

# Longitudinal Energy Waste Detection with Visualization

David Lachut<sup>†</sup>, Nilavra Pathak<sup>‡</sup>, Nilanjan Banerjee<sup>†</sup>, Nirmalya Roy<sup>‡</sup>, Ryan Robucci<sup>\*</sup>

<sup>†</sup>Computer Science, <sup>‡</sup>Information Systems, <sup>\*</sup>Computer Engineering

University of Maryland, Baltimore County

dlachut1, nilavra1, nilanb, nroy, robucci@umbc.edu

## ABSTRACT

Leaky windows and doors, open refrigerators, unattended appliances, left-on lights, and other sources subtly leak energy accounting for a large portion of waste. Formal energy audits are expensive and time consuming and do not capture many sources of leakage and waste. In this short paper, we present a hybrid IR/RGB imaging system for an end-user to deploy to perform longitudinal detection of energy waste. The system uses a low resolution,  $16 \times 4$  IR camera and a low cost digital camera mounted on a steerable platform to automatically scan a room, periodically taking low resolution IR and RGB images. The system uses image stitching to create an IR/RGB hybrid panoramic image and segmentation to determine temperature extrema in the scanned room. Finally, this data is combined with thermostat set-point information to highlight hot-spots or cold-spots which likely indicate energy leakage or wastage. The system obviates the need for expensive, time-consuming waste detection methods, for professional setup, and for more intrusive instrumentation of the home.

## CCS CONCEPTS

• **Computer systems organization** → **Embedded systems**; *Redundancy*; Robotics; • **Networks** → Network reliability;

### ACM Reference format:

David Lachut<sup>†</sup>, Nilavra Pathak<sup>‡</sup>, Nilanjan Banerjee<sup>†</sup>, Nirmalya Roy<sup>‡</sup>, Ryan Robucci<sup>\*</sup>. 2017. Longitudinal Energy Waste Detection with Visualization. In *Proceedings of ACM BuildSys conference, Delft, The Netherlands, November 2017 (BuildSys '17)*, 5 pages. [https://doi.org/10.475/123\\_4](https://doi.org/10.475/123_4)

## 1 INTRODUCTION

Estimates say 75,000 US homes in one year waste an amount of energy equivalent to the energy loss due to the British Petroleum oil spill calamity [1]. A portion of this energy wastage can be attributed to callous use of appliances—open refrigerator doors, open microwave doors, televisions and computers switched on when there is no one in the room, and open doors and windows. The other portion can be attributed to poor wall insulation as well as drafty doors and windows. While poor insulation can be detected using energy audits, it is more challenging to determine wastage

due to careless usage of appliances and opening or closing of windows/door. Continuous monitoring using a sensing system that is low-cost and stays out of a user's way is an ideal method for discovering this waste.

### 1.1 Related Work

Direct energy metering is a plausible seeming solution to the problem of waste detection. Every appliance could be instrumented with an energy meter, and analytics on the energy consumption data could determine appliance (mis-)use. However, blanketing a home with such energy meters can be prohibitively expensive and highly intrusive [15]. Energy disaggregation [4] on data from a single point measurement system is another plausible solution, but its accuracy is limited if the number of appliances is large and the sampling rate is low. To address this problem, we present the design and preliminary evaluation of a low-cost thermal imaging system that uses a single low resolution IR and RGB camera to determine energy consumption hot-spots in a room, motivated by techniques mentioned in [3, 7].

Energy disaggregation with data fusion from indirect sensing has shown to improve accuracy and convey more knowledge regarding user behavior. Acoustic sensing has been used to improve appliance identification in [12]. The applicability of using user WiFi data to localize appliance usages is discussed in [13]. In a similar vein location and activity has been used in [14] to reduce the search space of disaggregation from contextual information.

The authors of [8] approach the problem by using a steerable camera system superficially similar to ours, but which requires the users to create or obtain a 3D CAD model of his home's interior space. Our approach is distinguished from this and from the above approaches to waste detection by our typically lower cost, lower intrusiveness, and lower amounts of user-facing complexity.

