

Fast, Automated Indoor Light Detection, Classification, and Measurement

Craig Hiller, Avidah Zakhori; University of California, Berkeley; Berkeley, California

Abstract

Lighting is one of the largest power consumers in the United States and around the globe. To better understand how much energy lighting uses in a building, a lighting audit can be performed. Typically, this is a long and manual process, current solutions require significant effort on the part of the auditor. This paper develops a system using commercially available hardware and custom algorithms that enable a single human operator to quickly cover a large area while estimating light positions, type, and surface area. These tasks are accomplished with an error rate of 6.9% and 13.9%, respectively, with surface area estimation within about a factor of two.

Introduction

Building lighting is one of the largest consumers of electricity around the globe. In California, over twenty five percent of commercial energy consumption is used just for illumination [1]. To pave the way for a greener future, a building's energy usage and therefore lighting usage, must first be understood. Currently, a lighting audit is performed by an auditor who surveys an area and manually classifies, counts, and marks the lights on the building's floor plan [2]. This is a long and error-prone process which can take over a day to complete for a large building [3]. To circumvent this, some auditors install energy monitoring devices to monitor power consumption when the lights are on or off. This installation procedure requires invasive changes to the building's electrical system.

Prior research in automated object detection for energy auditing includes window [4] and computer [5] detection. In these, a human carried and operated backpack system [6] with multiple LiDAR sensors and cameras is used to generate localized 3D points clouds. The objects are then detected within the scenes using infrared and visible light images, surface normal maps, and depth information.

Unmanned aerial vehicles, UAVs, have been experimented with for the light auditing task. Bay and Terrill et. al. affix a camera, spectrometer, and distance sensor to autonomously detect and classify lights within a building [7, 8]. Light sizes are estimated by using the distance to the floor in an area with a known fixed height ceiling. Using their approach, global position of lights cannot be estimated. This paper utilizes similar methods for raw image processing and light spectra classification.

While an autonomous audit system is exciting, there are significant advantages to a portable human carried system. UAVs are limited by load carrying capacity, this limits camera quality and system battery life. A human carried system is less constrained by weight, meaning larger capacity batteries can be used for longer collection times and a higher quality, heavier camera can be used to aid in detection. The human operated device can ensure that



Figure 1. Data collection platform

the entire region is covered, even portions which would be impassable or difficult for a flying craft such as doorways, narrow passageways, and locations with variable height ceilings.

Since light sources are inherently the brightest objects in a scene, classical thresholding techniques for visible light imagery can be successfully used for detection. This eliminates the need for surface normal maps and points clouds meaning that expensive LiDAR systems need not be used. Instead, lower cost and readily available commercial hardware can be used to accomplish this task. A single camera, spectrometer, and localization device (Google Project Tango tablet) are combined into a portable unit to be carried by an ambulatory human operator. By using the Tango tablet with its precise localization abilities, an existing floor plan is not needed to locate the lights in a building.

Hardware Platform

To enable simple human operation of this system, all of the hardware components needed to be mounted together in an easy to carry fashion. Figure 1 shows how all components are assembled. For localization, a Google Project Tango Tablet was used. The camera, a Canon EOS 5D Mark III, was chosen for its ability to capture images at a high frame rate, six frames per second, until the memory card is full. In our experiments a 128GB card was used which gives approximately six hours of images. The lens, at its widest setting of 16mm, provides a 96.7-degree horizontal field of view and a 73.7-degree vertical field of view which is sufficient for this application. As shown in Figure 1 the camera and lens are pointed upwards to face the ceiling. The Ocean Optics Spark-VIS spectrometer was used with a diffuser to capture the emission

