to 20%[4] and for many companies, this would result in savings that run into the millions[5].

Most of the solutions to tackle this problem is converting closed ‘cubicle‘ work spaces to an open work space environment[2]. Together with smart planning, and optimising seat utilisation[6]–[8], big improvements in space utilisation can be achieved. Not only will optimisation improve in cost saving, when converting a closed office setting into an open office setting, collaboration between colleagues is encouraged and different work styles can be accommodated[2].

To improve upon the existing planning optimisation, sensors are used to track and detect occupancy. With the insight of occupancy patterns and trends, better decisions in space utilisation can be made. Occupancy Detection and People Tracking is a widely researched field. Many solutions are incorporating image sensors[9]–[12]. These solutions achieve a high accuracy in both occupancy detection and people tracking. The main problem with image sensors is the intrusive nature on privacy. And with new privacy laws in certain areas, image sensors cannot be used. Recently researched are the ultrasonic sensors[13], [14] and binary sensor networks[15], [16]. Sadly, the accuracy was found to be limited with these systems.

With the up-rise of cheaper thermopile arrays, low pixel infrared cameras are becoming more interesting in tracking people[17]–[22]. These sensors will increase heavily in price in relation to the amount of pixels.

1.1.2. Commissioning changes

Many network algorithms and systems rely on a known and static topology for an optimal functionality[23]. A good example is localisation using Time Difference of Arrival (TDA) methods where the measured TDA with the known locations of the sensor can accurately localise the source. If the locations of the receivers are to change, the localisation estimation will decrease in performance. Aside from a known topology, foreground segmentation requires a static background. When a background changes, the segmentation of the foreground will result in a larger foreground set, which results in greater errors.

These changes in topology or background are not uncommon to occur in open offices. New projects, changes in group hierarchy and employee numbers have the possibility to result in a commissioning change: the scenario where the sensor topology or the measured area of interest is changed.

Various smart building systems like lighting and HVAC (heating, ventilation and air conditioning) controls and space management applications rely on data from sensor systems [24], [25], [26], [27]. For instance, temperature data and people count estimates from thermopile sensor systems may be used to improve user comfort by adaptively controlling HVAC systems. People counting data may also be used in space management applications to optimise use of workspaces based on historic trends as well as real-time information. Commissioning information in the form of the sensor plan indicating locations of sensors in an area is critical in all these systems for proper functioning.

1.2. Scope of the research

Since image sensors are too intrusive, thermopile arrays are still too expensive and other sensor networks lack in accuracy, the research on this thesis will focus on tackling the problem with the utilisation of a single-pixel thermopile sensor. Previous works have only tackled the problem with thermopile arrays instead of a single-pixel sensor[17]–[21].

Single-pixel thermopile sensor is that it can be easily embedded in other systems such as lighting systems. Since PIR sensors are recently being embedded in lighting systems as well, its vacancy flag will be available for approaching a solution to the problem.

When embedding connected sensors in lighting, it becomes important to detect any commission changes in the setup to avoid estimation errors. This research had already been conducted using a network of PIR sensors with overlapping FoVs[28]. In this report, an implementation involving thermopile sensors will be tested.

In short, the key contributions of the thesis will consist of the following topics:

- Constructing a statistical signal model of a single-pixel thermopile sensor.
- Detecting people enter or leave events from the thermopile sensor.
- Estimate the number of people involved in this event.
- Embed the PIR vacancy flags and people estimated to track the number of people.
- Detect commissioning changes for sensors with overlapping FoVs

1.3. Setting of the project

This project was conducted during an Electrical Engineering master at the TU Delft. For the project, an internship at Signify was accepted. Signify is the new company name of Philips Lighting. They are the world leader in lighting and provide their customers with high-quality, energy-efficient lighting products, systems and services. Signify turns light sources into points of data to connect more devices, places and people through light, contributing to a safer, more productive and smarter world.

Signify leads in LED lighting innovation and is at the forefront of the industry's expansion into systems and services. The problem description and resources were provided by Signify. Some of the information is company confidential and will not be disclosed in this report.

1.4. Outline of the thesis

This thesis consists of 6 chapters. First, background information and the state of the art will be researched in chapter 2. Before determining an approach, the signal model of the sensor needs to be known. This will be constructed in chapter 3. After this the approach to the two problems will be discussed in chapter 4. Results of the approaches will be discussed in chapter 5. And lastly, the conclusions will be drawn and a discussion is opened in chapter 6 and chapter 7 respectively.

2

Background Information

2.1. Occupancy detection and people counting

2.1.1. Image Sensors

The problem of occupancy detection and people counting is widely researched for image sensors[9]–[12]. Tracking a human body has already been very accurately for cameras with a low resolution and low sampling frequency thanks to the Pfinder algorithm[29]. The algorithm is able to classify the head and hands from various angles. The algorithm uses a combination of foreground and background segmentation and blob finding/tracking methods.

Privacy and Legal Issues

Counting people with cameras is not a big problem anymore. The main focus in the field is tracking a single person in large groups [30]. From single camera setups, it is becoming more interesting to track a single person in a crowd over a multiple and non-overlapping camera network [31].

However, the use of cameras in public spaces has some privacy- and legal issues[32]. The debate on surveillance and privacy is increasing in relevance. Due to these concerns, some areas are not allowed to use cameras for tracking and surveillance purposes.

An approach to defeat the privacy problem is counting people in an image without tracking individual persons. To avoid explicit use of segmentation and tracking, a Gaussian Process model was proposed[9]. The algorithm extract motion parts from multiple frames of a image sensor. The features from blobs moving in a certain direction are captured and inserted into a Gaussian Process model which estimated the number of people. The features per motion direction are the area, perimeter, perimeter edge orientation and perimeter-area ratio. The Gaussian model is trained with the help of a regression of a training set. The accuracy of this approach resulted to be above 90%.

Social Issues

Guaranteeing privacy of people in the field of view of cameras using alternative algorithms does not solve all the problems. Firstly, it is hard to proof that privacy is contained and organisations collecting and processing the data have to be trusted. Secondly, there are many social effects on people's behaviour when installing visible cameras.

A study about the behaviour and consent of people exposed to cameras was conducted[33]. In the study it found that people acted differently when being recorded and there were occasions where participants showed concerns about being recorded.

Another study researched the effect of lifelogging your daily life and its legal and ethical issue[34]. The study claims that many people feel uncomfortable when getting filmed. There are even cases of physical assault to the filmer. It is therefor safe to say that the use of cameras for the thesis' problem is not desired.

2.1.2. Thermopile Arrays

Previous works of thermopile sensors have only used them in a network as an array, creating a very low resolution thermal image[17]–[22]. The most popular method of detection people in the images is thresholding the pixels[17]–[19]. Each blob or area of detected pixels is classified as one or more people depending on the size or shape of the area. Results varied in the range of 80% to 90% accuracy.

Other approaches included the use of background extraction of the room temperature and the use of a distinction map[20] or use an entropy-based machine learning classifier to detect occupancy[21]. These approaches had a slightly worse performance than the use of thresholding.

The last approach used only 4 sensors and could track a single person by exploiting the overlapping properties of the sensors[22]. This research was more focused on tracking a single person in location than to count the number of people. Its maximum error was about 30cm.

While thermopile sensors are decreasingly expensive, when the approach requires a 2D array, cost will increase substantially. As previous mentioned in the previous section, creating a 2D image will cause the same ethical problems of the image sensors.

2.1.3. Ultrasonic Sensors

An interesting approach for occupancy detection and people localisation is with the use of ultrasonic sensors[13], [14]. The first approach focuses on the localisation of a person using a ultrasonic sensor array. The setup is said to be energy efficient and cheap enough to embed in lighting. By analysing the direction of arrival, a localisation of a person can be estimated.

The second approach is able to detect occupancy for zones by using distributed transmission slot synchronisation. When a person moves, the measured echos are distinctive enough to detect occupancy in zones.

2.2. Thermopile sensor

A thermocouple generates a voltage dependent on the temperature difference between the joints[35]. The couple consists of two junctions, a cold and a hot junction. When this temperature difference increases, the voltage increases as well. A thermopile sensor connects these thermocouples in series. The cold junctions are placed at the thick chip substrate while the hot junction is placed over a thin membrane. This membrane is able to absorb the IR radiation radiated by objects. Dependent on the magnitude of the IR radiation, the membrane heats up or cools down.

When a reference temperature on one of the joints is known, the temperature measured on the other junction can be calculated based on this voltage difference. This reference temperature needs to be measured at the cold junction. The relation of the voltage and the temperature of the reference temperature and the object temperature is given in Equation 2.1.

$$V_{ir}\left(T_a, T_o\right) = A \cdot \left(T_o^4 - T_a^4\right) \tag{2.1}$$

In the equation, V_{ir} stands for the voltage difference measured, T_a for the reference temperature, also known as the ambient temperature. T_o stands for the average object temperature and A is the sensitivity of the sensor. A figure of the functionality of a thermal couple is shown in Figure 2.1.

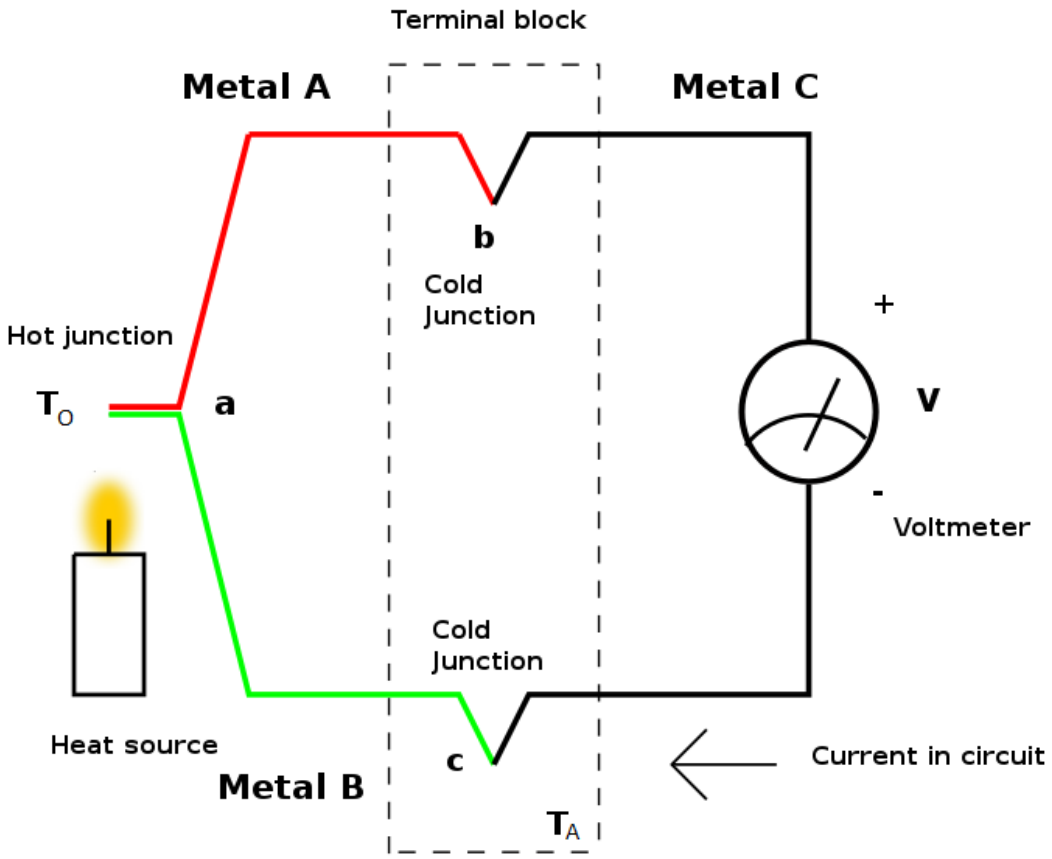

Figure 2.1: Schematic of a thermocouple generating a voltage potential.

2.3. Change point detection

In previous years, many research has been done on change point detection. A useful survey[36] gives a nice overview of all the current state-of-the-art methods for detecting change points. The survey splits the problem into supervised and unsupervised methods. For this thesis, we will consider the subcategories for the unsupervised methods which are listed below:

1. Likelihood Ratio Approach: by either calculating the probability densities of both windows or by calculating the direct ratio.
2. Subspace Model Approach: by finding parameter changes in a subspace model.
3. Probabilistic Approach: by finding the probability of a change point
4. Kernel Based Approach
5. Graph Based Approach
6. Clustering Approach

Another source for change point detection is Adaptive Filtering and Change Detection by Fredrik Gustafsson[37]. In the book, four types of change detection methods are explained:

1. Adaptive filtering and whiteness based change detection
2. Maximum likelihood based change detection
3. Multiple-model based change detection
4. Algebraic (parity space) change detection

2.4. Commissioning change detection

With the uprising of IoT devices and networks, ad-hoc networks become more common. Network theory and distributed (non-centralised) algorithms are increasing in popularity. For many old systems, a known topology is required for optimal performance of the network algorithms. For example, every network that exploits the knowledge of overlapping sensor measurements are heavily dependent on a known topology. Localisation algorithms often require topology and geological prior knowledge as well for localising and tracking object. For topology changes, re-calibration is required.

Another important method still used today is foreground extraction or background segmentation. Many (image) sensor first filter out background noise to collect relevant data. An example of these methods are the monitoring of certain objects[38], [39]. When the background changes, the algorithm needs to re-calibrate.

Often these re-calibration need manual interference from technical employees. The automatic re-adjustment or re-calibration will reduce the work required for the applications to perform optimally. An automatic detection of a commissioning change will aid in the initialisation of such an event. We will consider 2 different approaches to the detection of a commissioning change: (i) traditional algorithms and (ii) machine learning approaches.

Similarity Features The most common and traditional approach for detecting commissioning changes is by the property of similarity measures[40]–[45]. Most of these approaches are constructing a similarity measure for two 2D images of large landmarks. Many of these similarity features incorporate Shannon entropy or correlation and apply these values to a threshold for decision making. One fresh paper proposed a semi-supervised SVM measurement for classification.

Machine Learning Tackling the problem of change detection or commissioning change detection with the help of machine learning has been researched mainly for image applications[46]–[49]. While time series applications are less common[50], [51]. The most helpful resource was done using a binary network of which binary similarity features were used to use in a Random Forest Classifier[51]. Of all the used features, similarity features, correlations and differences between the images have been used to find a commissioning change.

3

Statistical Model

There has not been much research on the single-pixel thermopile sensor and its possibilities in different applications. The majority of the research in occupancy detection and people counting was with a setup of multiple thermopile sensors to mimic an infrared image[17]–[22]. In this chapter, we will present a signal model for a single-pixel thermopile sensor for people counting applications in an office setting. To our knowledge, this has not been done yet and will help in the approach to solve the people count problem.

The statistical model will be deducted from real measurements observations for multiple office scenarios. All the relevant features observed will be incorporated into the model. A balance between complexity and simplicity is very important. A more complex model could include more features and possible represent the data more accurately. However, deducting an optimal solution for a complex model is not always possible. Some features might not be relevant for finding the solution and thus simplicity might be preferred.

We've observed that the signal is constructed of three components:

1. the signal in vacancy
2. the signal during a transition
3. the attenuation model due to the sensor's Field of View (FoV)

In this chapter we will first gather information based on the sensor's datasheet in section 3.1. In section 3.2, we will construct a model of the signal in vacancy and add the transitional model in section 3.3. Finally, the resulting model is presented in section 3.4 and a framework for simulating the data is explained in section 3.5.

In Figure 3.1, the object temperature signal from a single-pixel thermopile sensor mounted at the ceiling is shown for varying occupancy levels. The temperature can be observed to gradually increase in the first hour, even under constant occupancy. This is due to the room temperature being controlled by the HVAC. It can also be observed that the mean object temperature under vacancy is not constant over time and depends on factors like HVAC control and occupancy levels prior to occupancy (e.g., the temperature decreases gradually when people count changes from 2 to 0 around 14:40).

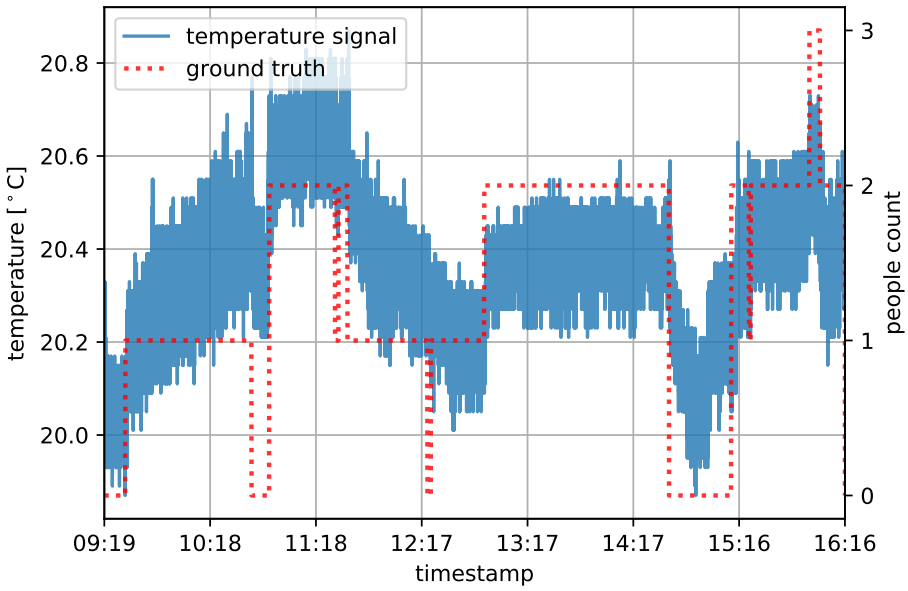

Figure 3.1: Example of an experimental measurement taken during working hours.

