Fine-grained measurement of the indoor built environment with robotic vacuum cleaners

Joseph McLaughlin[1], Michal Young[1], Siobhan Rockcastle[1]
[1]Univeristy of Oregon, Eugene, USA

Abstract

Research that links indoor environmental factors with human health and comfort has progressed significantly in recent decades, creating demand for high-quality data that can inform existing metrics and calibrate mechanical and electrical systems. A conventional sensor network that monitors light, temperature, and humidity has a spatial resolution proportional to the number of sensors within a space. This paper proposes a single roving sensor, constructed from inexpensive commodity parts and mounted on a commercially available robotic vacuum, that can record measurements with a uniquely fine spatial resolution. This device is capable of performing surveys of confined spaces to collect measurements with 10 cm spatial accuracy over sub-second time intervals. We share a regionally weighted interpolation method which we use evaluate this data. Using this device, we can monitor environmental conditions with a fine-grain spatial resolution that is resilient to changes in the position of furniture and occupants and is relatively low-cost.

Key Innovations

- Construction and evaluation of a low-cost and uniquely high-resolution and flexible roving environmental sensor.

Practical Implications

A low cost roving sensor permits detailed measurement in spaces that do not traditionally accommodate commercial sensor networks or where such networks are difficult to install. Measurements can be communicated in real-time, making this suitable for evaluating existing ambient conditions, driving control systems to improve occupant health and comfort, and commissioning an existing system to improve efficiency.

Introduction

Light, temperature, and humidity can have profound impacts on occupant comfort, productivity, alertness, and satisfaction (Amundadottir et al. (2017); Bluyssen (2013); Lan et al. (2010)). Mechanical and electrical systems that regulate these factors often rely on data from sensors to measure existing conditions and verify when a target set point or threshold is achieved. The use of sensors in the right locations within a building can improve operational efficiency and save energy by detecting anomalies and correcting operational problems (Fernandez et al. (2017)). Research behind the development of sensor systems for control of the built environment is well established (Sofos et al. (2020)), yet we lack an in-depth understanding of how environmental attributes vary over time and space within a majority of US building stock.

Conventional building sensors can either be wired directly into the building system or wirelessly mount to ceilings and wall surfaces or integrate within luminaires. While wireless installations are often less expensive to install, individual sensors have limited battery life and can be more expensive per unit due to networking and battery requirements (Cree et al. (2013)). Complex sensing platforms are most common in large buildings where the cost of operation is high and savings due to efficiency and occupant performance can result in significant long-term savings.

Environmental sensors are less common in residential buildings where improvements require a willingness to invest in new equipment, insulation, or changes in operational behavior. However, residential building typically consume more energy than commercial buildings. In the U.S. in 2019, the sum of energy consumed by all residential buildings was approximately 1.2×10^{20} J—22% more energy than all commercial buildings (U.S. Energy Information Administration (2020)). Lighting systems have become more efficient due to LED source technologies; outdated heating, cooling, and window technologies likely contribute to an increase in energy consumption and decrease in the indoor environmental quality (IEQ) in residential homes. Estimates from the U.S. Department of Energy (2011) suggest that up to 15% of heating and cooling costs could be saved by adjusting the setpoint for residential thermostats or by making non-invasive adjustment to filters, registers, window coverings, and radiant barriers.

As it were, robust and pervasive environmental sensing remains infeasible for many energy intensive sectors within the built environment. Even for buildings with complex sensing platforms, the resolution of available data is spatially constrained to the number of sensors in the system. Building systems with a fixed number of sensors need to make assumptions about the goals of building performance to meet those goals. People naturally reorganize the spaces they occupy over time, threatening any assumptions made when a sensor system was first installed and further threatening the system's ability to meet its goals. These conventional systems that rely on multiple, fixed pieces of hardware are coarse and inflexible to changes within the environment.