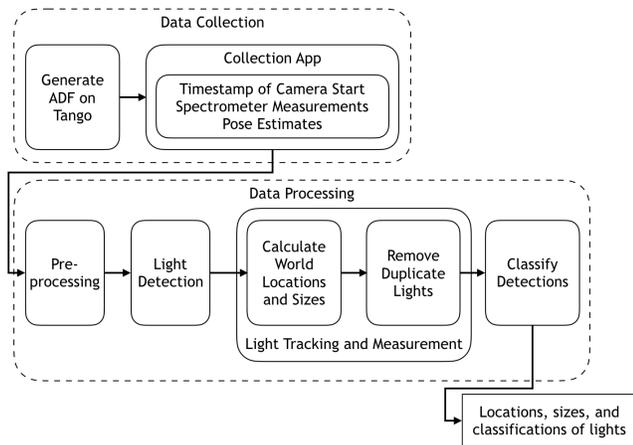


Figure 2. High level overview of the system.

spectra of lights. This sensor has a narrow field of view and can only accurately capture the spectrum when directly underneath a light. A custom controller circuit was built to interface the Tango tablet with the camera's shutter release over USB. All of these devices were attached together rigidly on a wooden board so that the locations of the components relative to each other was known.

Methods

An overview of the data collection and processing steps is shown in Figure 2.

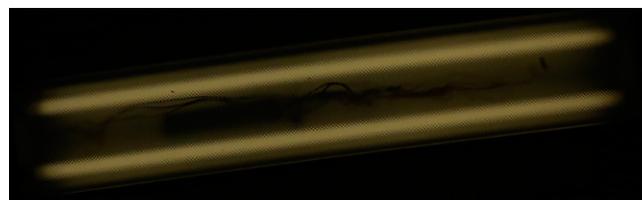
Data Collection

Data is collected in two phases. The first phase generates an Area Description File, ADF, using Google's Explorer application. This file contains a representation of visual landmarks within a given area. These landmarks are stored so that during data collection the Tango can more accurately localize itself. A custom Android application was written to then load ADF and connect the Tango to the spectrometer and camera for collecting data. When the user initiates the data collection process in the application, the camera begins capturing images at 6 Hz, the spectrometer records data at approximately 15 Hz, and the Tango records position and orientation data at 100 Hz. The camera is preset to capture images at a 1/4000 second shutter speed at f/22 aperture and ISO 3200 sensitivity. These settings were found to eliminate any motion blur and focus issues while having no significant noise in the images. All data measurements are recorded and timestamped according to the Tango's system clock. The exception is images captured on the camera. Here, the initial trigger time is recorded on the Tango and each image has a timestamp with respect to the camera's clock. These are synchronized in the pre-processing stage.

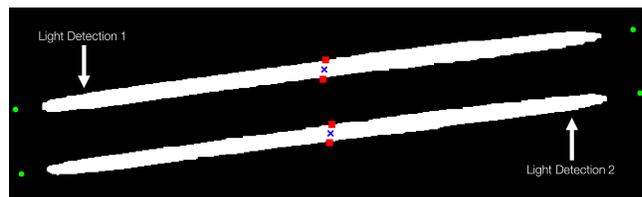
Data Processing

Pre-processing

After all of the data is collected, the raw data files are copied to a computer for processing. Here, the measurements from the spectrometer and Tango are converted into CSV format and the images are resized to 1500 × 1000 pixel JPGs to speed up the remainder of the processing steps. Using the timestamped



(a) Input light image



(b) Annotated light mask from blob analyzer

Figure 3. Light detection example. (a) Fluorescent tubes from the dataset (b) Light mask annotated with the centroid (blue x), end points of major axes (green dots), and minor axes (red squares) of each light

spectrometer and position data, each recorded emission spectra is tagged with its location in the real world. Similarly, each image is tagged with its location. However, since the image timestamps are based off of a different clock, they must first be aligned temporally. To do this, the timestamp when the camera was first triggered by the Tango added with a 61 millisecond shutter lag in the 5D Mark III camera is subtracted from the first image's timestamp. This time offset is then applied to each image to align them with the Tango's clock. Then, by finding the position measurement closest in time to the new timestamp, the image is tagged with position and orientation from the Tango measurements.

