CountIR: Occupancy Estimation using a Low-Pixel Count Thermal Imager

Author: Ash Tyndall (20915779)

School of Computer Science and Software Engineering

Supervisors: Rachel Cardell-Oliver, Adrian Keating

Introduction

Energy prices are rising. Occupancy detection, the ability to count people, offers one way to increase energy efficiency and save money, as knowing the number of people in a space allows more efficient heating and cooling.

Aim

We aimed to make an occupancy detection system that is;

- Low Cost: Less than $300 in prototype stage.
- Non-Invasive: Collects minimal data to perform goal.
- Reliable: At least 75% accurate.
- Energy Efficient: Can last more than 7 days on battery.

“ThermoSense,” a recent paper, used a low resolution thermal sensor to count occupants with acceptable accuracy. We based our algorithmic approach off this paper.

Prototype

A hardware prototype was constructed with a three-tiered architecture to investigate these criteria, costing $185. We used a sensor which has a narrower, rectangular field of view when compared with ThermoSense’s sensor.

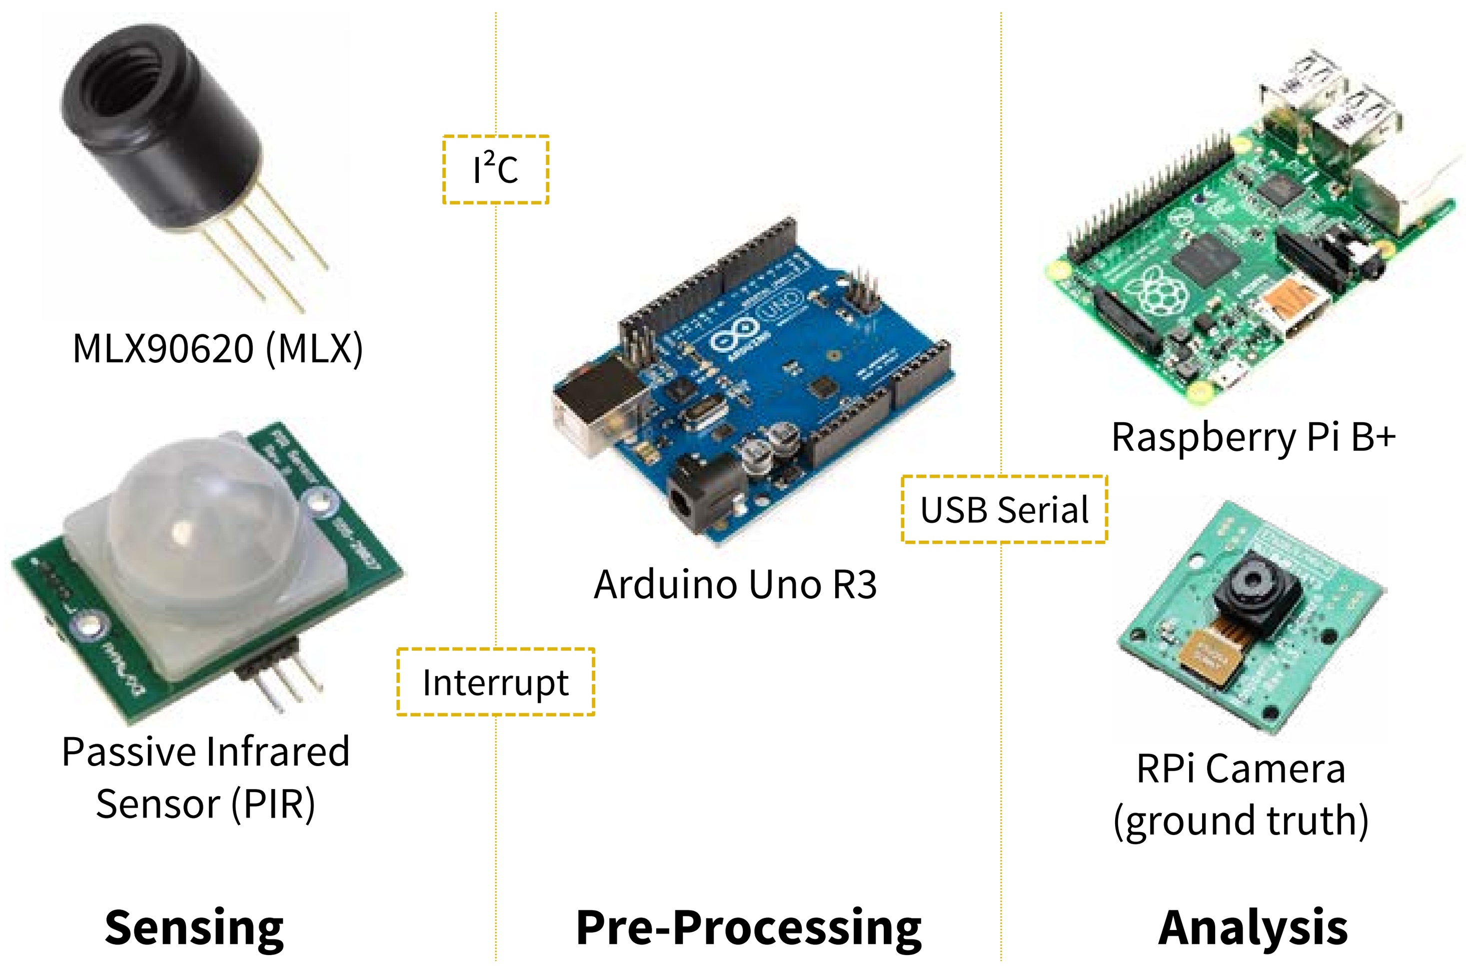

Prototype Architecture

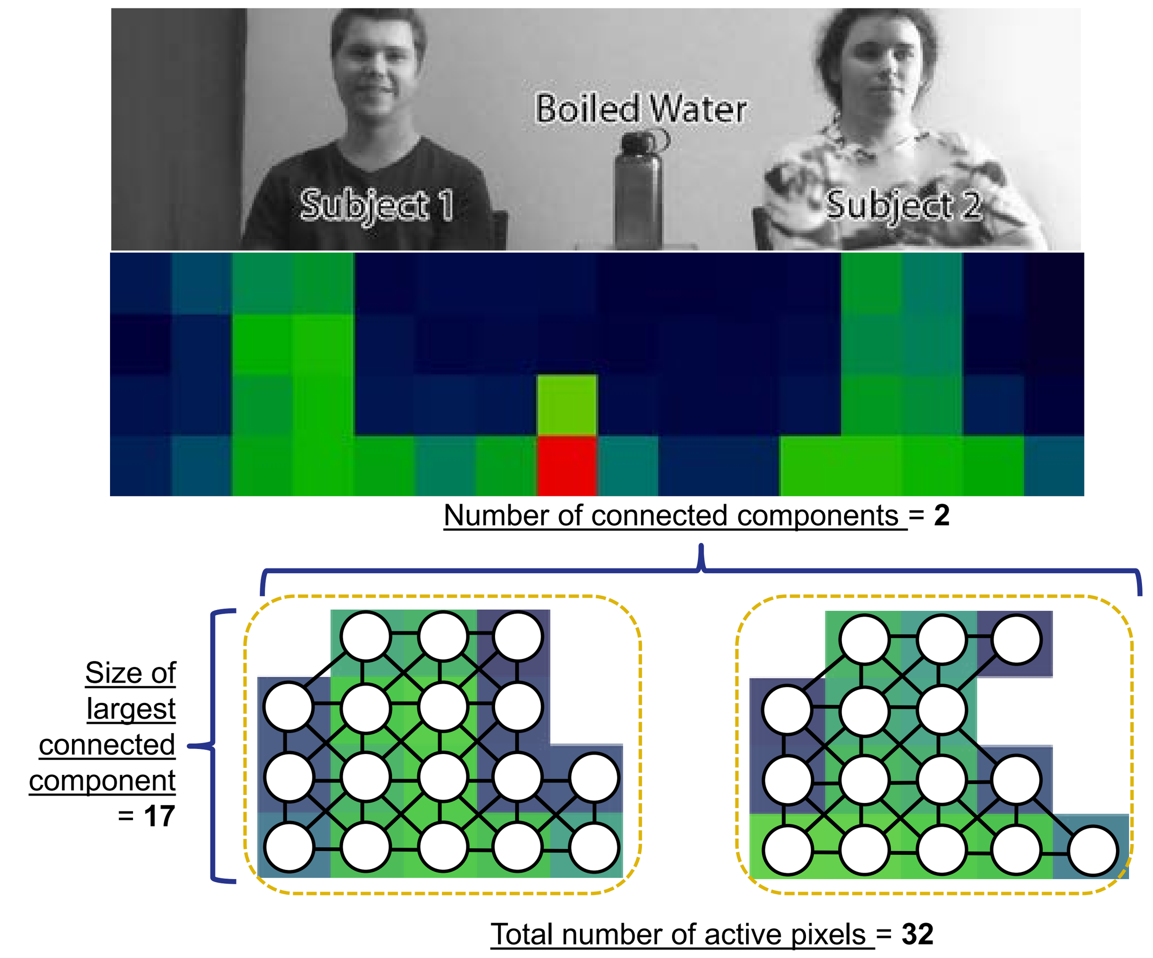

Taking a raw thermal image and extracting its three features

Algorithm

1. When no motion is detected, use thermal frames to calculate a background mean and standard deviation.
2. When motion is detected, any pixel >3 standard deviations above the mean is considered “active.”
3. Convert “active” pixels to a graph, using adjacency as edges, and extract three feature vectors:
   - Total number of active pixels.
   - Number of connected components. (pixel “islands”)
   - Size of largest connected component.