### 1.2 Contributions

This project is part of an overall effort to reduce the number of people who require a full, time-consuming, expensive, professional energy audit. In particular, this project seeks to detect energy leakage at a coarser level than the detailed professional service while providing the additional benefit of scanning for behavioral, rather than structural, energy waste.

Our system employs an IR and a RGB camera mounted on a steerable base. The system periodically scans a room, each scan panning the cameras through a complete revolution while taking images at regular intervals. The system then uses a combination of image stitching and segmentation to discover temperature hot-spots and cold-spots in the room. By overlaying and correlating panoramic views of the thermal and visual images, the system

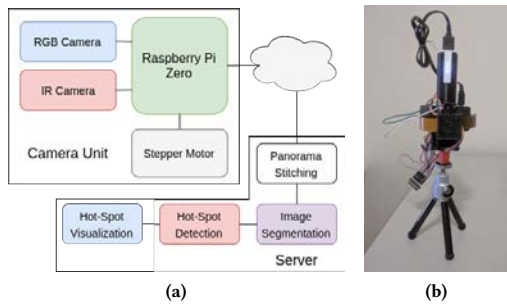
Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

*BuildSys '17, November 2017, Delft, The Netherlands*

© 2017 Copyright held by the owner/author(s).

ACM ISBN 123-4567-24-567/08/06...\$15.00

[https://doi.org/10.475/123\\_4](https://doi.org/10.475/123_4)



**Figure 1: (a) Block diagram of the energy waste detection system. (b) Camera module prototype viewed from the front. One can see the battery on top, the pair of cameras in the middle, and the tripod base below.**

provides a visualizations of the hot-/cold-spots. Using multiple scans, the system could correlate these hot-spots over time, and so reveal consistent patterns of wastage. Through a preliminary evaluation using an experimental setup in the lab and in a home, we show that the system can accurately find these hot-spots.

The design, implementation, and evaluation of this system presents the following research contributions:

- **A user-deployed energy audit system.** We present thermal imaging system that can discover likely sources of a user’s energy waste. It obviates the need to instrument appliances with individual energy meters. We develop an image processing algorithm that can reliably detect temperature hot-spots in the room. This provides the ability to collect longitudinal data on energy wastage of homes.
- **Functional prototype and evaluation.** We have developed a prototype system and present preliminary evaluation showing that we can detect likely sources of energy waste in a home environment.

## 2 SYSTEM DESCRIPTION

Figure 1a gives an overview of the system’s organization. The system has two major components—Camera Module and the Server component. The Camera Module consists of the IR camera, RGB camera, and the stepper motor and is pictured in Figure 1b. The three tools are controlled by a Raspberry Pi. The motor rotates the module by a set number of steps. Then the two cameras collect simultaneous images. The rotate-and-image loop continues until the module has panned one full rotation. After the set of images is collected, the Raspberry Pi uploads them to the server and waits until time for another scan. Once on the server, the images are pre-processed, stitched into panoramas, and segmented. The server then examines each segment looking for hot-/cold-spots and produces an isolated image of any such spot found.

### 2.1 Camera Module

The first component of our system is the hybrid IR/RGB Camera Module for longitudinal thermal scanning of an entire room which helps to find the potential sources of energy leak or waste. While commercially available systems to perform the scanning exist, they are expensive and warrant professionals to operate. Motivated by this, we integrate a relatively low-cost, low-power IR camera

(MLX90621) [6] with a Raspberry Pi to build a simple prototype which can be deployed at ease by the home owners, residential and commercial occupants at their own comfort. The total BoM for the prototype costs \$125 placing it well within the range of current smarthome products. While one camera module can only assess one room at a time, it is less expensive than many consumer smarthome solutions that would provide similar functionality for that one room and is far simpler to move to another room.

The IR camera measures the ambient temperature with a range of  $-40^{\circ}\text{C}$  to  $85^{\circ}\text{C}$  with a resolution of  $0.02^{\circ}\text{C}$  and object temperatures in a range of  $-50^{\circ}\text{C}$  to  $300^{\circ}\text{C}$ . These measurements are the bases for our thermal maps of a room. We simultaneously use a visible-spectrum camera to collect RGB images of the scanned area, which images are correlated with the IR data to produce the thermal maps of the room. The two cameras are kept in constant relative alignment by a 3D printed housing.