Light Detection

Since the images were captured with a low exposure value, the lights can be detected by thresholding the images. However, due to variations in fixtures, the bulbs can be partially obstructed or have bright reflections. To account for these and fill small gaps across the face of a bulb, the following morphological operations are applied:

1. Opening with a 3×3 pixel rectangle
2. Closing with a 20×20 pixel rectangle
3. Filling of holes
4. Dilation with a 30×30 pixel rectangle
5. Erosion with a 30×30 pixel rectangle

The resulting binary image represents a mask of where the lights are (white pixel) and aren't (black pixel). This mask is then inputted into Matlab's blob analyzer function which computes connected regions and returns the centroid, orientation angle, lengths of major and minor axes for each region. Each region is treated individually as a separate light. An example of the masking and blob detection is shown in Figure 3. If any part of the region touches the edge of the image, the detection is ignored and not used in tracking. This ensures that the entire light is contained in the frame for each image used.

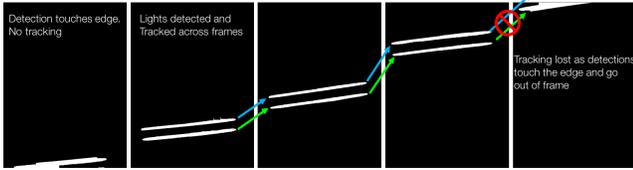


Figure 4. Light tracking example showing detected lights and their tracks, sequence shortened from original 15 frames.

Light Tracking

Since the images are captured at six frames per second, each light will appear in multiple frames. An example of the light tracking process is shown in Figure 4. In order to avoid over counting the lights, each is tracked between frames using the approach in [9]. Starting with the coordinate of the light centroid in pixel space, a Kalman filter based on a constant velocity model is used to predict the location of the light in a subsequent frame. Then, between the frames, the Munkres variant of the Hungarian assignment algorithm [10] is used to match the Kalman predicted centroids to detected centroids in the subsequent image in $O(n^3)$ time. The assignment algorithm runs on an $n \times n$ matrix of costs, in this case between the predicted lights from the Kalman filter and the newly detected lights. The cost is the pixel distance between the predicted and detected centroid. However, when a new light is detected or an old light goes out of frame, the number of detections can be mismatched. To account for this, additional rows or columns are added with a high cost, 1000 pixels. If a light is not currently being tracked or at any point leaves the frame, it is considered a new light. This is dealt with in the de-duplication section.

Light Localizing and Measurement

A key element of this project was determining the location of the lights in a building. In order to provide the location of the lights with respect to the real world, rather than pixels, pairs of images and their orientation and position information are used. Since accurate orientation and position information are collected on the Tango tablet, the change in position and orientation and can be treated as a stereo baseline for the two images. Using Matlab's triangulate function with the pair of corresponding light centroids, the coordinates of the light with respect to the first image's location are computed. By adding in the position of the first image, the real world position of the light is found. This same calculation is run over all pairs of images that contain the same light. For each light, the median centroid is found.

Using this same technique, but on the end points of the major and minor axes of each light the real world length and width of each light can be approximated. For a given detection with centroid $c = (c_x, c_y)$, orientation θ , major axis length M and minor axis length m , the four end points are:

$$\left(c_x - \frac{M \cos(\theta)}{2}, c_y + \frac{M \sin(\theta)}{2} \right) \quad (1)$$

$$\left(c_x + \frac{M \cos(\theta)}{2}, c_y - \frac{M \sin(\theta)}{2} \right) \quad (2)$$

$$\left(c_x - \frac{m \sin(\theta)}{2}, c_y + \frac{m \cos(\theta)}{2} \right) \quad (3)$$

$$\left(c_x + \frac{m \sin(\theta)}{2}, c_y - \frac{m \cos(\theta)}{2} \right) \quad (4)$$

The distances calculated along each axis can then be used to estimate the surface area of the lights.