The same constraints of fixed sensing systems limit research into IEQ and performance. While past studies have evaluated the impacts of dynamic light and temperature exposure on occupant health and comfort, they have been constrained to laboratory, controlled field, or simulation-based studies (Kim et al. (2019); Danell and Rockcastle (2020)). In any case, more robust data would reduce our reliance on assumptions that may prove detrimental to occupant health and energy performance in many buildings.

Alternative Sensing Models

Several mobile sensing systems have been explored for different applications such as IEQ monitoring or contaminant detection (Chen et al. (2017); Reggente et al. (2010); Jin et al. (2018)). A paper by Jin et al. (2018) sought to address the scarcity of dynamic indoor environmental data by creating an automated mobile sensing device to measure environmental data from several attached sensors. This device performed autonomous navigation using Simultaneous Location and Mapping (SLAM) with an infrared camera system to move about the space. The system recorded sensor data periodically, associating sensor data with approximate locations in space. The authors performed an experiment to evaluate the device's performance by activating it in a controlled space, gradually introducing CO_2 into the space, and measuring the result. They controlled their experiment by placing fixed sensors at known locations in the space to verify the fidelity of sensor data and of the location data. This contribution demonstrates the potential of an agile mobile device for sensing the built environment, but it relies on a computationally expensive method of navigation that may not be well-suited to practical indoor environments.

Research by Ulaganathan et al. (2017, 2019) proposed the use of indoor occupants as navigational hosts to collect data on illuminance. Using an Actiwatch 2, this method allowed researchers to consistently measure a participant's light exposure over several days. The authors use the duration of time where values greater than 1000 lux were recorded as an approximation of when participants were outside and exposed to bright light. This approach did not incorporate location tracking nor was it intended for building-scale control operations. It nevertheless demonstrates the capacity of a mobile host to support an IoT sensor system and help researchers understand temporal conditions experienced over time.

Method

The limited adoption, high cost, and low spatial resolution of commercial environmental sensing devices constrains research into the impacts of the built environment on occupants and limits the design of energy efficient buildings that support occupant well-being. However, IoT devices including smart lighting fixtures, thermostat controls, and autonomous robot vacuum cleaners have become common end-consumer products. These devices give users access to fairly robust control and analysis of their environment. For example, the autonomous robot vacuum cleaner is typically designed with the goal of thoroughly navigating and vacuuming the floor of its environment without user interaction. Autonomously navigating a dynamic environment remains an open area of research, but inexpensive and widely available consumer robot vacuum cleaners accomplish it well enough to fulfill their purpose.

Our approach consists of a sensing system mounted on a common robotic vacuum cleaner and minimally interferes with the vacuum cleaner's typical operations. Our sensing system has two necessary functions: (i) environmental sensor readout capabilities and (ii) real-time location estimation; on every readout, the measurement is annotated with location and time. The vacuum cleaner with the sensor system attached is free to perform its typical duties and navigate the environment while the sensor system collects measurements. The result is a time series capturing both location and sensor measurements. We call this approach a semi-parasitic construction because our sensing system interfaces with the vacuum cleaner as a parasite—benefiting directly from the movement of the vacuum while minimally interfering with its regular processes.

In this section we detail the design of such a device, the construction of a pilot study, and a method of interpolating the resulting data.

Operational Design

We constructed a sensor system using a single board computer, inexpensive commodity sensors, and an ultra-wide band (UWB) location estimation network. The sensor system is mounted on the vacuum as it runs its cleaning cycle. Under typical conditions the vacuum is configured to run once per day. However, to run it more frequently and on specific time intervals we programmed the sensor system to communicate with the vacuum cleaner. This was possible as our robotic vacuum cleaner exposed a serial port, allowing it to accept commands from a host device—in our case the mounted computer. We call this construction *semi-parasitic* because we interrupt the vacuum's typical programming, depending upon the computer to launch and terminate the cleaning cycles in-synchronous with the sensor readouts.

We attached two sensors to the computer, one measuring air temperature (C) and relative humidity (%) and another measuring illuminance (lx). Sensor readouts are only sought when a position is determined by the location estimation network. The measurements are buffered and associated with the location estimation and the current time. In practice, buffering is necessary because different sensors readout at different time intervals. In our case, the illuminance sensor's readout is slower than the accompanying sensor.