To achieve a  $360^{\circ}$  pan for generating panoramas, we integrate a stepper motor into the camera module. The 3D printed housing—containing the controller, battery, and both cameras—sits atop the stepper motor. A standard camera tripod holds the motor and housing steady. Previously, we had used a simpler analog DC motor, but the stepper motor allows us to precisely control the angle of rotation between each simultaneous IR and RGB image. In each scanning period, the device gathers 64 simultaneous IR/RGB image pairs, rotating  $360^{\circ}$  in  $5.625^{\circ}$  increments. Once a scanning period is complete, the device uploads the collected images to the server for processing.

### 2.2 Image Processing

Once uploaded to the server, the image data goes through six different steps:

*Preprocessing.* The first step turns the raw data into useful information. The RGB camera captures images in portrait orientation to maximize the vertical resolution of the horizontal panorama, so the server corrects this orientation. The IR camera captures readings as a flat array of raw byte values, which the server combines with values collected from that camera’s EEPROM to create a  $16 \times 4$  matrix of temperature readings in  $^{\circ}\text{F}$ . We then map the range of temperatures in the 64-image set to the range of whole numbers from 0 to 255, producing a set of low-resolution, greyscale thermal images.

Because the 3D-printed housing maintains the relation between the two cameras, we were able to experimentally determine a consistent homography between the two kinds of images. Using this, we project each thermal image on to the space of the corresponding RGB image. This warping is imperfect, because the cameras do have different fields of view, requiring us to crop out all areas where the images do not overlap. The high number of images we collect results in sufficient overlap to mitigate the potential information loss.

*Panorama Stitching.* The panorama stitching procedure is based on [9, 10]. We first perform a cylindrical warp on all the images. Then, because we use a stepper motor, we know the relative angular distance from each image to the next. Trigonometry converts this into a pixel distance between each image. Finally, we use a pyramid blending procedure to create smooth thermal and visual panoramas.

*Image Segmentation.* Once we have panoramas, we adapt [11] to take a hybrid approach to the segmentation task. Unlike typical segmentation through watershedding, we incorporate data from both the IR panorama and the RGB panorama. Watershedding is a way of separating an image into non-overlapping segments based on gradients between pixel values. The process depends heavily on the number and locations of the markers, or starting points. We first find the markers for our watershed starting points in the thermal panorama, effectively basing the locations on detected thermal zones. Taking these starting points, we then perform the watershedding on the visual image. This allows us to fuse the two different kinds of sensor data and account for both in the detection process.

*Hot-Spot Detection.* Each of these segments is then evaluated based on its temperature. Recall in the preprocessing step, the temperature scale was remapped to byte values. Taking the thermal panorama, we can do the reverse procedure and get a good-enough approximation of the real temperatures. We then take the mean and standard deviation of the temperatures as seen in the panorama. For each segment, we calculate the segment’s mean temperature. If segment’s mean temperature is different from the overall mean by more than one standard deviation, we mark the segment as a hot-spot or a cold spot. Our process then filters these segments in two ways. First, we ignore segments below a minimum size, these are invariably noise and not useful detections. Second, if the room’s thermostat target is known, we can optionally filter a class of detections. This second filter allows the system to avoid alerting on cold AC vents in the summer, or hot vents in the winter.

*Hot-Spot Isolation.* The next step in image processing is to isolate the detected hot-/cold-spots. We produce for the user a panorama overlaying the color blue on cold spots and red on hot spots. Further, we create cropped images isolating the segment of the panorama containing the potential source of waste.

*Image Matching.* The final step in the image processing uses longitudinal data to confirm sources of energy leakage and waste. The visualization step creates cropped images isolating potential leaks. If these are meaningful sources of waste, they will appear multiple times in the record as the system collects longitudinal data. By using computer vision to match keypoints in these images, we can detect when the same hot-/cold-spot appears too frequently in the record and bring that specific source to the user’s attention. As a longitudinal component, this depends on multiple scans of the same room, and the system must be retrained if moved to a different room.