3.1. Sensor datasheet

The datasheet[52] of the sensor is more focused on configuration and communication methods and does not go much into details with respect to the generation of the signal or the internal signal processing.

The benefits of the sensor is its small size and low cost. The datasheet claims to be easy integratable with the help of a SMBus communication protocol. This protocol is similar to the I2C protocol, designed by Philips Semiconductors. A high accuracy of 0.5°C and a measurement resolution of 0.02°C is guaranteed for normal operational temperatures between 0°C and 50°C. The module also offers a sleep mode to reduce power usage and the sensor claims to have an in-built digital signal processing block to reduce noise and guarantee an accurate measurement.

The sensor measures two components:

- the average object temperature in the sensor's FoV
- the ambient temperature of the sensor

Of interest in the datasheet is the usage of the digital signal processing (DSP) block and the behaviour of the field of view of the sensor.

3.1.1. DSP block

The sensor uses a 17-bit Analog to Digital Converter which is directly connected to a powerful DSP unit. The sensor collects the raw data of an IR sensor. This signal is amplified by a chopper amplifier with low noise, low offset settings and dynamic attenuation which can be configured by the user. The signal is then converted to a digital signal with the help of a Sigma Delta modulator. The output of this ADC is connected to the input of the DSP unit. The block diagram of the DSP block is shown in Figure 3.2.

The DSP unit consists of two filters: (i) a FIR filter and (ii) an IIR filter which are both configurable. These configurations can reduce the bandwidth of the signal to get a desired noise performance but will increase the refresh rate as well. For this project, the factory settings will be used provided by the manufacturer. However it was shown that the noise level could be reduced by a large factor while keeping an acceptable refresh rate.

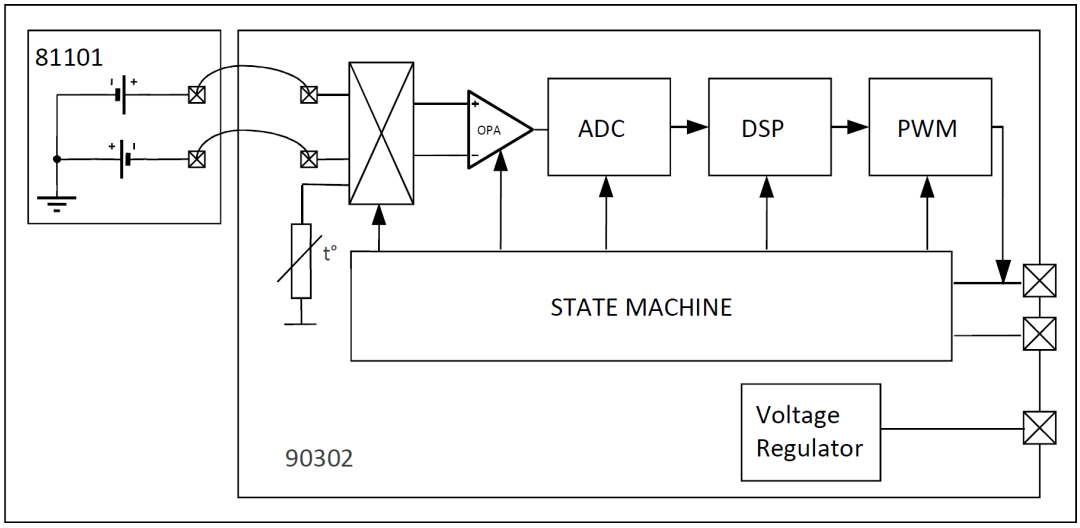

Figure 3.2: The digital signal processing block of the sensor from the datasheet[52].

The output of the IIR filter is then stored in an EEPROM register. This result can be fetched by a Master device using the SMBus protocol.

The last piece of information given by the datasheet is the configurable settings of this attenuation, FIR filter and IIR filter. All of the configurations have 8 options which can be selected by 3 bits each. This configuration can be shown in Table 3.1.

bit1	bit2	bit3	gain	FIR	IIR
0	0	0	1	8	a1=1, b1=0
0	0	1	3	16	a1=0.8, b1=0.2
0	1	0	6	32	a1=0.666, b1=0.333
0	1	1	12.5	64	a1=0.571, b1=0.428
1	0	0	25	128	a1=0.5, b1=0.5
1	0	1	50	256	a1=0.25, b1=0.75
1	1	0	100	512	a1=0.1666, b1=0.833
1	1	1	100	1024	a1=0.125, b1=0.875

Table 3.1: Configuration of the DSP Block.

3.1.2. Field of view

Another important feature mentioned in the datasheet is the angle of incidence attenuation for the computation of the average object temperature. Assuming the sensor is installed on a ceiling and directed downwards. This angle is defined between the horizontal line from the sensor and the line from the object towards the sensor. The attenuation for the given angle was measured by the manufacturer and is shown in Figure 3.3.

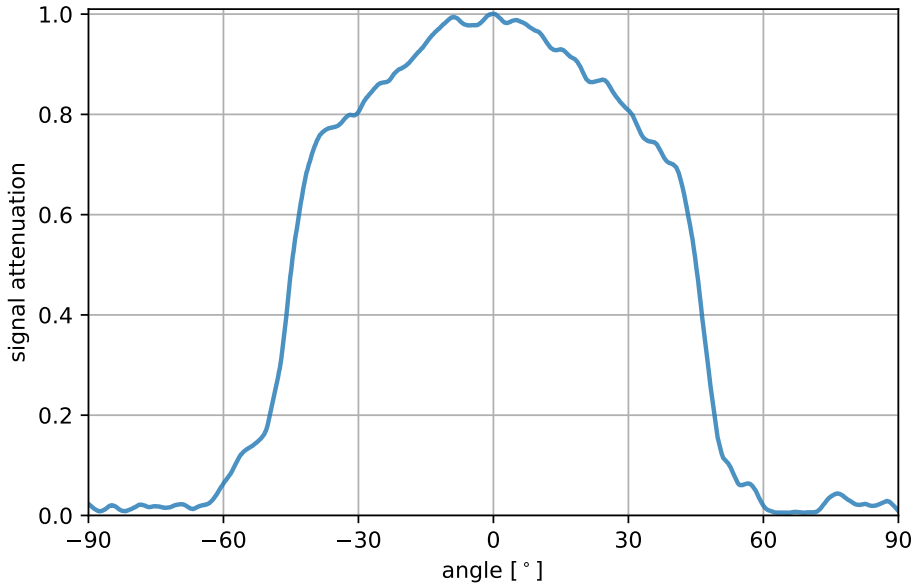

Figure 3.3: Measured attenuation per angle according to the datasheet.

For the graph, the measurements were done in two directions and is notated from -90° to 90°. From Figure 3.3, we see that the -3dB attenuation threshold is found at +- 45°. Concluding that the sensor has a effective field of view of 45° in total. It can also be seen that the x-axis runs from -90° to 90° degrees to display symmetry in the signal attenuation based on the angle. This means we only require one angle parameter to calculate the attenuation from 0° to 90°.

Instead of using the measured attenuation to incorporate in the model, a functional approach is constructed to decrease the complexity of the model while still mimicking the dynamics of this feature. A good function to model this -3dB threshold and a slope approaching unity which is also symmetrical around $0°$ would be the raised cosine function. Which is defined in Equation 3.1.

$$f(x) = \begin{cases} 1, & |f| \leq \frac{1-\beta}{2T} \\ \frac{1}{2}\left[1 + \cos\left(\frac{\pi T}{\beta}\left[|f| - \frac{1-\beta}{2T}\right]\right)\right], & \frac{1-\beta}{2T} < |f| \leq \frac{1+\beta}{2T} \\ 0, & \text{otherwise} \end{cases} \tag{3.1}$$

The raised cosine function has two parameters. $1/2T$ denotes the -3db intersection which can naturally set to be 45. The β parameter affects the slope magnitude at this intersection. A lower β will result in a steeper slope. The effect of this parameter is plotted in Figure 3.4.

We fit the β parameter of the raised cosine function to the measured data of the manufacturer. We specifically fit around the -3dB threshold to mimic this slope at 45°. A value of $\beta = 0.4$ is found to represent this slope very well as can be found in Figure 3.5. We only use the (3.1) from 0° to 90° since we assume symmetry in the Z-axis.

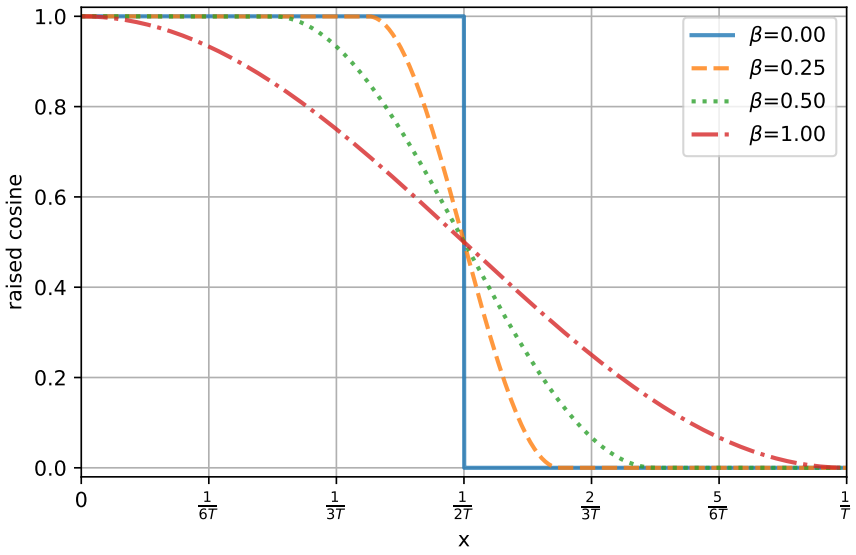

Figure 3.4: Raised Cosine function for different β.

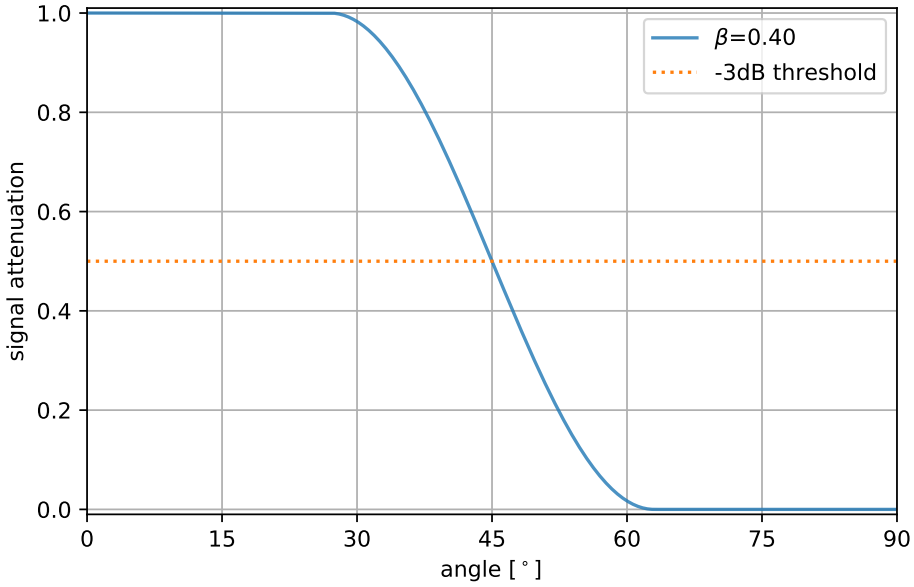

Figure 3.5: Raised Cosine function for $\beta = 0.4$ which represent the datasheet's attenuating curve.

3.2. Vacancy model

During periods of vacancy, the temperature variations are typically slow and happen due to factors like daylight changes and HVAC control. We model the thermopile temperature signal over vacancy by a slowly varying value $T_0[n]$ with an additive noise component $\nu[n]$,

$$x[n] = T_0[n] + \nu[n]. \tag{3.2}$$

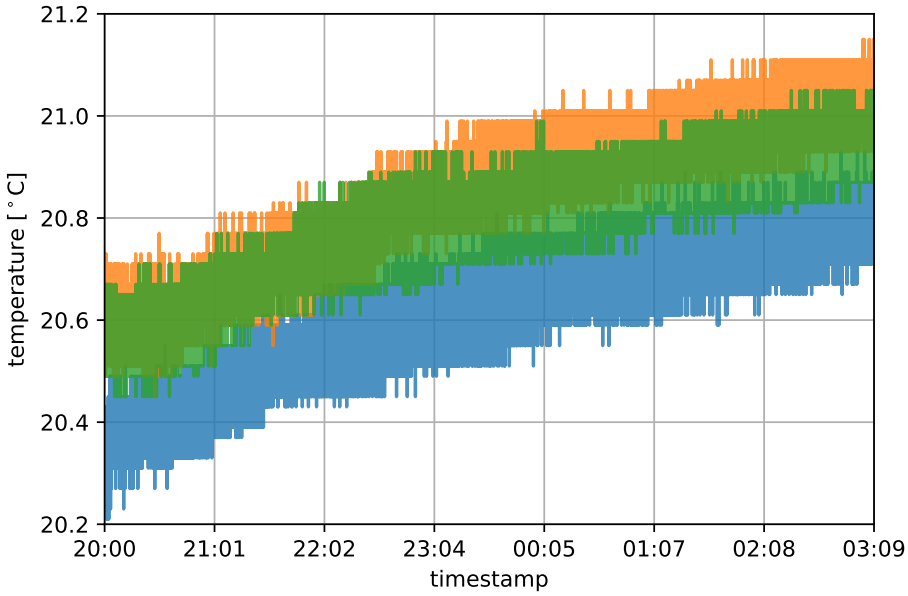

Figure 3.6: Example of vacancy data for 3 sensors measured during night.

In Figure 3.6, we show object temperature data from three thermopile sensors over a 7 hour period under vacancy. The gradual temperature rise is on account of the HVAC system being disabled at night. The measured object temperature was tested for normality by testing the null hypothesis if the sample given comes from a normal distribution. The test combines the skew and kurtosis to test this property based on the D'Agostino-Pearson test[53], [54]. The normality test returns the probability that a null hypothesis can be rejected. We choose this decision for a threshold of 10^{-3}. When the null hypothesis rejection probability is lower than this threshold, a normal distribution can be assumed. In Figure 3.7, we display the results of this normality test. We observed that $\nu[n]$ can be modelled as an additive white Gaussian noise (AWGN) component when considering a short duration of time. For a window of 9000 samples (15 min), 97% of the windows could be considered normally distributed.

Due to the low resolution of the sensor, only a few values will be measured in a timed window. When there are only a few distinct values measured, the normality test will fail while the noise could still be normal. To increase the resolution, the samples will be smoothed over a very short window allowing intermediate values between the samples. After smoothing, the normality test is done over a dataset with vacancy.

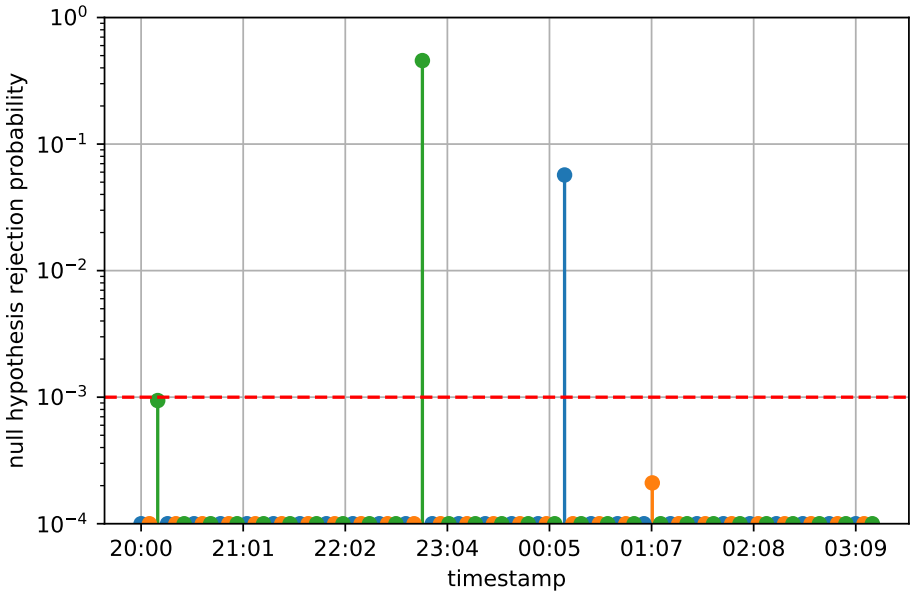

Figure 3.7: Normality test for 3 sensors. For samples below the threshold, a normal distribution can be assumed for that window.