We constructed the location estimation network using inexpensive, commercially available UWB DWM1001 radios. These networks produce location estimations with 10cm accuracy under typical conditions (Decawave (2020)). Our UWB network architecture follows an anchor-tag model, wherein anchor radios are fixed in the space at known locations and a tag radio is free to move around the space. We attached the anchor radios to the walls of the space and the tag radio to the vacuum. UWB sensors operate best with line-of-sight communication, so we elevate the sensors 1m high to increase the chance that line-of-sight is maintained between the anchor and tag.

The computer communicates with a server hosted on a cloud platform, which stores the sensor readouts in a SQL database. This places an additional requirement that the space have a wireless network. However, if network access was unavailable we could store the measurements locally on the device and process them later. An hour of continuous sensing from the device yields roughly 3MB of data comprising roughly 12000 readouts stored in an uncompressed text format.

The physical structure of the vacuum was modified to accommodate its new capabilities. An elevated tray was attached to the top of the vacuum to store the computer and a lithium-ion battery which acts as the computer's power source. We elevated the tray to not interfere with the manual button interface of the vacuum. An antenna arm extends off of the tray to anchor the environmental sensors and the UWB radio. Finally, the bumper was extended vertically with moulded ABS plastic to allow for collisions to engage the bumper and to protect the antenna arm from collisions. Figure 1 shows the completed device.

Pilot study

We conducted a pilot study in a 70 m^2 room at the University of Oregon. The room has two south facing windows, one ceiling-mounted lighting fixture, wall-mounted radiant heat, and ceiling-mounted air circulation. This space contains several desks, chairs, and laboratory equipment, and is frequently occupied by research assistants. The dynamic and unpredictable organization of furniture in the space provides a realistic context for evaluating the performance of the

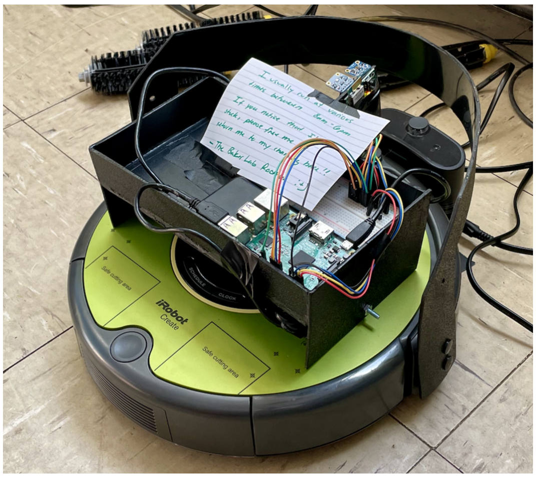

Figure 1: The completed device used in the pilot study. The structural modifications to the vacuum cleaner and completed sensor system is shown.

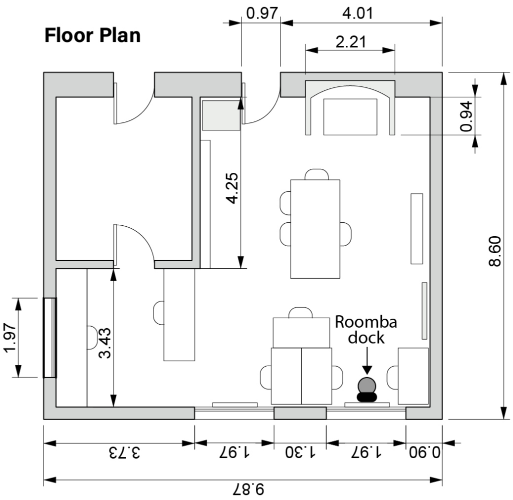

Figure 2: The floor plan of the room used in the pilot study, given in meters.

device. Figure 2 is a detailed floor plan of the space.