*Alert Delivery.* As more alert images match to previous alert images over time, the probability increases that these images have all detected the same device wasting energy. The system should then select one of these images and present it to the user as an indication that the object pictured is likely wasting energy, having been left on—as a heat-generating appliance—or open—as a defective refrigerator door—or allowed to leak—as a drafty window. The user will then know where she can look for energy savings.

### 3 EVALUATION

To begin our evaluation, we asked a volunteer to test the camera at his home. We deployed it, setting the module to collect and



(a) Example visual panorama.



(b) Example panorama with thermal data overlaid.




(c) Example panorama with detected hot-spot segments colored in red.

**Figure 2: A generated panorama of the system successfully detecting a bright window and an unattended computer keeping the room at 90°F while the thermostat is set to 77°F.**

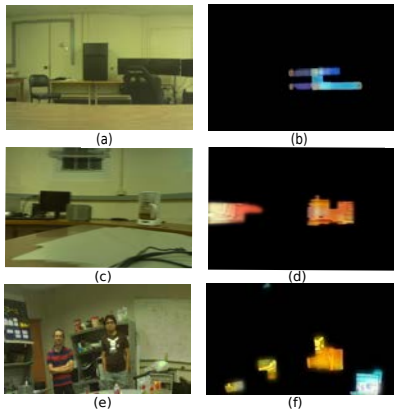


**Figure 3: Two separate examples of the longitudinal detection system detecting an unattended computer. The top two images were isolated from their visual panoramas and identified as images of the same object via 125 matched key-points shown in green.**



**Figure 4: Longitudinal temperature of window, compared to outdoor and mean room temperature.**

upload image data every 15 minutes the device was active. Data was collected for approximately 25 total hours over the course of a few days. We wanted to see if our data collection and image processing pipeline could detect potential energy leaks in an actual residential environment. While this is only a preliminary evaluation, the results are promising.



**Figure 5: Preliminary appliance and object recognition applied to a refrigerator (a, b), a coffee maker and a monitor (c, d), and humans with an open refrigerator (e, f).**

Figure 2 shows examples of the image panoramas our system creates. On top is the visual panorama created from only the RGB images. In the middle we present that same visual panorama overlaid with a colorized version of the thermal panorama. Temperatures in the thermal image range from  $80^{\circ}F$  to over  $100^{\circ}F$  with warmer colors representing warmer measured temperatures. These two panoramas were produced by *Panorama stitching* and are used in *Image Segmentation* as described in 2.2. In the third panorama, the system has completed *Image Segmentation* drawn a red mark over each segment found during *Hot-Spot Detection*. One sees the system alerted on a window that may need resealing or better shade and on the computer which appears to have been left on and unattended.

This was not the only time our subject left the computer on, unattended, wasting energy, and heating up his already hot room. Figure 3 shows two *Hot-Spot Isolation* images created from processing different scans of the room. That image processing step, described in 2.2, isolated and stored a record of the unattended computer in two separate scans. By *Image Matching* those alerts, our matching system identified 125 visual keypoint matches, highlighted in green in the third image of Figure 3, between the two images. The high number of keypoint matches indicates that both alerts are for the same appliance. Altogether, this longitudinal information suggests the user's unattended computer is a source of energy waste.

Longitudinal data can likewise rule-out things that would appear from a single scan to be leakages. Figure 4 plots the measured temperature over time of the window on which the system alerted in Figure 2, comparing this temperature to the mean temperature of the room and the outdoor temperature. The window alert in Figure 2 was generated upon that first sweep where the window temperature is especially high. Subsequent data shows that alert to likely be atypical, ruling out the Window as a source of waste or leakage.

#### 4 NEXT STEPS

Though the capability was not ready by time of publication, we have begun work on a presence detection and object recognition stage in processing, which will allow us to determine room occupancy and content. Visualized in Figure 5, the method is based on

comparing the thermal measurements to known surface temperatures of various appliances and pets. Once complete, this will give the system access to semantics about the objects it recognizes and allow it to pass that information to the user.