3.3. Transitional model

The presence of one or more people produces a temperature shift in the object temperature of the thermopile sensor. The shift depends on the number of people, their locations and secondary aspects like orientation. Furthermore, the shift in temperature transitions faster during entry of people and slower when people leave, due to occupants having heated the environment and leaving a thermal footprint that fades with time.

The magnitude of this mean change can have a few explanations. First thing that is apparent is the angle of incidence discussed in section 3.1. Based on the location of the person, the effect that person has can be limited by the lower attenuation caused by the angle of incidence. Another big impact on this magnitude is the number of people involved. However, from observations, the magnitude change of multiple people will not be constant for the same number of people. We expect that the previous number of people and new number of people is explicable for this magnitude shift. During the approach of the problem, we will go into more detail on this expectation.

When observing the transitions from a closer perspective, a small smooth transition can be found. The transitional speed is different for an entry and leave event. This can be explained by the heat that is introduced and left behind respectively. When a person enters, the only heat to be introduced is from that person. Very slowly over time, it will heat up its environment. When leaving, a person will leave a lot of heat behind on the workspace and seat. In general, this cooling will cause for a longer transition, since the heated parts only get exposed during a leave event.

Zooming into a transition at the samples, we see a similar characterisation of an exponential function $e^{-\alpha n}$. This can be explained by the attenuation from the angle of incidence. The attenuation will increase when entering the FoV and it will decrease when leaving. The small difference per sample for later samples can be explained due to the heating and cooling events.

Construction the transitional model results in:

$$d_i[n] = \Delta T_i \left(1 - e^{-\alpha_i (n - n_i)}\right) u[n - n_i]. \tag{3.3}$$

In this equation, $d_i[n]$ is the change function for change event i, ΔT_i is the total temperature change over the full transition. $\left(1 - e^{-\alpha_i (n - n_i)}\right)$ is the exponential function of the change which approaches ΔT_i over time with a transitional speed parameter α. For large α, the transition is quicker. The change function is relevant at a change point $n - n_i$ at which the function is initialised by a step function $u[n - n_i]$.

3.4. Resulting model

Putting together the earlier signal components, we can now write the thermopile signal model for K_m events followed by initial vacancy as

$$x[n] = T_0[n] + v[n] + \sum_{i=1}^{K_m} \sum_{j=1}^{K_i} f(\theta_{ij}) d_i[n]. \tag{3.4}$$

In (3.4), till time sample n_1, there is vacancy and the temperature at an initial temperature T_0 with AWGN $v[n]$. At time sample n_i, we consider K_i occupants with the j-th occupant at angle θ_{ij} with respect to the sensor. At event i, there is a temperature shift ΔT_i with a transition speed α_i.

A simulated example of the dataset using the model in (3.4) with 10 occupancy events is shown in Figure 3.8.

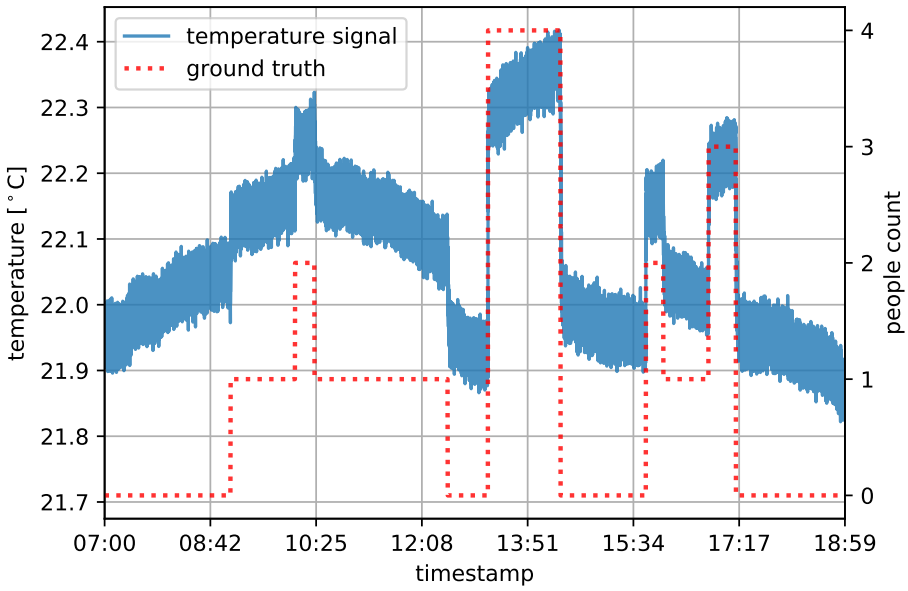

Figure 3.8: Example of a simulated dataset over a day.

3.5. Simulations

Not only can we find a good approach for the solution of the problem with this model, but we can also create many simulation based on the given parameters. A performance test per parameter can be achieved by exhaustively iterating over the range of this given parameter. By creating many simulations for different parameter values, the algorithm can be tested and sensitivities to certain parameters can be found.

The first thing to simulate is the binary occupancy per workspace. An important trend is that the first activity can occur from 07:00 and will last until 19:00. There should be vacancy outside these timestamps. Each workspace has an angle of incidence with respect to the sensor. The easiest way to simulate the activity of a workspace is to have an even number of events per workspace.

The initialisation of the signal will be a constant with AWGN. Per workspace, the transitions will modify the initial signal. Do this for every event and every workspace will result in the final simulated signal. The simulation will be stored together with the ground truth which is created by summing all the ground truths of the individual workspaces.

For the simulations, a standard and physical setup will only be considered for benchmark where the parameters will closely resemble the measurements. When performing an exhaustive performance test, these parameters will be iterated through a certain range uniformly. An example of the simulation code is given in Algorithm 1.

Data: N, M, T_0

$events \leftarrow generateEvents(N, M)$
$x \leftarrow random(N) + T_0$
for $event$ ***in*** $events$ **do**
    $cp \leftarrow event.cp$
    $g(\theta_i) \leftarrow event.g$
    $d_i \leftarrow event.d$
    for g ***in*** $g(\theta_i)$ **do**
        $x[cp:] \leftarrow x[cp:] + g \cdot d_i$
    end
end

Result: $\mathbf{x}$

Algorithm 1: Simulation dataset generation pseudo-code.

4

Approach and Implementation

With the statistical model presented in chapter 3, we can more easily approach the solution for the two problems: (i) people counting using a single sensor and (ii) detecting a commissioning change by classifying the sensor layout.

The most important feature found in the model is the transitional model during events. Our main focus will be to emphasise our approach on this feature to increase the performance. In section 4.1, we will approach the first problem and propose an implementation. Our approach for problem (ii) will be discussed in section 4.2.

4.1. People counting

The approach to the people count problem will be divided into a few sections. Firstly, spike anomalies from the sensor requires filtering in subsection 4.1.1. A quick and simple implementation of an interquartile filter will be discussed for this. Secondly, a change point detection algorithm is discussed in section 4.1.2. At the change point, its magnitude needs to be estimated in section 4.1.3. This change magnitude is converted into a people count. Eventually, the merging of the PIR sensor will be explained in section 4.1.4, where the vacancy flags of the PIR sensor is used to more accurately track the number of people. A design of the full system can be seen in Figure 4.1. For each function block, a subsection is created. The arrows depict data flow at the speed of the sampling rate of 10Hz.

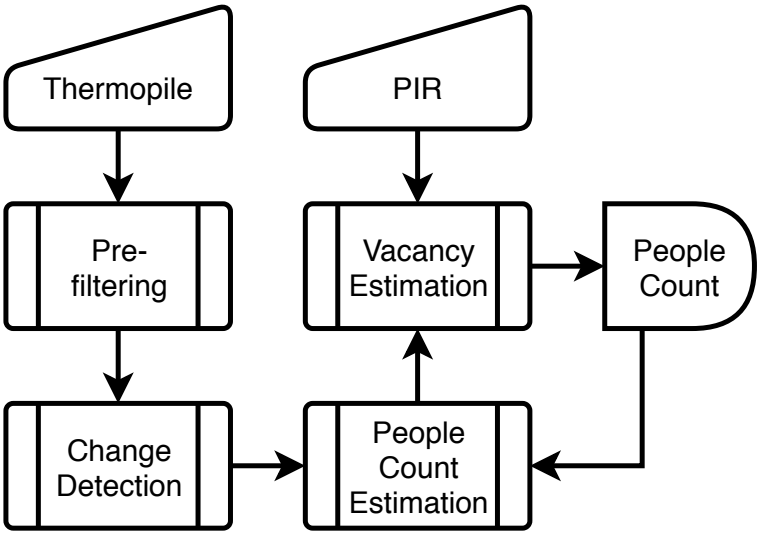

Figure 4.1: System design for people counting using a single-pixel thermopile sensor and a PIR sensor.

4.1.1. Pre-filtering

From observations of the sensor measurements, anomaly spikes were detected on occasion. The spikes consisted of a single sample being much greater than the noise. There are many options of filtering a spike. Sadly the spikes are not always big or small enough to use a static threshold filter. The spikes are often too big to be filtered through a moving average filter since it would affect the filtered signal too much. A good approach for filtering these anomalies are through the means of z-score filtering[55], also known as the standard score which is the number of standard deviations a value is away from the mean given a sample set. For the windowed set, the mean and standard deviation is calculated. The calculation of the z-score of a single sample given the sample set is given as:

$$z[n] = \left| \frac{x[n] - \mathrm{E}\left[\mathbf{x}\right]}{\sqrt{\mathrm{VAR}\left[\mathbf{x}\right]}} \right| \tag{4.1}$$

When the z-score exceeds a certain threshold, we would identify it as an anomaly. To keep the sampling intervals uniform, instead of dropping the sample, the mean value is given to that sample. The resulting output of the filter becomes:

$$y[n] = \begin{cases} x[n], & z[n] \leq 3 \\ \mathrm{E}\left[\mathbf{x_{w_n}}\right], & z[n] > 3 \end{cases} \tag{4.2}$$

In this instance, x_{w_n} is the selected window surrounding sample n. An example of the functionality of this filtering can be seen in Figure 4.2.

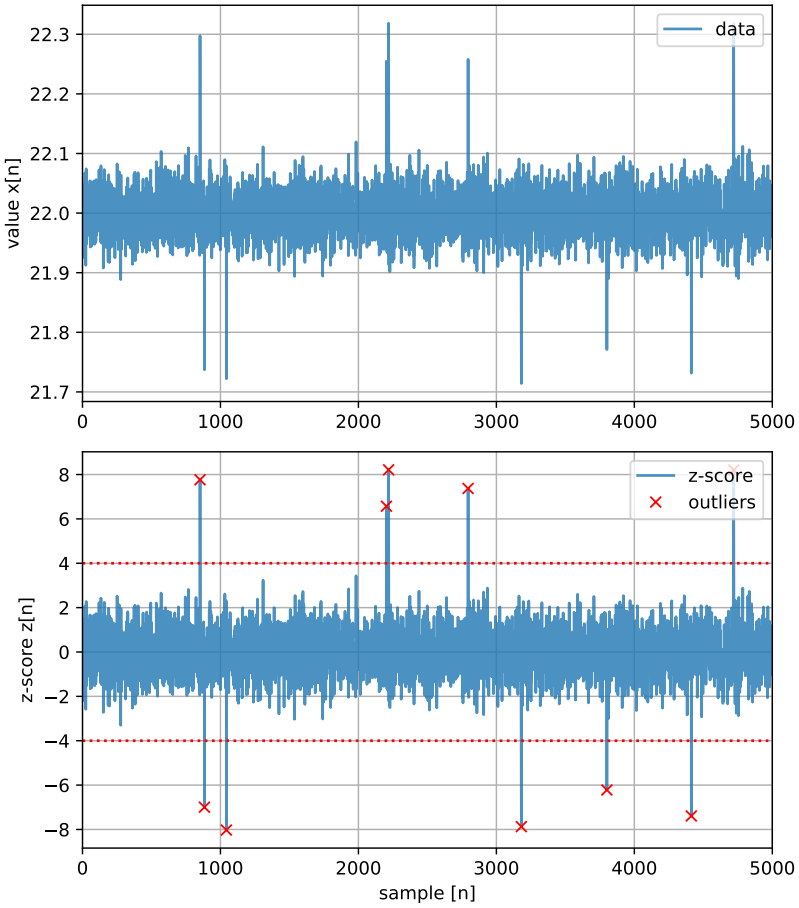

Figure 4.2: Example of z-score filtering of anomalies and spikes.

4.1.2. Change detection

After filtering the anomalies, change points should be detected. There are many change point detection algorithms as found in chapter 2. The goal is to find an algorithm to exploit its properties and adapt it to accurately estimate the magnitude of the change and eventually estimate the number of people involved in the event. The chosen algorithm for this algorithm is the well known CUSUM algorithm. While many of the applications of the algorithm are used in the medical field of work. The CUSUM algorithm can accurately detect mean changes in data.

There are many reasons for preferring the CUSUM algorithm. Benefits include the low complexity of the algorithm and the online nature to detect change points. Another property of the CUSUM, the creation of a two-sided score signal will be useful for the problem of commission change detection as will be discussed later. In section 4.1.3, a method is proposed to exploit the score properties of the CUSUM to directly and accurately estimate the magnitude of the change.

There are many versions of the CUSUM algorithm. The algorithm with the least amount of parameter tuning and the one which is independent on initial constant value is the CUSUM RLS proposed by Gustaffson[37]. This algorithm is shown in Algorithm 2.

Data: $\mathbf{y}$, λ, v, h

$rls \leftarrow y[0]$
$e \leftarrow 0$
$g_p \leftarrow 0$
$g_m \leftarrow 0$
ids $\leftarrow [\,]$
for $i \leftarrow 1$ **to** N **do**
  $rls \leftarrow \lambda\, rls + (1-\lambda) y[i]$
  $e \leftarrow rls - y[i]$
  $g_p \leftarrow \max\left[0,\ g_p + e - v\right]$
  $g_m \leftarrow \min\left[0,\ g_m + e + v\right]$
  if $g_p > h$ ***or*** $g_m < -h$ **then**
    $rls \leftarrow y[i]$
    $g_p \leftarrow 0$
    $g_m \leftarrow 0$
    ids$.append(i)$
  end
end

Result: **ids**

Algorithm 2: Traditional CUSUM RLS algorithm

The algorithm introduces three parameters available for tuning: (i) λ, the forgetting factor of the RLS estimate, (ii) v, the drift for filtering noise spikes and (iii) h the detection threshold. For each iteration, the following variables are updated:

$$\hat{T}[n] = \lambda \hat{T}[n-1] + (1-\lambda) y[n] \tag{4.3}$$

$$\epsilon[n] = y[n] - \hat{T}[n] \tag{4.4}$$

$$g_p[n] = \max\left[0, g_p[n-1] + \epsilon[n] - v\right] \tag{4.5}$$

$$g_m[n] = \min\left[0, g_m[n-1] + \epsilon[n] + v\right] \tag{4.6}$$

When the scores $g_p[n]$ and $g_m[n]$ is larger than h or smaller than $-h$ respectively, a change point would be detected. The algorithm is displayed in Figure 4.3 on the bottom left side where the data (top left) and the error (top right) is plotted. As can be observed in the score, peaks are generated during a transition.