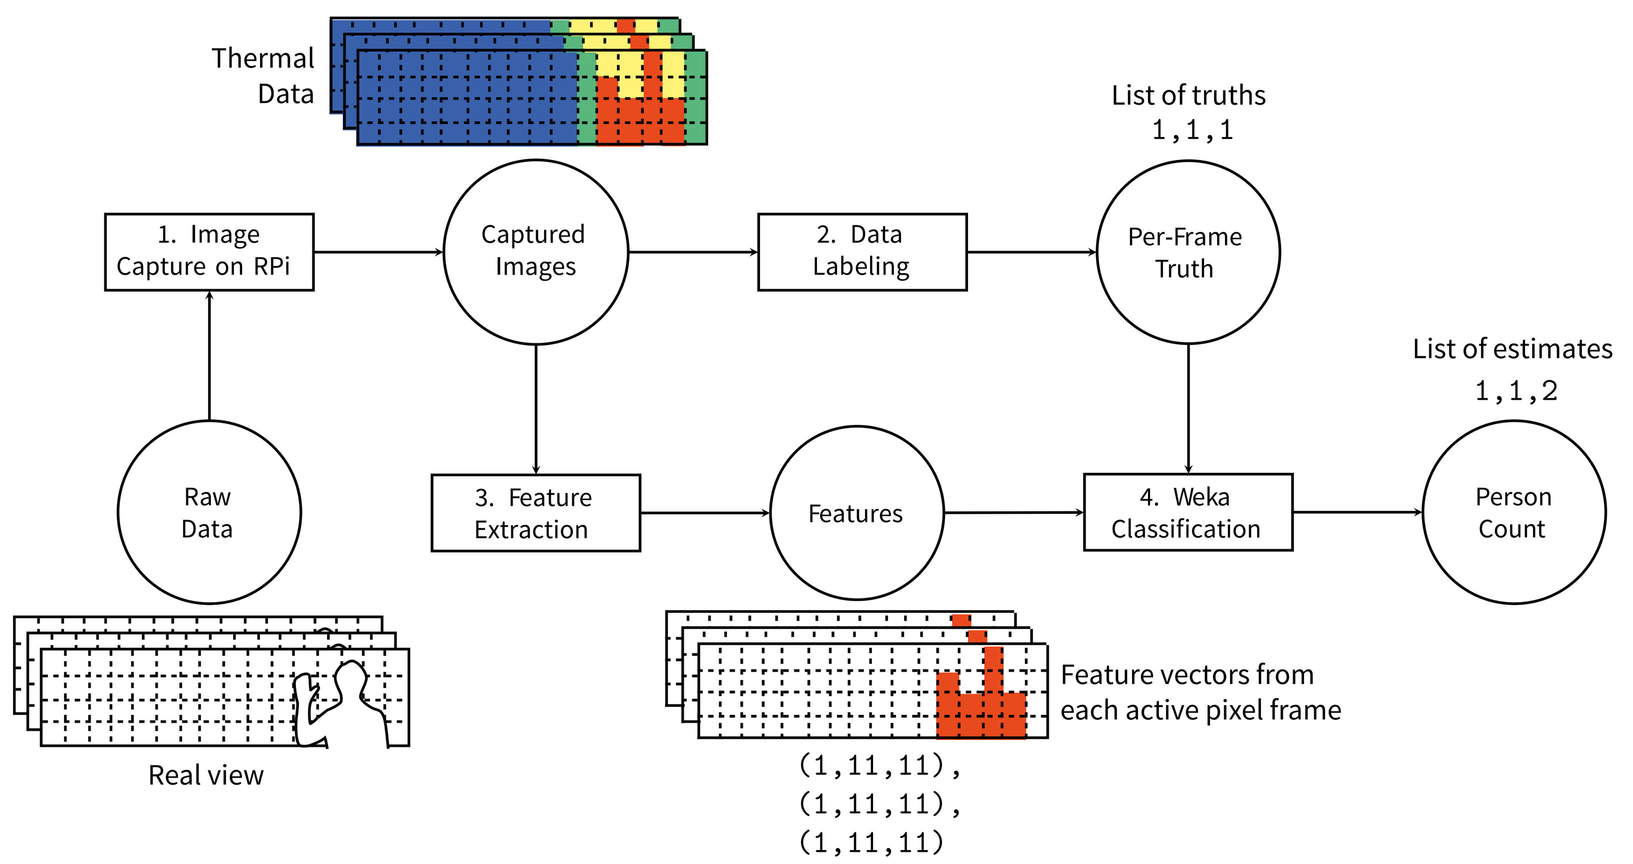

Processing Pipeline

Evaluation

To evaluate our reliability and energy efficiency criteria, a series of experiments were required. To that end, we created a software processing pipeline for converting raw data into occupancy predictions.

We collected ~1,000 frames of data with corresponding ground truth using a roof-mounted prototype with a floor-facing sensor. We imported this data into Weka to run a series of machine learning classification experiments.

The following algorithms were trailed:

- *Numeric Multi-Layer Perceptron*
- *Linear Regression*
- *K-Nearest Neighbours*
- Nominal Multi-Layer Perceptron
- K*
- C4.5
- Support Vector Machine
- Naive Bayes

Algorithms in *italics* were used in the ThermoSense paper.

We found that our best performing techniques were C4.5 and K*, which were both ~82% accurate. We were unable to replicate the high accuracies of ThermoSense’s classifiers, suggesting differences in our experimental setups.

When capturing, the prototype drew ~255 mW, suggesting it would last 8 days with a 50 Wh battery. With modifications, draw could theoretically be reduced to only ~0.4 mW, increasing life to several years.

Conclusions

- ☑ Low Cost? $185 is sufficiently low, with the price of all components trending downwards in the future.
- ☑ Non-Invasive? Low resolution thermal imaging gathers sufficiently little information to be non-invasive.
- ☑ Reliable? 82% accuracy is sufficient for our purposes.
- ☑ Energy Efficient? Current prototype predicted to last 8 days. With modifications, it could last several years.

Beltran, A., Erickson, V. L., and Cerpa, A. E. ThermoSense: Occupancy thermal based sensing for HVAC control. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings* (2013), ACM, pp. 1–8.