A  $98.6^{\circ}F$  human body is a particularly challenging object to detect as it is both a thermal outlier in a typical room, causing false positive alerts, and it comes in varieties of morphology and apparel that can blend its parts into the background. Our longitudinal approach to data collection will help as users' movements relative to the room help them stand out. Human detection is especially important as it is primarily humans who interact with appliances and so will allow us to further filter false positives from our scan results as humans in a room can be presumed to be interacting with an appliance. After integrating this process with the system, we will scale up our system deployments to do a comprehensive evaluation and system refinement. This refinement should include doing as much computation on the device as possible, to improve users' privacy protection.

#### 5 CONCLUSION

We presented a thermal imaging system prototype, which is capable of scanning an entire room and identifying potential sources of energy leakage and waste. The system is relatively user-friendly and non-intrusive, as it obviates the need to blanket a room with energy meters. The system is also cost-effective, being less expensive than a professional energy audit by combining an essential hardware device with back-end software to extract information from data through sensor fusion and computer vision. The system's design to remain in place and collect data over time should allow it to give the user unique insights into their energy use or misuse.

#### REFERENCES

- [1] EIA 2014 [http://www.eia.gov/electricity/sales\\_revenue\\_price/pdf/table5\\_a.pdf](http://www.eia.gov/electricity/sales_revenue_price/pdf/table5_a.pdf), 2014.
- [2] Matthew Brown and David G. Lowe. 2007. Automatic Panoramic Image Stitching using Invariant Features. *Int. J. Comput. Vision* 74, 1 (August 2007), 59-73.
- [3] P. X. Gao and S. Keshav. SPOT: A Smart Personalized Office Thermal Control System. In *Proc. ACM e-Energy* 2013, 2013.
- [4] Kolter, J. Zico, Siddharth Batra, and Andrew Y. Ng. Energy disaggregation via discriminative sparse coding. *Advances in Neural Information Processing Systems*. 2010.
- [5] M. Sezgin and B. Sankur. Survey over image thresholding techniques and quantitative performance evaluation. *Journal of Electronic Imaging* 13, 2004.
- [6] [http://www.mouser.com/ds/2/734/LX90621\\_product\\_flyer-706890.pdf](http://www.mouser.com/ds/2/734/LX90621_product_flyer-706890.pdf)
- [7] <http://spectrum.ieee.org/geek-life/hands-on/ir-eye/>
- [8] Rao, Vaishakh, and Anthony Rowe. Low-cost continuous thermal sensing. *Third ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*. ACM, 2011.
- [9] R. Szeliski. *Image Alignment and Stitching: A Tutorial*. Technical Report MSR-TR-2004-92. Microsoft Research, 2006.
- [10] [http://docs.opencv.org/3.1.0/dc/dff/tutorial\\_py\\_pyramids.html](http://docs.opencv.org/3.1.0/dc/dff/tutorial_py_pyramids.html)
- [11] [https://opencv-python-tutorials.readthedocs.io/en/latest/py\\_tutorials/py\\_imgproc/py\\_watershed/py\\_watershed.html](https://opencv-python-tutorials.readthedocs.io/en/latest/py_tutorials/py_imgproc/py_watershed/py_watershed.html)
- [12] N. Pathak, M. A. A. H. Khan, and N. Roy. Acoustic based appliance state identifications for fine-grained energy analytics. *PerCom*, 2015.
- [13] M. Saha, S. Thakur, A. Singh, and Y. Agarwal. Energylens: combining smartphones with electricity meter for accurate activity detection and user annotation. *ACM E-Energy* 2014.
- [14] Nirmalya Roy, Nilavra Pathak, Archan Misra, Fine-grained appliance usage and energy monitoring through mobile and power-line sensing, *Pervasive and Mobile Computing*, Volume 30, August 2016.
- [15] Lachut, David, et al. "Minimizing intrusiveness in home energy measurement." *Proceedings of the Fourth ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings\**. ACM, 2012.