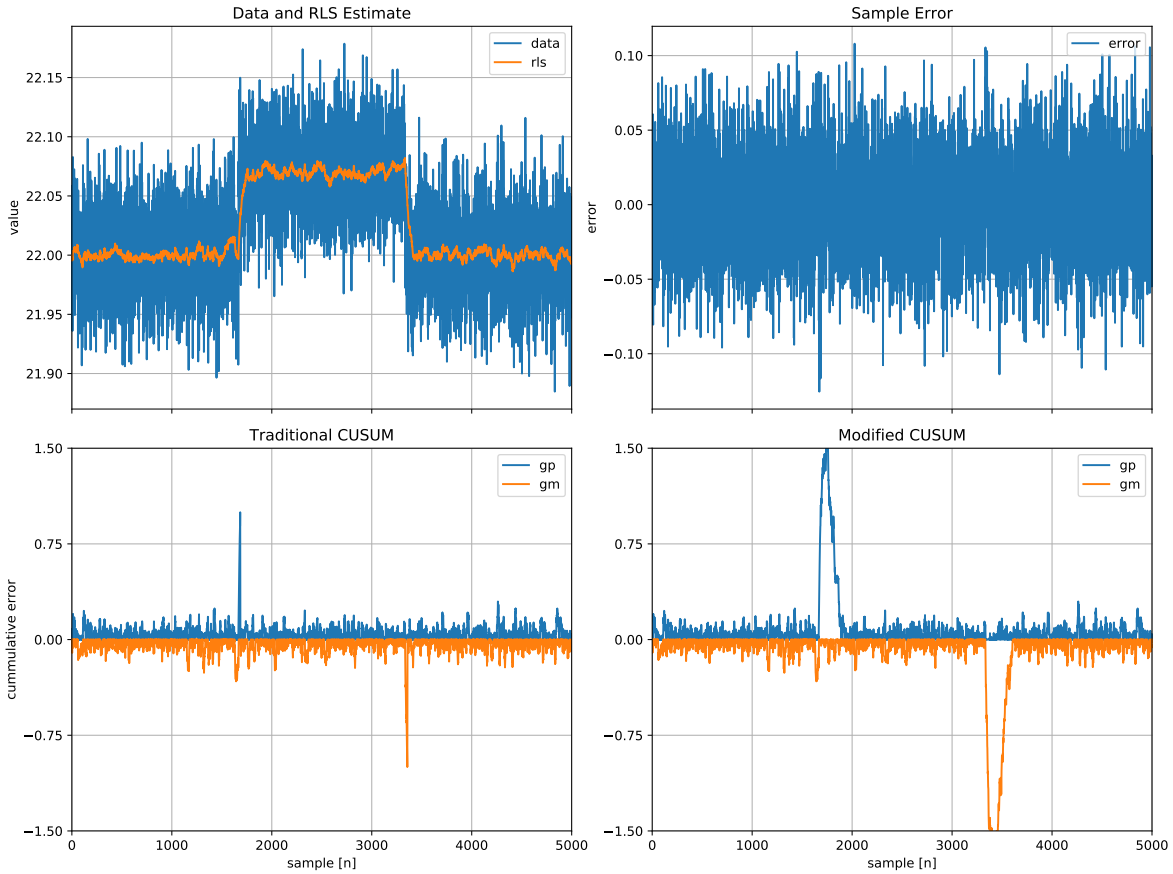

Figure 4.3: Comparison of the traditional CUSUM algorithm versus the modified CUSUM algorithm. top left: data with RLS estimation. top right: error per sample. bottom left: traditional CUSUM score. bottom right: modified CUSUM score.

The three parameters need tweaking. For the forgetting factor, the parameter is bounded by $0 < \lambda < 1$. When approaching one, the estimation is very accurate and removes most of the noise. During a transition, the RLS estimation will converge very slow to the new magnitude when approaching one. For the forgetting factor we thus use $\lambda = 0.97$. Gustaffson[37] proposes the drift to be fitted such that at least half of the score values equals to zero. We find this point when we set the drift equal to the deviation of the noise in vacancy. Lastly the threshold will be set to $h = 0.8$ to include all subtle changes as well.

4.1.3. People estimation

Mean change estimation

Given the change points, the objective in this step is to estimate the magnitude of the change. Originally, the CUSUM algorithm has not been used to directly calculate the mean change magnitude. In this part, we will propose a method to exploit the score of the CUSUM and the RLS estimation to directly calculate the magnitude of the change.

Due to the lag the RLS introduces, there is a small latency for detecting a change. In the notes on scientific computing for bio-mechanics and motor control[56], it was suggested that the start of the change point can be found in the score of the CUSUM. This start is found at the last zero in the score before the change point detection. At this point, we would find the initial temperature calculated by the RLS estimation.

To find the end of a transition we would naturally be looking for the first next zero in the score. Sadly, the original algorithm resets the score upon detecting a change point. If this reset mechanism would be dropped, the score would be increasing during the rest of the transition. At the end of the transition, the score would approach zero again when the RLS settles to the new constant value since the difference is smaller than the drift. When the score returns to zero, we can safely assume that this RLS estimation has settled near the new constant value again. We will use this RLS value as the new constant value.

Together with the RLS value capture before the change point, we simply take the difference to immediately receive the accurate mean change. This temperature estimate is written as:

$$\hat{\Delta T} = \hat{T}[n_r] - \hat{T}[n_l]. \tag{4.7}$$

where n_l and n_r are found with:

$$n_l = i, \text{ when } g[i] = 0, g[j] \neq 0 \text{ for } i < j < n_p, \tag{4.8}$$

$$n_r = i, \text{ when } g[i] = 0, g[j] \neq 0 \text{ for } n_p < j < i. \tag{4.9}$$

with g the selected score which exceeds the threshold:

$$g = \begin{cases} g_p, \text{ if } g_p[n_p] > h \\ g_m, \text{ if } g_m[n_p] < -h \end{cases} \tag{4.10}$$

Here, n_l and n_r denote the start and end respectively of the transition and the settling point of the RLS estimate to the new temperature level and n_p is the sample at detection.g is either g_p or g_m depending on which of the scores exceeded the threshold.

In short, by dropping the reset mechanism in the original CUSUM algorithm, the differences of the RLS estimations at the first zero before and after the change point detection, we obtain an accurate estimation of the change magnitude. The new and proposed algorithm is shown in Listing 3. The effect can be seen in Figure 4.3 on the bottom right.

In Figure 4.4, we can see a signal with mean changes on the top. The bottom left shows the CUSUM signal. The algorithm detects change points and their magnitudes. On the bottom right, we see the true mean in green and estimated mean based on the mean change estimates in orange.

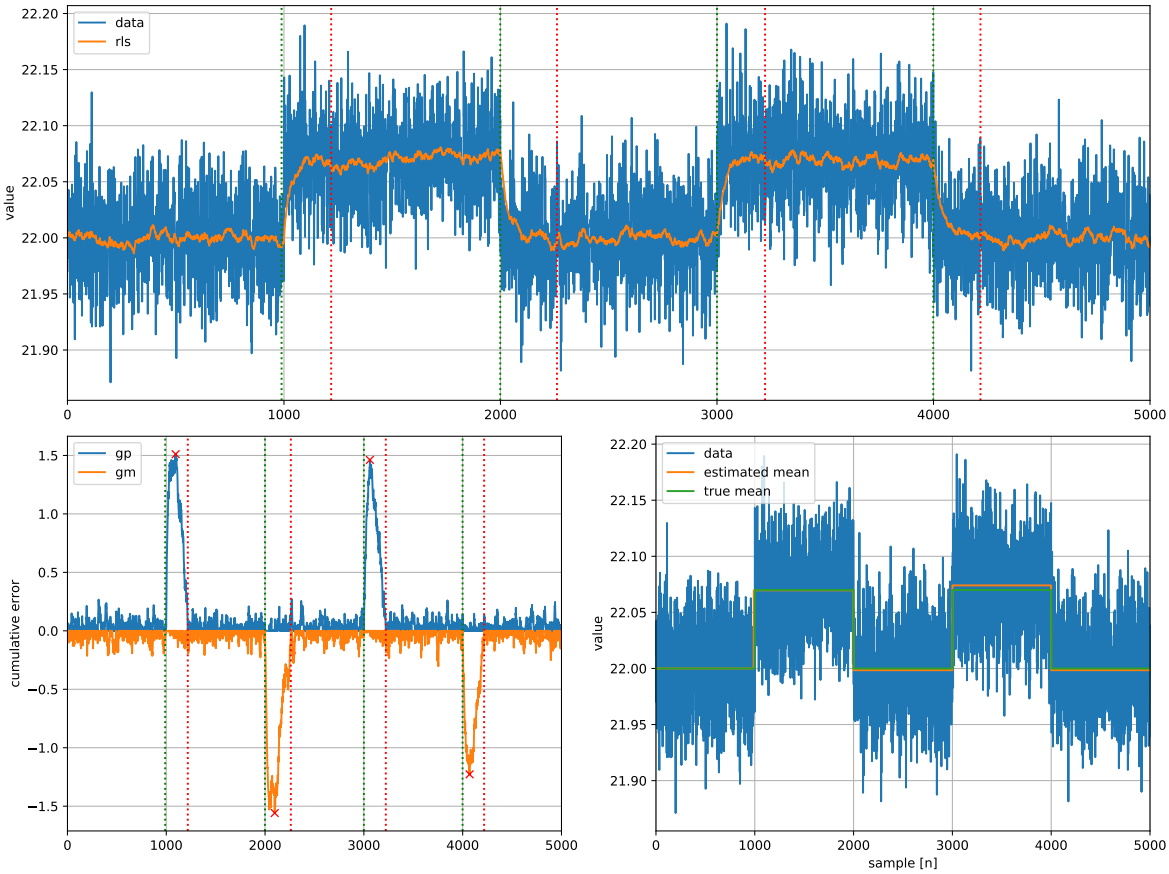

Figure 4.4: Example of constant mean estimation by estimating the mean change per change point. top: data with RLS estimation. bottom left: CUSUM score and detected change point. bottom right: true constant mean and estimated mean.

People count estimation

The temperature differential in (4.7) is used to estimate the people count. This temperature change is dependent on various factors: the number of people leaving or entering the room, the position at which the person(s) are located in the FoV of the sensor, and the previous number of people in the FoV.

For every possible transition for a changing people count, we measure the temperature difference observed during that transition. From these measurements, a probability distribution of the temperature differential under various people count states is estimated. In Figure 4.5, we show the kernel density estimate with a Gaussian kernel for a sensor over an area with two people. The three people count states ($p = 0$, $p = 1$ and $p = 2$) have each two possible transitions for which the estimated probability distributions using an experimental dataset are shown. In general, for an area covered by a sensor with capacity of C people, we would have C states and a total of $C(C-1)$ probability distributions. The distributions can be built using a combination of experimental data and simulated data using the signal model in (3.4).

The people count estimation algorithm works as follows. For the current estimated people count, we consider the corresponding distributions. When a change is detected, for temperature differential $\hat{\Delta T}$, compute the likelihood probabilities $\mathcal{L}(c|p[n-1], \hat{\Delta T})$. We then estimate the number of people $\hat{c}[n]$ corresponding to the temperature change by choosing the highest likelihood and update the people count,

$$\hat{c}[n] = \arg_c \max \mathcal{L}(c|\hat{p}[n-1], \hat{\Delta T}), \tag{4.11}$$

$$\hat{p}[n] = \hat{p}[n-1] + \hat{c}[n]. \tag{4.12}$$

This method is also known as maximum likelihood ratio testing.

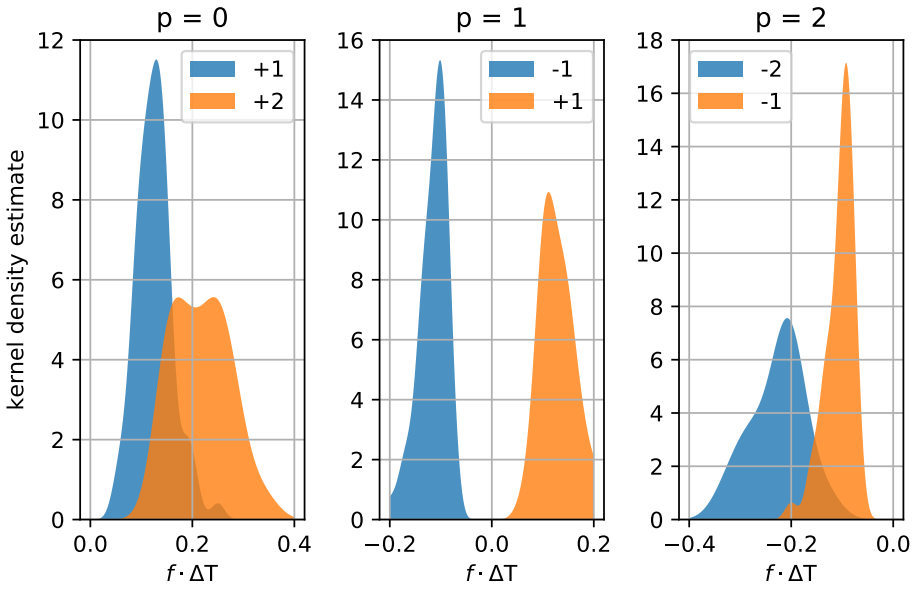

Figure 4.5: Kernel density estimation of the base temperature change for a given number of people and number of people change.

Data: $\mathbf{y}$, λ, υ, h

$\mathbf{rls} \leftarrow \mathbf{y}$
$g_p, g_m, e \leftarrow 0$
$z_p, z_m \leftarrow 0$
$s_p, s_m \leftarrow 0$
$\mathbf{ids} \leftarrow []$
$\mathbf{dts} \leftarrow []$
for $i \leftarrow 1$ **to** N **do**
  $rls \leftarrow \lambda\, rls + (1-\lambda) y[i]$
  $e \leftarrow y[i] - rls$
  $g_p \leftarrow g_p + e - \upsilon$
  $g_m \leftarrow g_m + e + \upsilon$
  if $g_p < 0$ **then**
    if $s_p = 1$ **then**
      $\mathbf{ids}.append(z_p)$
      $\mathbf{dts}.append(rls[i] - rls[z_p])$
      $s_p \leftarrow 0$
    end
    $z_p \leftarrow i$
    $g_p \leftarrow 0$
  end
  if $g_m > 0$ **then**
    if $s_m = 1$ **then**
      $\mathbf{ids}.append(z_m)$
      $\mathbf{dts}.append(rls[i] - rls[z_m])$
      $s_m \leftarrow 0$
    end
    $z_m \leftarrow i$
    $g_m \leftarrow 0$
  end
  if $g_p > h$ **then**
    $s_p \leftarrow 1$
  end
  if $g_m < -h$ **then**
    $s_m \leftarrow 1$
  end
end

Result: $\mathbf{ids}$, $\mathbf{dts}$

Algorithm 3: Modified CUSUM RLS algorithm.

4.1.4. Vacancy estimation

One of the challenges with people counting with the thermopile sensor is slow variation of temperature under vacancy. Furthermore, the dependence of mean temperature differential on human orientation characteristics may lead to errors in people count estimation, and thus possible error propagation. On the other hand, the PIR sensor is typically configured to have low false alarms (at the expense of high missed detections) [57], i.e., it can detect vacancy reliably. Instead of using the raw signal of a typical PIR sensor, the processed signal that controls lighting with built-in delay will be used. We adopt the following decision rule to combine the people count estimate of the thermopile sensor with the PIR sensor,

$$\hat{p}[n] = \begin{cases} (\hat{p}[n-1] + \hat{c}[n]) r_{PIR}[n], & \text{otherwise} \\ 0, & \hat{p}[n-1] < 0.1 \end{cases} \tag{4.13}$$

where

$$r_{PIR}[n] = \begin{cases} 1, & \text{occupancy} \\ r, & \text{vacancy} \end{cases} \tag{4.14}$$

with $0 < r < 1$ a decay factor when vacancy is determined, and where $\hat{c}$ is the estimated people count from the temperature differential. The result of (4.13) will be that the people count is gradually set to 0 when the processed PIR signal determines vacancy. Since the people count estimate with the decay factor will never be zero, we force it to zero when small enough as indicated in the second condition in (4.13).

An example of the functionality of the merging can be seen in Figure 4.6.

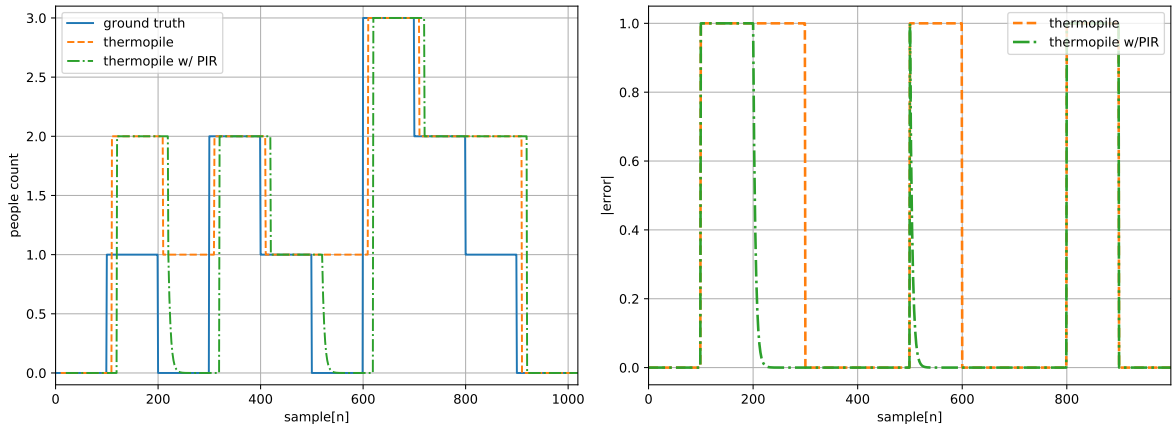

Figure 4.6: Example of the people count estimation with and without the PIR sensor.

4.2. Commissioning change detection

Given data from various thermopile sensors from an indoor space, we consider a 2-class random forest classifier to determine which of the sensor plans exists in that space. We first describe the signal pre-processing and then describe the signal features used in the random forest classifier.

A simple and intuitive approach to a binary classification problem when a large dataset can be provided is with the means of machine learning . During the time of the project, not many data could be collected with respect to commissioning change data. However, with the help of the signal model, the scenarios and classifications can be simulated.

In this section, we will first go over the physical setup and the workflow in subsection 4.2.1. This will give a good overview of the two classifications and an example of the effect per sensor.