De-duplication of Lights

While traveling through a building, the same light may be encountered in nonsuccessive frames. Heuristics are used to remove these duplicates. First, if two detections are within one meter of each other they can either be two bulbs in the same fixture or duplicates. When bulbs are in the same fixture, they will appear together in subsequent frames together. Once the duplicates are identified, they are treated as additional tracked frames of the original light to aid the classification and measurement procedures. Next, a second heuristic removed false detections which do not correspond to any light. This uses the fact that an operator must be directly underneath a light for it to be properly identified by the spectrometer. Therefore, if there is no image with the detected light within 200 pixels of the image center, the detection is discarded.

Light Classification

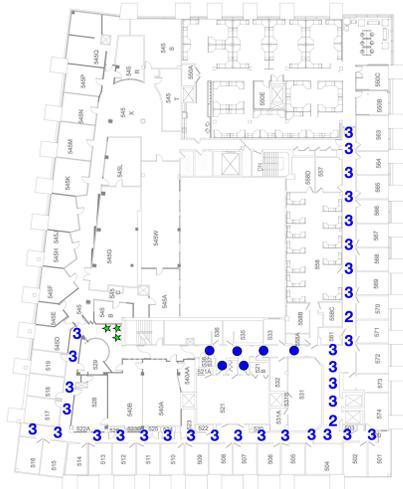
The final step is to classify the spectra of the detected lights based on their light source: LED, fluorescent, or incandescent/halogen. The spectrometer outputs a 1024×1 vector of intensity values for wavelengths across the visible spectrum. To account for slight variations in samples, every four elements of the vector are averaged together resulting in a 256×1 feature vector. This is then fed into a single layer feed-forward neural network with 50 neurons which was trained on over 7,000 spectra from commercially available light bulbs of varying wattages, styles, color temperatures, and chemistries. The closest five spectra in world-space to the detected light centroid were run through this classifier for each light detection. The light is assigned the label of the most frequent class from those five.

Results

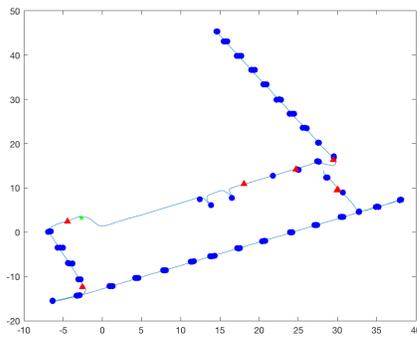
To evaluate this system, four datasets were collected on the campus of the University of California, Berkeley. The paths collected totaled 937 meters and contained 319 lights. Of those, 2 were incandescent bulbs, 32 were LED bulbs, and 285 were fluorescent lights. Figure 5a shows manually collected ground truth data for one dataset similar to what an auditor would create. Figure 5b shows the collection path as well as the detected lights.

For light detection, the positions of the algorithmically detected lights were visually compared against the ground truth and a correspondence between them denotes a correct detection. Twelve lights were missed in the data set and there were ten false positive detections. This means that 297 light were correctly detected. For light classification the number of correctly classified spectra is compared against the spectra for the 317 collected lights. Finally, measurement results compare the actual surface area of the lights to the estimated, in square meters. The results are summarized in Table 1.

In total, light detection had a 6.9% error rate and classification a 13.9% error rate. Light measurement estimation was the most erroneous at 105% error. In one of the regions traversed, there was a maze of cubicles. In order to collect the data, the device was extended at arm's length, an unstable pose, to maneuver



(a) Manually annotated ground truth



(b) Algorithmically detected lights and collection path

Figure 5. One of the datasets collected containing 106 fluorescent lights (blue circles and numbers) and 3 LED (green stars) lights with a 300 meter long path. Incandescent light detections are marked as red triangles in (b).

around desks and filing cabinets. This led to increased noise in the position and orientation measurements which are crucial to light measurement. Since the images and the Tango are only synchronized within 10 milliseconds, the position and orientation at the instant the picture was taken could be significantly different than those assigned to the image, further exasperating the problem.