When designing a solution for the problem there are two important factors: (i) the features that will play an important role in detecting the commissioning change and (ii) the machine learning method to detect the change. In subsection 4.2.2, we will display two approaches for finding the features. In subsection 4.2.3, we will propose the final implementation by choosing a machine learning method.

4.2.1. Setup

We consider a sensor system with thermopile sensors co-located in luminaires. Sensor and lighting plans may respectively be designed such that there are no blind areas in the sensor coverage region and the illumination from the luminaires is uniform. We shall use a simple illustrative example wherein there are two possible sensor plans, with four sensors in two rows of sensors shown in Figure 4.7 and Figure 4.8. We consider two situations for commissioning detection. One, wherein in an area sensors are situated as per one plan and then moved to a different location as per another design plan. Two, wherein in different areas of a building, sensors are situated according to one of the two sensor plans. The problem in either case is to determine which commissioning exists in an area.

In the first sensor plan shown in Figure 4.7, sensors 1 and 2 have a field of view that covers five workspaces labelled 1, 2, 3, 4 and 7, and 1, 2, 3, 4 and 5 respectively. It can be expected that occupancy over workspaces 1, 2, 3 or 4 would result in similar temperature measurements at thermopile sensors 1 and 2. In the second sensor plan shown in Figure 4.8, sensors $2'$ and $4'$ are displaced locations compared to sensors 2 and 4 in the previous plan. Sensor $2'$ has a FoV that covers five workspaces labelled 2, 5, 6, 7 and 8. In this sensor plan, the measurements of sensors 1 and $2'$ are influenced by quite different occupancy patterns, given their FoV covers quite distinct workspaces.

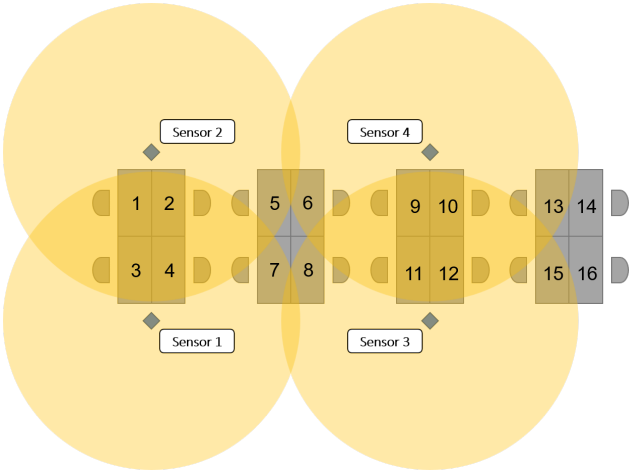

Figure 4.7: Sensor commissioning of an area with 4 sensors covering 12 workspaces. Classification 0

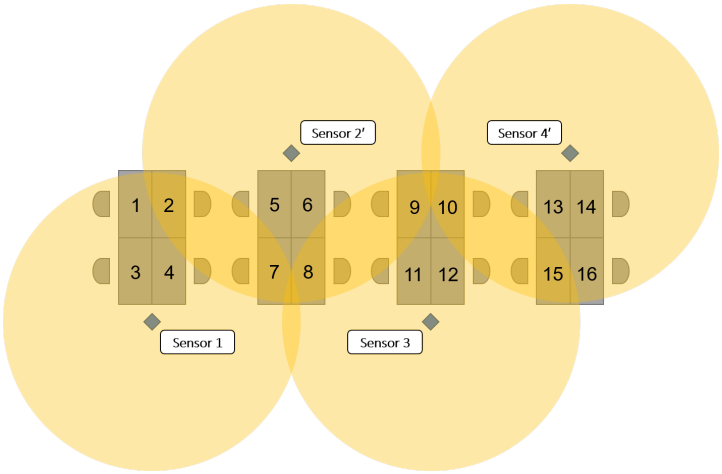

Figure 4.8: Sensor commissioning where the locations of sensors 2 and 4 are displaced. Classification 1

4.2.2. Features

In order to effectively utilise a machine learning approach, the data will be pre-processed. By aggregating the data into a few relevant features, training time and number of required training data will be significantly reduced. This subsection is divided in three paragraphs: (i) the introduction of the CUSUM control chart, (ii) the use of baseline features and (iii) the proposed feature set.

CUSUM RLS processing

In a subsection 4.1.3, a CUSUM algorithm was used to detect events and estimate the magnitude of these changes to track the number of people. The distinct features in the signal model becomes easier to estimate when incorporating this CUSUM algorithm. The CUSUM algorithm will result in two signals, one for negative and positive changes.

The generation of two signals doubles the number of features for the random forest. Adding, the enter and leave events are separated. From the observations, we know that the leave events are more subtle and more difficult to detect. The subtle leave events wont be overshadowed by the enter events.

The two-sided score of the CUSUM also has two built-in filters: (i) the RLS estimation used to calculate the sample error and (ii) the subtraction of the drift suppressing the noise. This way, correlation between the signals are more accurate and less spoiled by the noise.

All these properties of the CUSUM control chart have the potential to increase the machine learning problem solution by a surprisingly large margin. It is therefor interesting to compare the raw signal to the two-sided CUSUM signal when designing the features and random forest.

It is therefore that we have chosen to test two different kind of features: (i) the most generic features based on common signal statistics and (ii) features that are instinctively relevant to solving the problem.

Baseline features

It is common to use statistical features to characterise a signal by considering measures of central tendency, variability, shape and impurity, by considering features like the mean, variance, minimum value, maximum value, skewness, kurtosis and entropy [58], respectively given by:

$$\mu_a = \frac{1}{N}\sum a[n] \tag{4.15}$$

$$\sigma_a^2 = \frac{1}{N}\sum (a[n]-\mu)^2 \tag{4.16}$$

$$a_{min} = \min a \tag{4.17}$$

$$a_{max} = \max a \tag{4.18}$$

$$\text{Skew}[a] = \frac{1}{N}\sum (a[n]-\mu)^3/\sigma^3 \tag{4.19}$$

$$\text{Kurt}[a] = \frac{1}{N}\sum (a[n]-\mu)^4/\sigma^4 - 3 \tag{4.20}$$

$$H(a) = -\sum p(A)\cdot \log p(A) \tag{4.21}$$

Given signals a and b from two sensors, we use the difference of features computed according to (4.15-4.21) and use this feature in the random forest classifier.

Proposed features

We employ different features based on similarity/distance between two signals [59]. Given two signals a and b, the first feature is the Pearson correlation coefficient,

$$\gamma(a,b) = \frac{E[(a-\mu_a)(b-\mu_b)]}{\sigma_a\sigma_b}. \tag{4.22}$$

The other features are based on spatial distance between signals. These are respectively the City block, Cosine, Euclidean, and the Jensen-Shannon, given by

$$d_{city}(a,b) = \sum |a[n]-b[n]|, \tag{4.23}$$

$$d_{cosi}(a,b) = 1-(a\cdot b)/(||a||_2||b||_2), \tag{4.24}$$

$$d_{eucl}(a,b) = ||a-b||_2, \tag{4.25}$$

$$d_{jens}(a,b) = \frac{1}{2}\sum_{w=1}^{W} D(A_w||M_w) + D(B_w||M_w) \tag{4.26}$$

where

$$D(A||B) = \sum p(A)\cdot \log\frac{p(A)}{p(B)} \tag{4.27}$$

in (4.26) is the Kullback-Leibler divergence and M_w is the point-wise average of A_w and B_w. The Jensen-Shannon divergence is calculated for multiple windows to limit the number of samples per calculation. Features (4.22-4.26) are used in the random forest.

4.2.3. Random Forest

A random forest classifier is a collection of decision trees which all make a separate classification decision based on the input features [60]. The final class is then determined by majority voting within the classifier. There are two key parameters - number of trees and their depth, in the random forest classifier. A visual representation of a Random Forest Classifier can be seen in Figure 4.9.

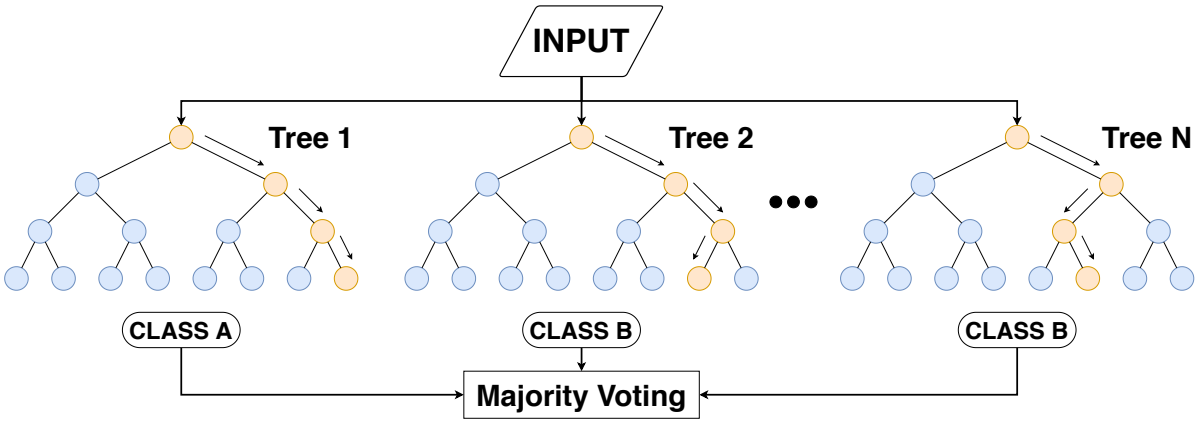

Figure 4.9: Representation of the functionality of a random forest.

5

Evaluation

5.1. Error statistics

5.1.1. People counting

We evaluate the proposed method through exhaustive simulations using the signal model in (3.4) and using experimental data. Influence of various model parameters on the performance of the method is evaluated in in section 5.2. The experimental evaluation is done in a cell office/meeting room environment in section 5.3 and characterises the performance of the method in a realistic indoor office environment.

Consider the error over a window of M samples,

$$\epsilon_{kj} = p_{kj} - \hat{p}_{kj}, \;\; j = 1, \cdots, M. \tag{5.1}$$

We consider an error metric by discarding a portion of the errors, specifically the largest and smallest errors. This choice is driven by application considerations - small latencies or errors do not matter (e.g. a workspace where a user is away for a few minutes' break need not be considered vacant). Define the trimmed average counting error as

$$\epsilon_k^{(M)} = \frac{1}{(1-2\delta)M} \sum_{j=\delta M}^{(1-\delta)M} \left|\epsilon_{k[j]}\right| \tag{5.2}$$

where $\epsilon_{k[j]}$ is the ordered sequence of error terms ϵ_{kj} within the k-th window. The average counting error (ACE) is then given by

$$ACE(M) = \frac{1}{K} \sum_{k=1}^{K} \epsilon_k^{(M)} \tag{5.3}$$

For $M = 1$, ACE(1) is simply the point-wise counting error. For other values of M, (5.3) provides an average error with the largest and smallest δ fraction of values removed; in the performance evaluation to follow, we choose $\delta = 0.1$. For instance, ACE(600) gives the average counting error for windows of duration 1 minute ignoring a few errors within this window. Such a performance metric better captures requirements for space management applications where a people counting value needs to be delivered every few minutes.

5.1.2. Commissioning change detection

In order to evaluate the performance of the proposed approach, we consider a commissioning change scenario wherein the commissioning plan changes from the one depicted in Figure 4.7 to Figure 4.8. For this commissioning plan detection scenario, we shall use the true positive rate (TPR) - ratio of the number of detected plan changes to the number of actual plan changes, and false positive rate (FPR) - ratio of the number of detected plan changes to the total number of no plan changes, as performance metrics.

5.2. People count simulations

We use the signal model in (3.4) under different occupancy patterns to generate simulated data to evaluate the impact of various model parameters on performance. Three parameters are of interest: (i) the base temperature, $f(\theta_{ij})\Delta T_i$, which depends on angle of object with respect to sensor, user orientation and the sensor height, (ii) the transition speed, α, of an event, and (iii) sensor noise level, under normality assumption.

The distribution of the base temperature change $f(\theta_{ij})\Delta T_i$ values are used to characterise the distributions based on which likelihood probabilities are computed and people count estimation done as per (4.11). In Figure 5.1, we see the performance based on this temperature change. The example is given for a zero to one transition event with varying values of base temperature change. If the base temperature change increase is too small, then it is difficult to detect an increase in people count leading to possible underestimation, and if the base temperature change is too large then an overestimation in people count is likely to occur. We thus observe that ACE first decreases with the base temperature change and then again increases.

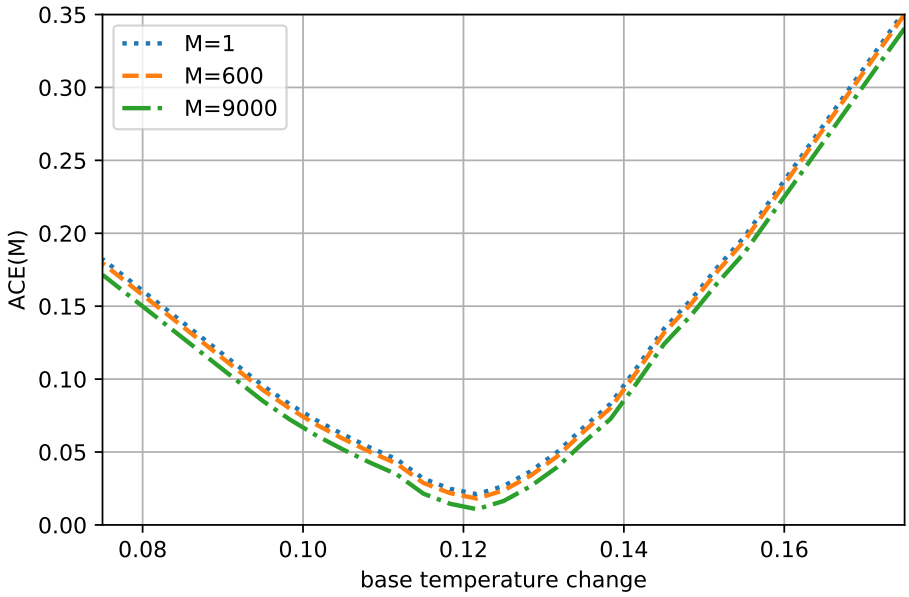

Figure 5.1: ACE vs base temperature change.

In Figure 5.2, we evaluate ACE over a range of transition speed, α, values. We see that ACE reduces for higher values of α. The transition speed determines how fast the temperature level increases or decreases under an occupancy event. A higher value of transition speed indicates that the temperature change under an occupancy event is faster. The transition speed affects the size of peaks in the score of CUSUM RLS processing. For slower transitions, the RLS has a smaller error with the current sample, resulting in smaller peaks. It was found that the performance for slower transitions could be increased by increasing the forgetting factor. Another affect found when dealing with slower transitions was the increase in underestimating the mean temperature change since the score could prematurely converge to zero.

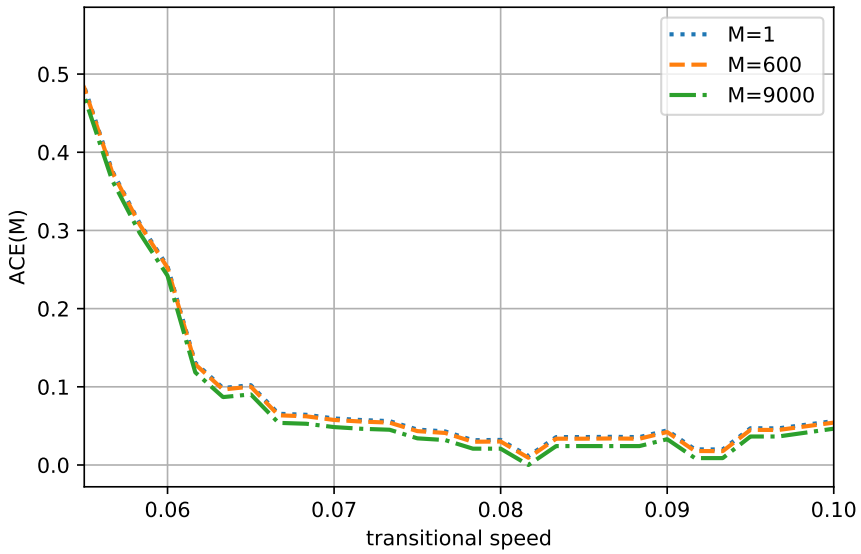

Figure 5.2: ACE vs transition speed.

The noise levels which were encountered from real sensor data measurements were found to have minimal effect on the performance of the algorithm as can be seen in Figure 5.3. The ACE is lower for higher values of M; errors for $M = 1$ are largely due to the inherent latency of the algorithm.

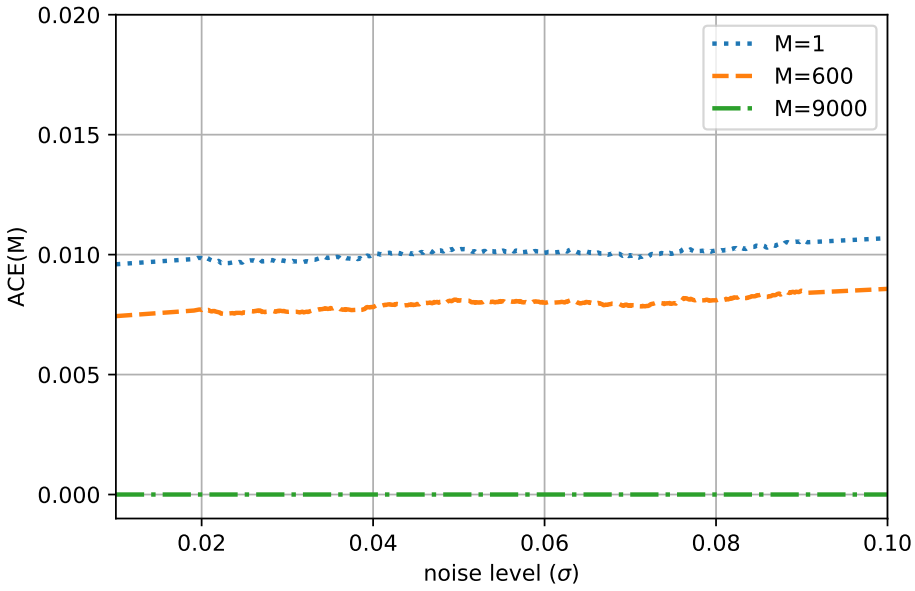

Figure 5.3: ACE vs noise level.

People counting estimation results of the example simulation in Figure 3.8 for a single day is shown in Figure Figure 5.4.

Lastly, we show the cumulative distribution function (cdf) of ACE using simulated data in Figure 5.5. The range for the transition speed is used between 0.07 and 0.1 and the range of the base temperature change is between 0.1 and 0.15. The noise level was chosen to be 0.05. From Figure 5.5, we see that with probability 0.9, ACE(1) = 0.23, ACE(600) = 0.19 and ACE(9000) = 0.11.

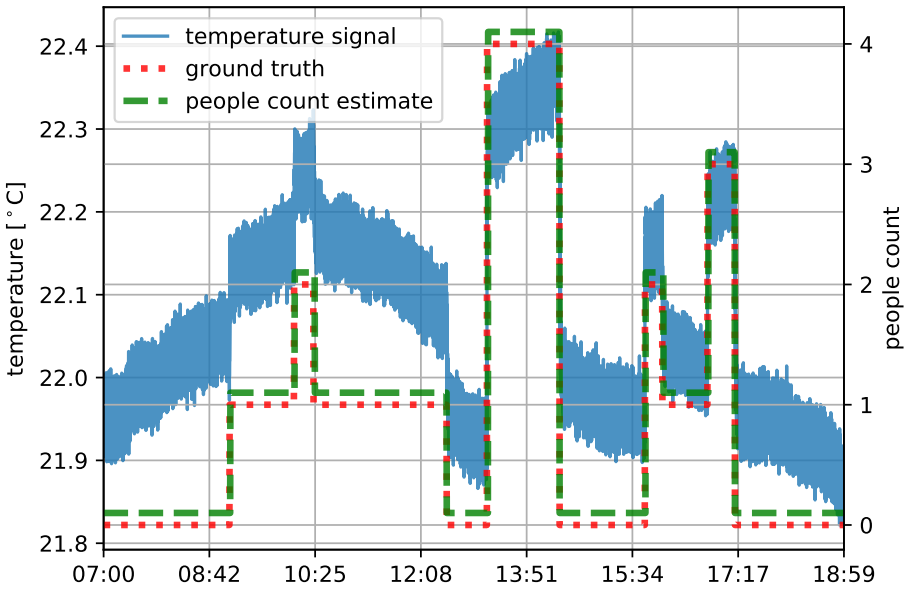

Figure 5.4: Results of a simulated dataset.

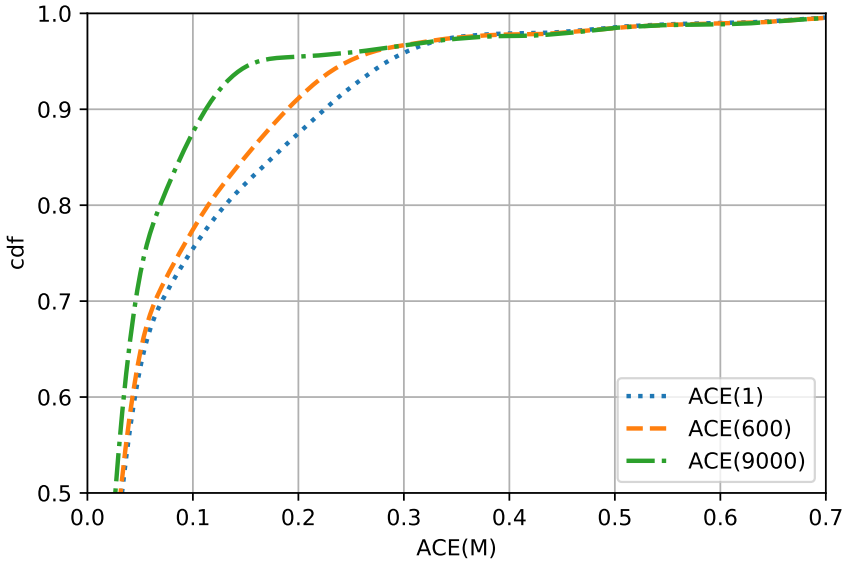

Figure 5.5: Cumulative distribution for the results of the simulation for parameters in the range of the most common observed values.

5.3. People count experiments

Experiments were conducted in a room emulating a cell office/meeting room environment, with the setup shown in Figure 5.6. There were 4 combi-sensors at the ceiling at a height of 3.2 meter from the floor. There were 4 workspaces within the FoV of the sensors.

During measurements, occupancy events were recorded manually each day and served as the ground truth. Measurements were also collected overnight to collect sufficient data under vacancy conditions. Data collection was done over three weeks, with diverse people count patterns. At maximum capacity, 4 people were present in the room.

Figure 5.6: Experimental setup of indoor space with four workspaces, with a zoomed view of the sensor setup at the ceiling.

Results of these experiments are shown in Figure 5.7. From this figure, we see that with probability 0.9, ACE(1) = 0.28, ACE(600) = 0.25, and ACE(9000) = 0.19. When increasing the window, the performance improves due to instantaneous errors being removed as well as temporal filtering.

An example of results from a dataset over a day is shown in Figure 5.8. The example shows a good estimation of the number of people and also depicts situations where the algorithm results in errors. As we can see, short events like getting coffee and small bathroom breaks are not affected by the algorithm. At the end, we do see an overestimation of such a short event.

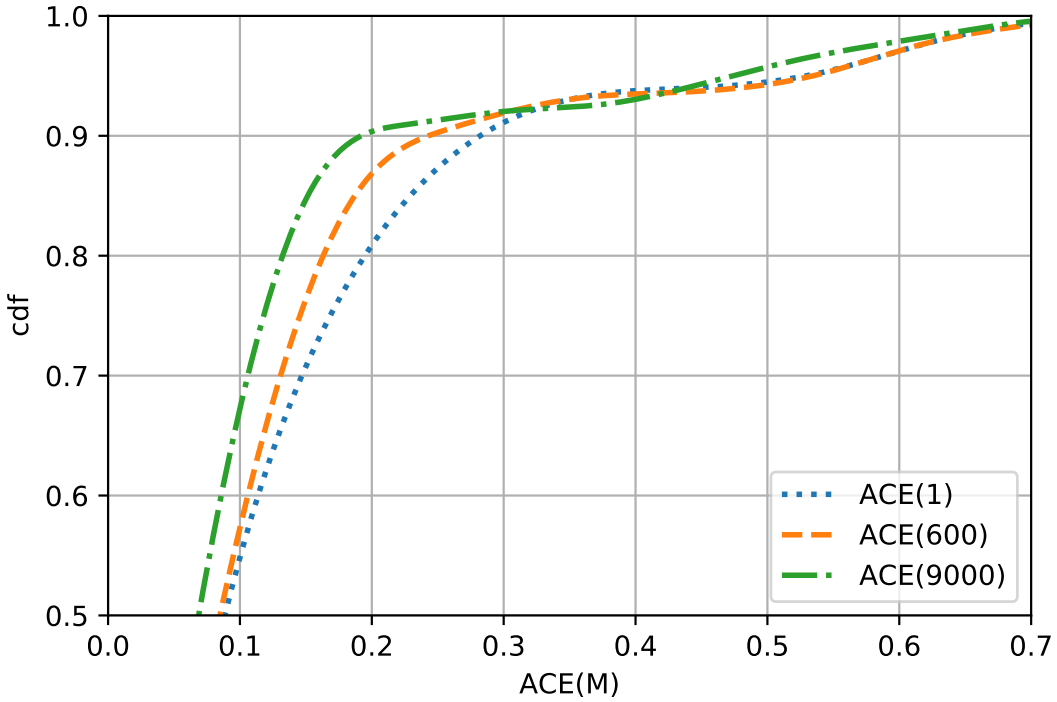

Figure 5.7: Cumulative distribution for the results of the real experiment results.

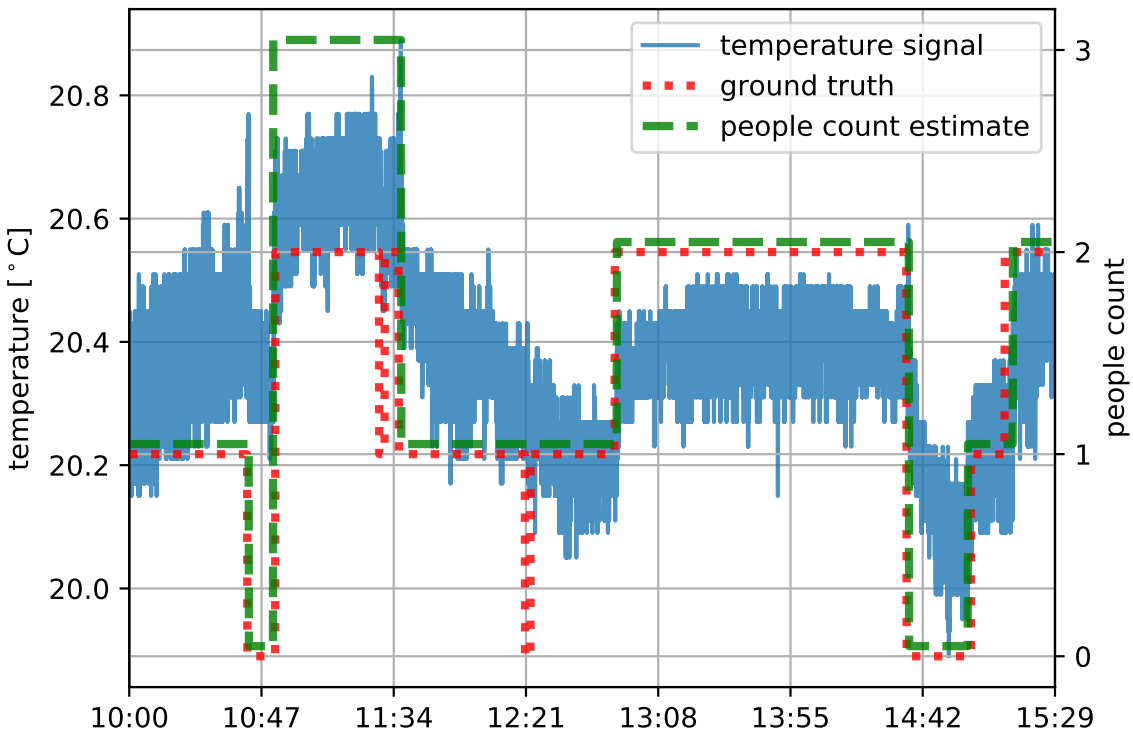

Figure 5.8: People count estimation of an experiment dataset.

5.4. Commissioning change simulations

We consider the setup shown in Figure 4.8 with occupancy generated over workspaces 1-8 over 10000 days, and a working time between 07:00 and 19:00. Sensor data was generated for sensors 1, 2 and $2'$ according to the model in (3.4), with a sampling frequency of 10 Hz. The noise was assumed to be AWGN with zero mean and a standard deviation of 0.04.

The attenuation due to sensor FoV is determined based on physical locations of user workspaces with respect to the sensors. Values for the transition speed α were chosen uniformly between 0.06 and 0.1 and for effective temperature change according to a normal distribution with mean 0.12 and standard deviation of 0.015. These values were determined based on experimental data collection. We use a simple occupancy model to emulate workspace occupancy over a day by generating a number of enter/leave events and the duration between events was chosen to be at least 5 minutes. The reader is referred to [61] for occupancy modelling in office buildings.

5.4.1. Random forest classifier parameters

We start by analysing the number of trees used in the Random Forest. In Figure 5.9, we see that the two curves start to converge after about 20 trees. From this point, the training data is perfectly fitted and adding more trees does not improve the performance of the test data. Since we have a binary classification and only 4 features selected. The required number of trees was not expected to be large. We set the number of trees to be 20.

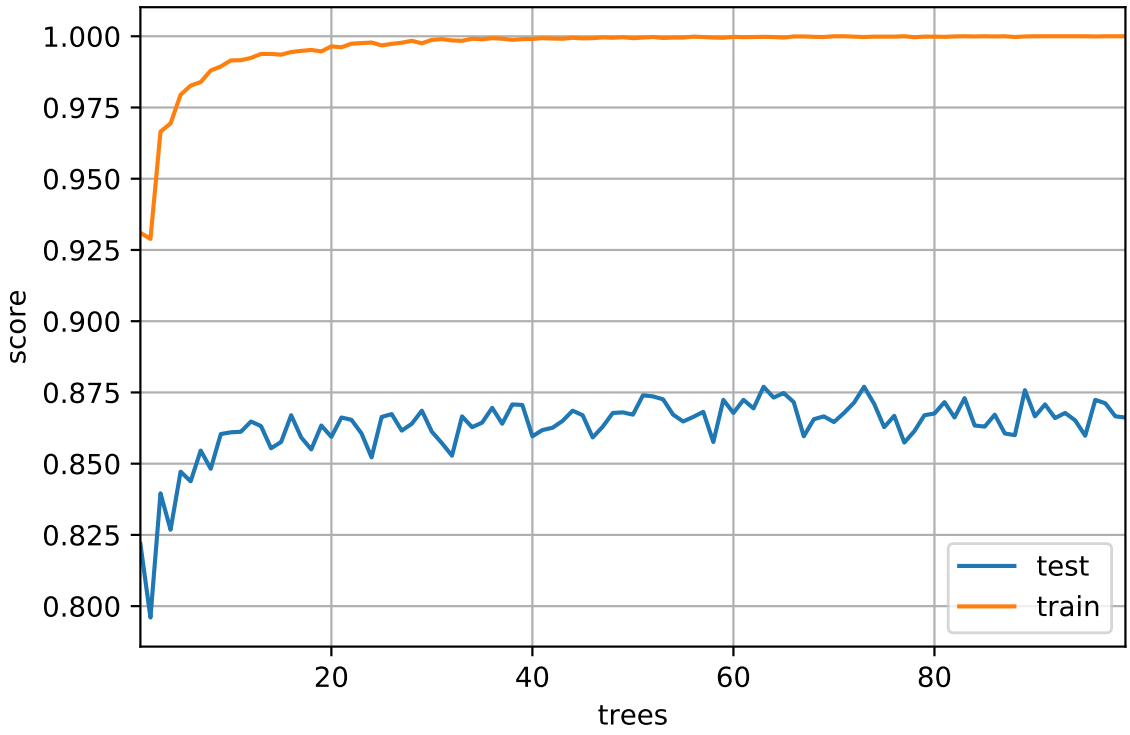

Figure 5.9: Response of the accuracy on the test and training data based on the number of trees in the Random Forest.

The second parameter to be tested is the maximum nodes per tree. In Figure 5.10, we see that the training set vastly increases in accuracy and approaches a perfect fitting of the training data. However, starting from 9 to 10 max depth, the accuracy of the test set is decreasing. This indicates that at a higher maximum node count, the Random Forest starts over-fitting the training data. We set the maximum number of nodes to be 9.