Further work on this topic would involve a more stable and easier to maneuver platform for collecting data, such as a steadicam. While this project provides estimates of total lighting surface area and type of bulb which could be used to determine actual power consumption, additional sensors such as an infrared camera, could be added to measure actual energy consumption and provide more accurate detections and classifications.

Table 1. Summary of Results

Experiment	Ground Truth	Calculated	Error
Detection	319 lights	297 lights	6.9%
Classification	317 spectra	273 correct	13.9%
Measurement	11.06 m ²	22.27 m ²	105%

References

- [1] Cori Jackson and Konstantinos Papamichael, LIGHTING ELECTRICITY USE IN CALIFORNIA Baseline Assessment to Support AB 1109 (Task 2.18), California Lighting Technology Center, Davis, CA, 2014, pg. 22.
- [2] Moncef Krarti, Energy Audit of Building Systems: An Engineering Approach, Second Edition, CRC Press, Boca Raton, FL, 2011, pp. 1-1-4.
- [3] K&M Electric Supply, Inc., Commercial Lighting Energy Audits, Online, August 1, 2016, retrieved from <http://www.kmelectric.com/energy-audits/>.
- [4] Richard Zhang and Stefan A Candra and Kai Vetter and Avidesh Zakhor, Sensor fusion for semantic segmentation of urban scenes, Proc. IEEE International Conference on Robotics and Automation (ICRA), pg. 1850-1857. (2015).
- [5] Satarupa Mukherjee and PS Hariprasad and Omar Oreifej and Brian Pugh and Eric Turner and Avidesh Zakhor, Automatic Computer Detection and Power Estimation in Indoor Environments from Imagery, retrieved from http://www-video.eecs.berkeley.edu/papers/satarupa/global_sip.pdf. (2015).
- [6] Nicholas Corso and Avidesh Zakhor, Indoor Localization Algorithms for an Ambulatory Human Operated 3D Mobile Mapping System, MDPI, Remote Sensing, 5, pp. 6611-6646 (2013).
- [7] Christopher J Bay and Trevor J Terrill and Bryan P Rasmussen, Autonomous lighting assessments in buildings: part 1-robotic navigation and mapping, Taylor & Francis, Advances in Building Energy Research, volume 11 number 2, pp. 260-281 (2017).
- [8] Trevor J. Terrill and Christopher J. Bay and Bryan P. Rasmussen, Autonomous lighting assessments in buildings: part 2 - light identification and energy analysis, Taylor & Francis, Advances in Building Energy Research, volume 11 number 2, pp. 227-244 (2017).
- [9] Mathworks, Motion-Based Multiple Object Tracking, Online, August 1, 2016, retrieved from <http://www.mathworks.com/help/vision/examples/motion-based-multiple-object-tracking.html>.
- [10] James Munkres, Algorithms for the Assignment and Transportation Problems, Society for Industrial and Applied Mathematics, Journal of the Society for Industrial and Applied Mathematics, volume 5 number 1, pp. 32-38 (1957).

Author Biography

Craig Hiller received his BS (2015) and MS (2016) in electrical engineering and computer sciences from the University of California, Berkeley. His work there focused on indoor localization and applications to building efficiency.

Avidesh Zakhor joined the faculty at UC Berkeley in 1988 where she currently holds the Qualcomm Professor of Electrical Engineering and Computer Sciences. Her areas of interest include theories and applications of signal, image and video processing, 3D computer vision, and multimedia networking. She was a Hertz fellow from 1984 to 1988, received the Presidential Young Investigators (PYI) award in 1992. In 2001, she was elected as IEEE fellow and received the Okawa Prize in 2004.