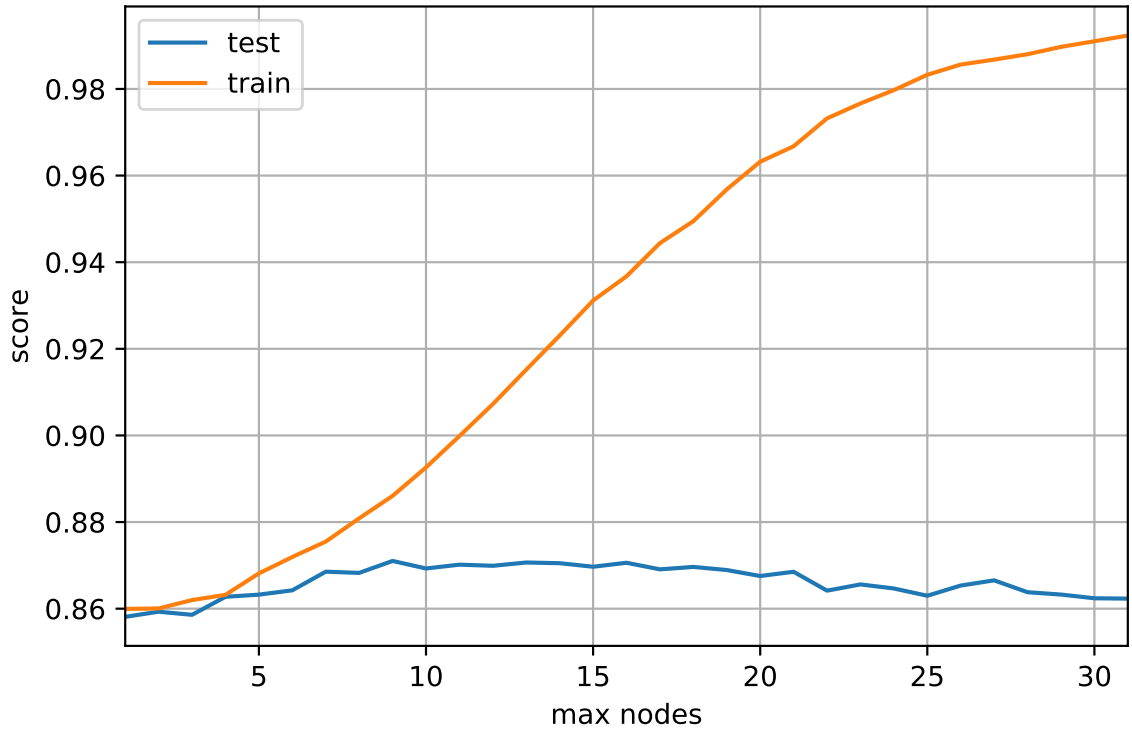

Figure 5.10: Response of the accuracy on the test and training data based on the maximum number of nodes per tree.

5.4.2. Performance results

The confusion matrix of the results are shown in Figure 5.11. From this confusion matrix, the performance does not seem overwhelmingly positive. The majority of the classifications is done correctly. From the confusion matrix, we find a TPR of 53.8% and 52.7% for the raw signal and CUSUM signals respectively. The FPR is at 46.1% and 47.2% respectively.

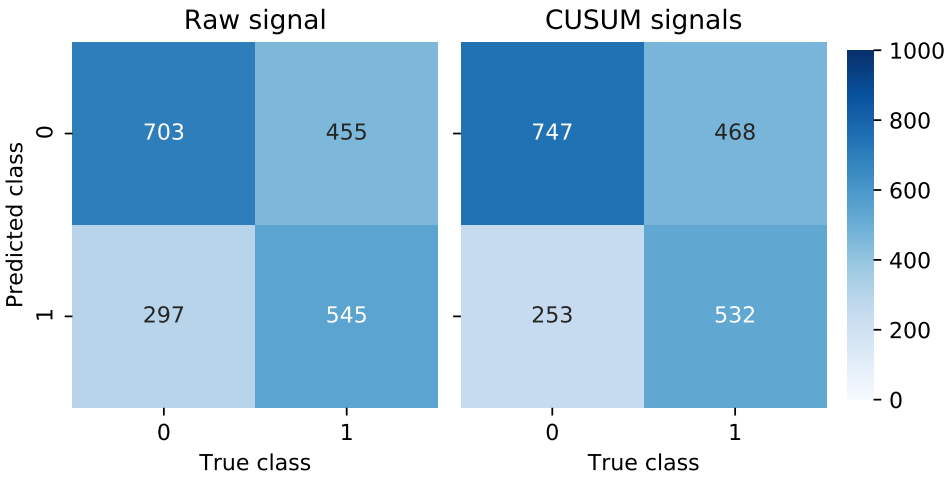

Figure 5.11: Confusion of matrix when using baseline features.

The feature importance spread seem almost uniformly divided in importance except for the CUSUM signal features where the minimum and maximum value features contribute less to the performance based on the positive or the negative side. In Figure 5.12, the importance spread is plotted for the raw signal and positive CUSUM signal.

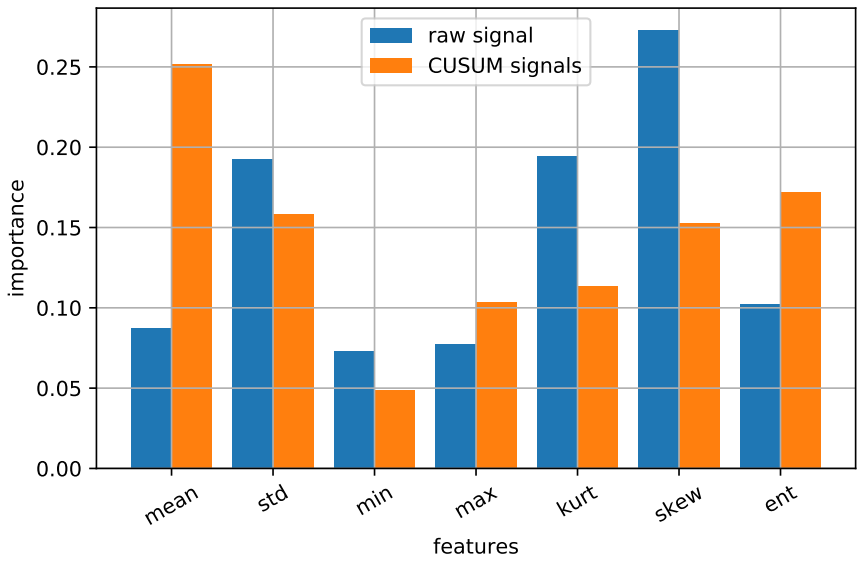

Figure 5.12: Importance of each baseline feature.

In Figure 5.13, we conclude a TPR of 74.0% and 90.2% for the raw signal and CUSUM signals respectively. The FPR is 8.6% and 1.3% respectively.

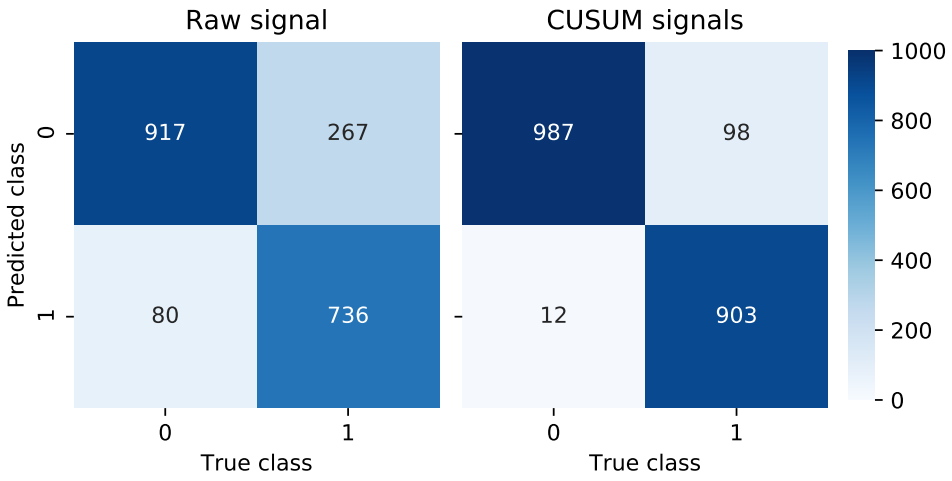

Figure 5.13: Confusion matrices with raw signal and CUSUM signal features using similarity features.

When using these features in Figure 5.14, the Pearson coefficient is the most important feature for both signals.

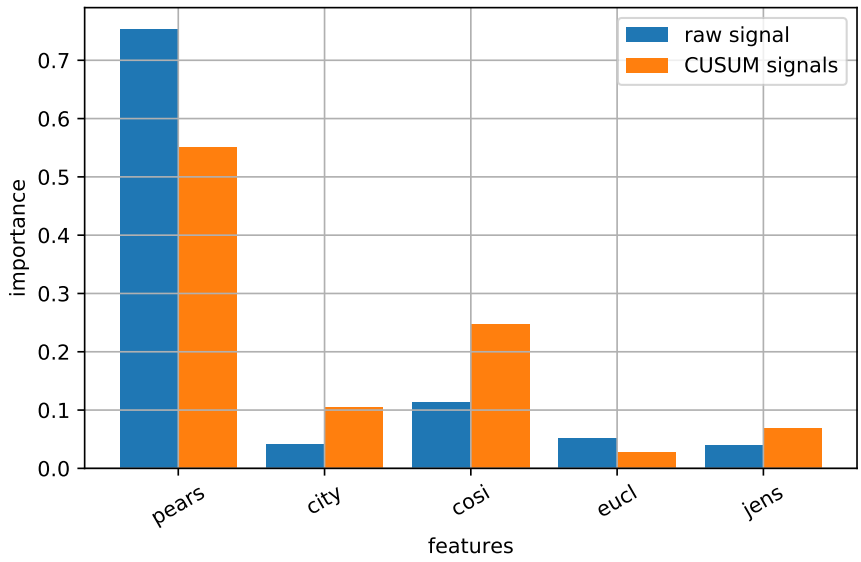

Figure 5.14: Importance of each similarity feature.

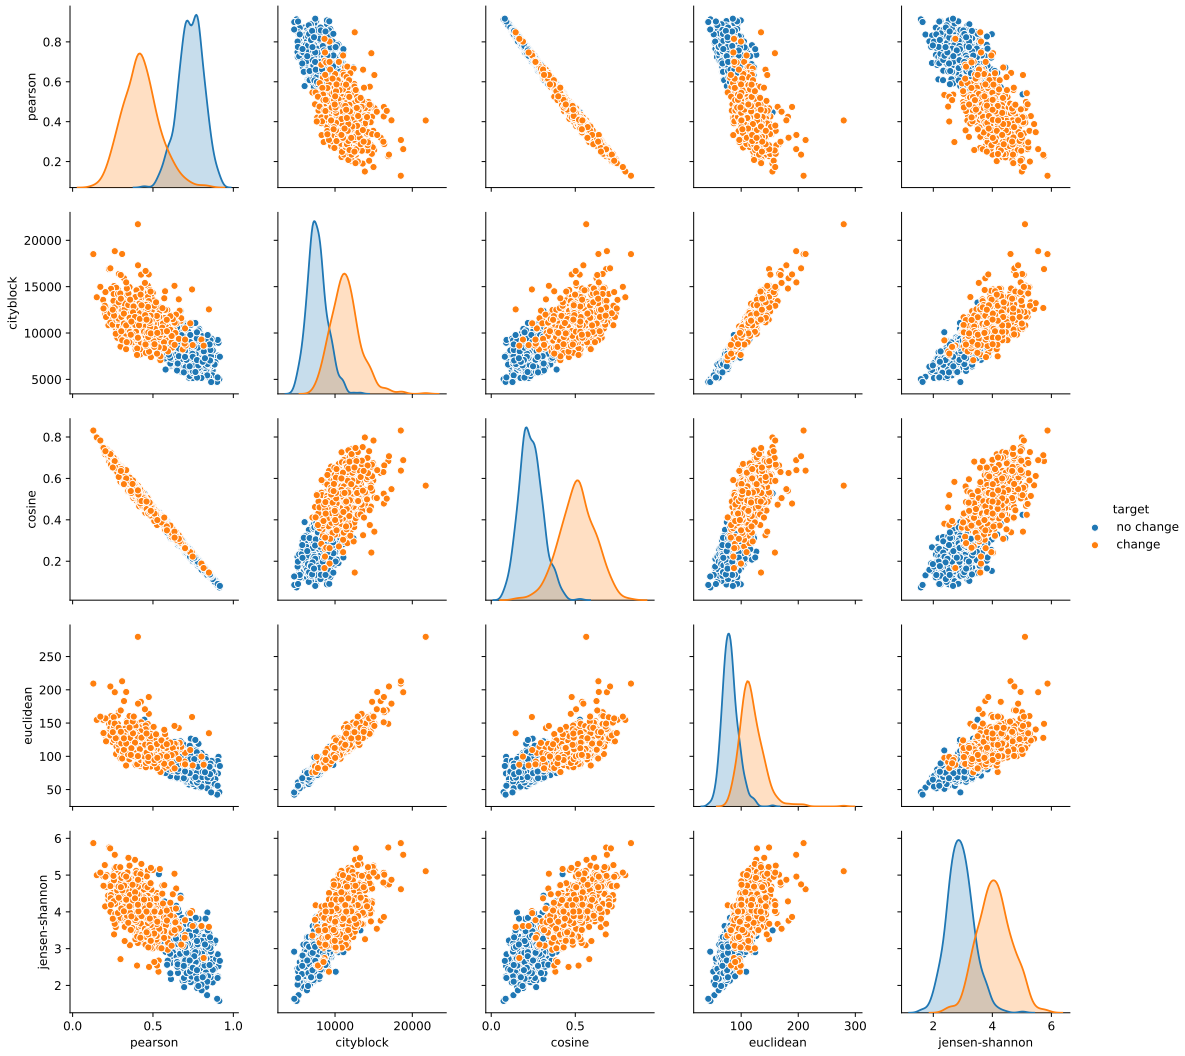

Figure 5.15: Pair-plot of the testing data. The kernel density estimation is plotted diagonally and a scatter plot per feature pair is shown outside the diagonal plots.

6

Conclusion

6.1. People counting

We showed that a single-pixel thermopile infrared sensing solution can provide people counting data with high accuracy. We extended the CUSUM-RLS filter by considering temperature change estimation, in turn using this to estimate people count. The proposed approach was shown to achieve an ACE of 0.11 and 0.19 for 90% of instants respectively when considering 15 minute windows for simulated and experimental datasets.

With the help of the statistical model, exhaustive parameter evaluation on three parameters was conducted. We find that performance increases for faster transitions. Experimental results also showed that small breaks (e.g., getting coffee or bathroom breaks) were mostly ignored by the algorithm preventing possible error propagation issues.

6.2. Commissioning change detection

We considered a random forest classifier to determine thermopile sensor commissioning in terms of their location plan. The classifier used various signal distance features for pairs of sensors. We proposed these features using CUSUM RLS processed object temperature signals from the thermopile sensors. Using simulated occupancy conditions and an analytical thermopile signal model, daily datasets were generated for evaluating the proposed approach. We obtained a TPR of 90.2% and an FPR of 1.3% using the proposed approach, in comparison to 53.2% and 25.3% when using basic statistical features.

By calculating 5 similarity features between the CUSUM signals of each sensor, a Random Forest Classifier was able to accurately detect a commissioning change. More importantly, the thesis showed a clear and relevant improvement when incorporating the CUSUM algorithm with respect to not using the CUSUM algorithm. When using the most basic signal statistics, an improvement of 3% accuracy was found probably due to the two-sided nature of the CUSUM algorithm. And more than 10% improvement was found when using relevant similarity features of the CUSUM signals instead of the raw signal.

7

Discussion

During the evaluation and run of the project, many ideas and questions were created. From the results of the People Count Estimation using the single-pixel thermopile sensors, we see that there is room for improvements. For example, the most sensitive model parameter for the performance of the algorithm is the base temperature change. This value highly depends on two factors: the position of the person and the orientation. Since the position is modelled, parametrical estimation could increase a correct estimation of the number of people from the temperature change estimation. When estimating the position in combination with the temperature change, a more accurate likelihood mapping is possible. Having estimated the positions allows for the localisation of work spaces as well.

Another interesting possible improvement for the involved people estimation is by implementing a dynamic distribution for calculating the likelihood. Since the probability densities should be different for different setups, by training the distribution per setup, the performance should increase. A method for dynamically calibrating and estimating this distribution would be required. Since we know from the PIR sensor when vacancy occurs and when activity is present, with the help of the number of events. We could fit the steps of temperature change to a fixed number of people and update the probability distribution.

Since we've only explored the possibilities of a single sensor, it is interesting to find if the performance could be increased when merging multiple overlapping sensors. Since the thermopile sensor is easily integratable into luminars, the option of sensor merging and performance improvements arrive.

When looking at the Commissioning Change Detection problem, we obtained high classification accuracy for simulated data. Of course, the next step would be to test the performance on real measured data. This however requires data availability for a large number of scenarios.

We found that the use of CUSUM RLS pre-processing led to a substantial improvement in classification for our problem. It would be worthwhile to generalize this approach for a larger class of problems wherein the information features lie in changes in a signal.

Bibliography

[1] C. Beyrouthy, E. K. Burke, D. L. Silva, B. McCollum, P. McMullan, and A. J. Parkes, “Towards improving the utilization of university teaching space”, *Journal of the Operational Research Society*, vol. 60, pp. 130–143, 2009.

[2] I. Cisco Systems, *How cisco designed the collaborative connected workplace environment*, 2007. [Online]. Available: https://www.cisco.com/c/dam/en_us/about/ciscoitatwork/downloads/ciscoitatwork/pdf/Cisco_IT_Case_Study_Connected_Workplace_POC.pdf.

[3] M. Apgar IV, “Managing real estate to build value”, *Harvard Business Review*, pp. 2–11, Nov. 1995.

[4] K. International, *The evolution in real estate: Insights from global real estate consulting case studies*, https://assets.kpmg.com/content/dam/kpmg/pdf/2016/07/the-evolution-in-real-estate.pdf, 2016.

[5] P. M. D. U.S. General Services Administration: Office of Real Property Management, *Workspace utilization and allocation benchmark*, https://www.gsa.gov/graphics/ogp/Workspace_Utilization_Banchmark_July_2012.pdf, 2012.

[6] P. Gupta and G. Parija, “Efficient seat utilization in global it delivery service systems”, in *2009 IEEE International Conference on Services Computing*, Sep. 2009, pp. 97–103.

[7] E. K. Burke, P. Cowling, and J. D. Landa Silva, “Hybrid population-based metaheuristic approaches for the space allocation problem”, in *Proceedings of the 2001 Congress on Evolutionary Computation (IEEE Cat. No.01TH8546)*, vol. 1, May 2001, 232–239 vol. 1. DOI: 10.1109/CEC.2001.934394.

[8] R. Pereira, K. Cummiskey, and R. Kincaid, “Office space allocation optimization”, in *2010 IEEE Systems and Information Engineering Design Symposium*, Apr. 2010, pp. 112–117. DOI: 10.1109/SIEDS.2010.5469670.

[9] A. B. Chan, Zhang-Sheng John Liang, and N. Vasconcelos, “Privacy preserving crowd monitoring: Counting people without people models or tracking”, in *2008 IEEE Conference on Computer Vision and Pattern Recognition*, Jun. 2008, pp. 1–7. DOI: 10.1109/CVPR.2008.4587569.

[10] M. Eldib, F. Deboeverie, W. Philips, and H. Aghajan, “Towards more efficient use of office space”, in *Proceedings of the 10th International Conference on Distributed Smart Camera*, ser. ICDSC ’16, Paris, France: ACM, 2016, pp. 37–43, ISBN: 978-1-4503-4786-0. DOI: 10.1145/2967413.2967424. [Online]. Available: http://doi.acm.org/10.1145/2967413.2967424.

[11] T. Teixeira and A. Savvides, “Lightweight people counting and localizing for easily deployable indoors wsns”, *IEEE Journal of Selected Topics in Signal Processing*, vol. 2, no. 4, pp. 493–502, Aug. 2008. DOI: 10.1109/JSTSP.2008.2001426.

[12] M. Eldib, N. B. Bo, F. Deboeverie, J. Nino, J. Guan, S. Van de Velde, H. Steendam, H. Aghajan, and W. Philips, “A low resolution multi-camera system for person tracking”, in *2014 IEEE International Conference on Image Processing (ICIP)*, Oct. 2014, pp. 378–382. DOI: 10.1109/ICIP.2014.7025075.

[13] D. Caicedo and A. Pandharipande, “Ultrasonic arrays for localized presence sensing”, *IEEE Sensors Journal*, vol. 12, no. 5, pp. 849–858, May 2012.

[14] ——, “Distributed ultrasonic zoned presence sensing system”, *IEEE Sensors Journal*, vol. 14, no. 1, pp. 234–243, Jan. 2014.

[15] K. Modepalli and L. Parsa, “Dual-purpose offline led driver for illumination and visible light communication”, *IEEE Transactions on Industry Applications*, vol. 51, no. 1, pp. 406–419, Jan. 2015. DOI: 10.1109/TIA.2014.2330066.

[16] P. M. Djuric, M. Vemula, and M. F. Bugallo, “Target tracking by particle filtering in binary sensor networks”, *IEEE Transactions on Signal Processing*, vol. 56, no. 6, pp. 2229–2238, Jun. 2008. DOI: 10.1109/TSP.2007.916140.

[17] V. L. Erickson, A. Beltran, D. A. Winkler, N. P. Esfahani, J. R. Lusby, and A. E. Cerpa, “Thermosense: Thermal array sensor networks in building management”, in *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems*, ser. SenSys ’13, Roma, Italy: ACM, 2013, 87:1–87:2, ISBN: 978-1-4503-2027-6. DOI: 10.1145/2517351.2517437. [Online]. Available: http://doi.acm.org/10.1145/2517351.2517437.

[18] J. Tanaka, H. Imamoto, T. Seki, and M. Oba, “Low power wireless human detector utilizing thermopile infrared array sensor”, in *SENSORS, 2014 IEEE*, Nov. 2014, pp. 462–465. DOI: 10.1109/ICSENS.2014.6985035.

[19] S. Kumar, T. K. Marks, and M. Jones, “Improving person tracking using an inexpensive thermal infrared sensor”, in *2014 IEEE Conference on Computer Vision and Pattern Recognition Workshops*, Jun. 2014, pp. 217–224. DOI: 10.1109/CVPRW.2014.41.

[20] M. Kuki, H. Nakajima, N. Tsuchiya, and Y. Hata, “Multi-human locating in real environment by thermal sensor”, in *2013 IEEE International Conference on Systems, Man, and Cybernetics*, Oct. 2013, pp. 4623–4628. DOI: 10.1109/SMC.2013.787.

[21] A. Tyndall, R. Cardell-Oliver, and A. Keating, “Occupancy estimation using a low-pixel count thermal imager”, *IEEE Sensors Journal*, vol. 16, no. 10, pp. 3784–3791, May 2016. DOI: 10.1109/JSEN.2016.2530824.

[22] H. M. Ng, “Poster abstract: Human localization and activity detection using thermopile sensors”, in *2013 ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, Apr. 2013, pp. 337–338. DOI: 10.1145/2461381.2461436.

[23] K. A. Agha, I. G. Lassous, and G. Pujolle, Eds., *Challenges in Ad Hoc Networking*. Springer US, 2006. DOI: 10.1007/0-387-31173-4. [Online]. Available: https://doi.org/10.1007/0-387-31173-4.

[24] A. Pandharipande, D. Caicedo, and X. Wang, “Sensor-driven wireless lighting control: System solutions and services for intelligent buildings”, *IEEE Sensors Journal*, vol. 14, no. 12, pp. 4207–4215, Dec. 2014, ISSN: 1530-437X. DOI: 10.1109/JSEN.2014.2351775.

[25] A. Kumar, A. Singh, A. Kumar, M. K. Singh, P. Mahanta, and S. C. Mukhopadhyay, “Sensing Technologies for Monitoring Intelligent Buildings: A Review”, *IEEE Sensors Journal*, vol. 18, no. 12, pp. 4847–4860, Jun. 2018, ISSN: 2379-9153. DOI: 10.1109/JSEN.2018.2829268.

[26] A. Pandharipande, M. Zhao, and E. Frimout, “Connected indoor lighting based applications in a building IoT ecosystem”, *IEEE Internet of Things Magazine*, vol. 2, no. 1, pp. 22–26, Mar. 2019. DOI: 10.1109/IOTM.2019.1900016.

[27] A. Aryal, B. Becerik-Gerber, F. Anselmo, S. C. Roll, and G. M. Lucas, “Smart desks to promote comfort, health, and productivity in offices: A vision for future workplaces”, *Frontiers in Built Environment*, vol. 5, Jun. 2019. DOI: 10.3389/fbuil.2019.00076. [Online]. Available: https://doi.org/10.3389/fbuil.2019.00076.

[28] G. Violatto, A. Pandharipande, S. Li, and L. Schenato, “Classification of occupancy sensor anomalies in connected indoor lighting systems”, *IEEE Internet of Things Journal*, vol. 6, no. 4, pp. 7175–7182, Aug. 2019, ISSN: 2372-2541. DOI: 10.1109/JIOT.2019.2914937.

[29] C. R. Wren, A. Azarbayejani, T. Darrell, and A. P. Pentland, “Pfinder: Real-time tracking of the human body”, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 19, no. 7, pp. 780–785, Jul. 1997. DOI: 10.1109/34.598236.

[30] K. A. Shiva Kumar, K. R. Ramakrishnan, and G. N. Rathna, “Distributed person of interest tracking in camera networks”, in *Proceedings of the 11th International Conference on Distributed Smart Cameras*, ser. ICDSC 2017, Stanford, CA, USA: Association for Computing Machinery, 2017, pp. 131–137, ISBN: 9781450354875. DOI: 10.1145/3131885.3131921. [Online]. Available: https://doi.org/10.1145/3131885.3131921.

[31] ——, “Inter-camera person tracking in non-overlapping networks: Re-identification protocol and on-line update”, in *Proceedings of the 11th International Conference on Distributed Smart Cameras*, ser. ICDSC 2017, Stanford, CA, USA: Association for Computing Machinery, 2017, pp. 55–62, ISBN: 9781450354875. DOI: 10.1145/3131885.3131912. [Online]. Available: https://doi.org/10.1145/3131885.3131912.

[32] V. Kumar, *Promoting Social Change and Democracy through Information Technology*. IGI Global, Jul. 2015, ISBN: 9781466685024. [Online]. Available: https://www.xarg.org/ref/a/1466685026/.

[33] S. Singhal, C. Neustaedter, T. Schiphorst, A. Tang, A. Patra, and R. Pan, “You are being watched: Bystanders’ perspective on the use of camera devices in public spaces”, in *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, ser. CHI EA ’16, San Jose, California, USA: ACM, 2016, pp. 3197–3203, ISBN: 978-1-4503-4082-3. DOI: 10.1145/2851581.2892522. [Online]. Available: http://doi.acm.org/10.1145/2851581.2892522.

[34] K. Wolf, A. Schmidt, A. Bexheti, and M. Langheinrich, “Lifelogging: You’re wearing a camera?”, *IEEE Pervasive Computing*, vol. 13, no. 3, pp. 8–12, Jul. 2014, ISSN: 1558-2590. DOI: 10.1109/MPRV.2014.53.

[35] S. Weckmann, “Dynamic electrothermal model of a sputtered thermopile thermal radiation detectorfor earth radiation budget applications”, Master’s thesis, Virginia Polytechnic Institute and State University, Sep. 1997, ch. 2. [Online]. Available: http://scholar.lib.vt.edu/theses/available/etd-8497-205315/.

[36] S. Aminikhanghahi and D. Cook, “A survey of methods for time series change point detection”, *Knowledge and Information Systems*, vol. 51, Sep. 2016. DOI: 10.1007/s10115-016-0987-z.

[37] F. Gustafsson, *Adaptive Filtering and Change Detection*. John Wiley & Sons, Ltd, Oct. 2001. DOI: 10.1002/0470841613. [Online]. Available: https://doi.org/10.1002/0470841613.

[38] E. Hagenaars and M. Berkers, *Smart sensors and communication using iot in supermarkets*, http://resolver.tudelft.nl/uuid:d3a4a0ac-a965-4fef-a4a4-628772c54432, Jul. 2017.

[39] T. den Boer and B. de Vos, *Smart sensors and communication using internet of things in supermarkets*, http://resolver.tudelft.nl/uuid:d8ccd59f-ca8b-4f6a-beea-02df1a454f88, Jul. 2017.

[40] F. Bovolo, L. Bruzzone, and M. Marconcini, “A novel approach to unsupervised change detection based on a semisupervised svm and a similarity measure”, *IEEE Transactions on Geoscience and Remote Sensing*, vol. 46, no. 7, pp. 2070–2082, Jul. 2008, ISSN: 1558-0644. DOI: 10.1109/TGRS.2008.916643.

[41] M. Jahari, S. Khairunniza-Bejo, A. R. M. Shariff, H. Z. M. Shafri, and H. Ibrahim, “Change detection using a local similarity measure”, in *2008 IEEE Conference on Innovative Technologies in Intelligent Systems and Industrial Applications*, Jul. 2008, pp. 39–43. DOI: 10.1109/CITISIA.2008.4607332.

[42] J. .-.-. Tourneret, V. Poulain, M. Chabert, and J. Inglada, “Similarity measure between vector data bases and optical images for change detection”, in *2009 IEEE International Geoscience and Remote Sensing Symposium*, vol. 2, Jul. 2009, pp. II-992-II–995. DOI: 10.1109/IGARSS.2009.5418268.

[43] S. Cui, M. Datcu, and L. Gueguen, “Information theoretical similarity measure for change detection”, in *2011 Joint Urban Remote Sensing Event*, Apr. 2011, pp. 69–72. DOI: 10.1109/JURSE.2011.5764721.

[44] J. Zhao, J. Yang, Z. Lu, P. Li, and W. Liu, “Change detection based on similarity measure and joint classification for polarimetric sar images”, in *2017 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, Jul. 2017, pp. 1896–1899. DOI: 10.1109/IGARSS.2017.8127348.

[45] M. Aktar, M. A. Mamun, and M. A. Hossain, “Statistical similarity based change detection for multitemporal remote sensing images”, *Journal of Electrical and Computer Engineering*, vol. 2017, pp. 1–8, 2017. DOI: 10.1155/2017/3123967. [Online]. Available: https://doi.org/10.1155/2017/3123967.

[46] F. Di, X. Li, and C. Zhu, “A new method in change detection of remote sensing image”, in *2009 2nd International Congress on Image and Signal Processing*, Oct. 2009, pp. 1–4. DOI: 10.1109/CISP.2009.5304661.

[47] H. Yang, J. Gao, C. Xu, Z. Long, W. Feng, S. Xiong, S. Liu, and S. Tan, “Infrared image change detection of substation equipment in power system using random forest”, in *2017 13th International Conference on Natural Computation, Fuzzy Systems and Knowledge Discovery (ICNC-FSKD)*, Apr. 2017, pp. 1745–1751. DOI: 10.1109/FSKD.2017.8393030.

[48] C. Huo, Y. Zhang, J. Yu, Y. Ling, and C. Pan, “Learning deep relationship for image change detection”, in *IGARSS 2018 - 2018 IEEE International Geoscience and Remote Sensing Symposium*, Apr. 2018, pp. 1918–1921. DOI: 10.1109/IGARSS.2018.8518693.

[49] L. Xu, W. Jing, H. Song, and G. Chen, “High-resolution remote sensing image change detection combined with pixel-level and object-level”, *IEEE Access*, vol. 7, pp. 78 909–78 918, 2019, ISSN: 2169-3536. DOI: 10.1109/ACCESS.2019.2922839.

[50] M. Botsch and J. A. Nossek, “Feature selection for change detection in multivariate time-series”, in *2007 IEEE Symposium on Computational Intelligence and Data Mining*, Mar. 2007, pp. 590–597. DOI: 10.1109/CIDM.2007.368929.

[51] S. Li, A. Pandharipande, B. Masini, and D. Caicedo, “Automated detection of commissioning changes in connected lighting systems”, *IEEE Internet of Things Journal*, vol. 6, no. 1, pp. 898–905, Feb. 2019, ISSN: 2372-2541. DOI: 10.1109/JIOT.2018.2865315.

[52] Melexis MLX90614 datasheet, https://www.melexis.com/-/media/files/documents/datasheets/mlx90614-datasheet-melexis.pdf, Accessed Feb 1, 2020.

[53] R. B. D’Agostino, “An omnibus test of normality for moderate and large size samples”, *Biometrika*, vol. 58, no. 2, pp. 341–348, 1971, ISSN: 00063444. [Online]. Available: http://www.jstor.org/stable/2334522.

[54] R. D’Agostino and E. S. Pearson, “Tests for departure from normality. empirical results for the distributions of b2 and sqrt(b1)”, *Biometrika*, vol. 60, no. 3, pp. 613–622, 1973, ISSN: 00063444. [Online]. Available: http://www.jstor.org/stable/2335012.

[55] B. Iglewicz, *How to Detect and Handle Outliers (Asqc Basic References in Quality Control)*. Asq Pr, Dec. 1993, ISBN: 087389247X. [Online]. Available: https://www.xarg.org/ref/a/087389247X/.

[56] M. Duarte and R. Watanabe, *Notes on scientific computing for biomechanics and motor control*, https://github.com/BMClab/BMC, 2018.

[57] W. Shen, G. Newsham, and B. Gunay, “Leveraging existing occupancy-related data for optimal control of commercial office buildings: A review”, *Advanced Engineering Informatics*, vol. 33, pp. 230–242, 2017, ISSN: 1474-0346. DOI: https://doi.org/10.1016/j.aei.2016.12.008. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1474034616301987.

[58] D. Michie, D. J. Spiegelhalter, C. C. Taylor, and J. Campbell, Eds., *Machine Learning, Neural and Statistical Classification*. USA: Ellis Horwood, 1995, ISBN: 013106360X.

[59] S.-H. Cha, “Comprehensive survey on distance/similarity measures between probability density functions”, *International Journal of Mathematical Models and Methods in Applied Sciences*, vol. 1, no. 4, pp. 300–307, 2007.

[60] L. Breiman, “Random forests”, *Machine Learning*, vol. 45, no. 1, pp. 5–32, Oct. 2001, ISSN: 1573-0565. DOI: 10.1023/A:1010933404324. [Online]. Available: https://doi.org/10.1023/A:1010933404324.

[61] C. Liao and P. Barooah, “An integrated approach to occupancy modeling and estimation in commercial buildings”, in *Proceedings of the 2010 American Control Conference*, Jun. 2010, pp. 3130–3135. DOI: 10.1109/ACC.2010.5531035.