The operational device we constructed uses the Create 2 vacuum cleaner platform and a Raspberry Pi model 3 single board computer. We chose two low-cost sensors, a BME280 sensor to measure temperature and relative humidity, and a TSL2561 sensor to measure lux. The two sensors are connected to the computer through its GPIO interface and rest 15cm above the floor on the device's antenna arm. Three UWB radios were configured as anchors in the space. The radios were fixed 1 m above the floor and placed on three different walls of the room. A final radio is connected to the computer and is fixed to the antenna arm.

Interpolation Method

Sensor readings from a roaming sensor will only have one value at any point in time, but over longer time intervals the data over space can emerge to cover the entire area at a high resolution. This kind of data is distinct from data collected in fixed sensor networks and should be evaluated differently. This was first recognized by Jin et al. (2018) who describe this as an optimization problem where the goal is to find a function that best estimates the value of an unexplored location at a certain time, given a time series of locations and values $\mathcal{T} = \{t_1, p_1, v_1\} \ldots \{t_n, p_n, v_n\}$. To accomplish this, they base their estimation for a specific time and location, t_k, p_k from the LOWESS smoothed value taken at that time, $LOWESS(v_k)$, and refined with the residuals from the smoothing method, $r_i = LOWESS(v_i) - v_i$ captured by an inference method, e.g. KNN or random forest (Cleveland (1979)). The authors recognize that their method assumes that the dynamics of a room are most influenced by global trends such as those caused by outdoor weather or HVAC systems as well as local trends like those due to occupants or furniture. Their method encodes the global trends in the value from the smoothing method and the local trends in its residuals.

We observed that distinct regions within the space were more homogeneous than the space as a whole. For example, the lower left corner of the room was consistently cool and the locations near the window were consistently brighter. Adding to this, the motion of the vacuum is effectively random, leading to the possibility that it may survey a single region for extended periods. We therefore chose to forgo global trends and instead interpolate based on regional trends. The procedure for our method follows:

1. Scale the data to unit variance and perform k-means clustering to classify the data into j distinct regions over space, time, and value, $\mathcal{Y} = KMEANS(\mathcal{T})$. In practice, we chose to optimize j with the silhouette method. (Rousseeuw (1987)).
2. Each of the j regions is smoothed to determine a trend for each class, $LOWESS(v_i^y)$ for $y \in \mathcal{Y}$.
3. The value at any location and time is the k-nearest values over space and time, smoothed by their respective region.

This allows the residuals from multiple regions to influence a particular inference, particularly when samples come from several classes and if k is large.

Results

We recorded approximately 70 000 illuminance measurements and 100 000 temperature and relative humidity measurements over 10 hours of operation, spanning 5 days, from 6 January to 10 January 2021. The 10 hours is divided proportionally across each day. Each day comprised 6 collection periods spanning 20 minutes each. Each collection period consists of 2000–4000 annotated measurements, forming a map of the measured attributes across space and over the time interval of the collection period. This readout rate is approximately 2–3 measurements per second; however, lux measurements are consistently read at a lower rate than temperature due to limitations of our sensor hardware.

Figure 3 illustrates our raw, annotated measurements over a single collection period. In this instance the vacuum did not reach the lower left corner of the room in its 20 minute period. We found that this area of the room was missed in roughly half of all collections. We attribute this to the complexity of the space, and in particular that this area is only accessible by a roughly 1.5 m wide path. In general, we found that the room as a whole behaved as a suitable representation of a real-life space. Also notable in Figure 3 are the flux of relatively low measurements throughout. In this instance, anomalous measurements on the order of 0–10 lx can be attributed to interactions with occupants or with furniture objects such as a chairs or desks.

Acting as a parasite on the vacuum cleaner yields some uncertainties that manifest in the results. The 20 minute collection periods did not provide enough time for the vacuum to circumnavigate the entire room, and in general we found that areas with fewer obstructions were included in more collections. Figure 4 illustrates the frequency of measurements with respect to location across different collection periods in a single day. Here we see that 49% of the navigable regions (those with floor access) appeared in

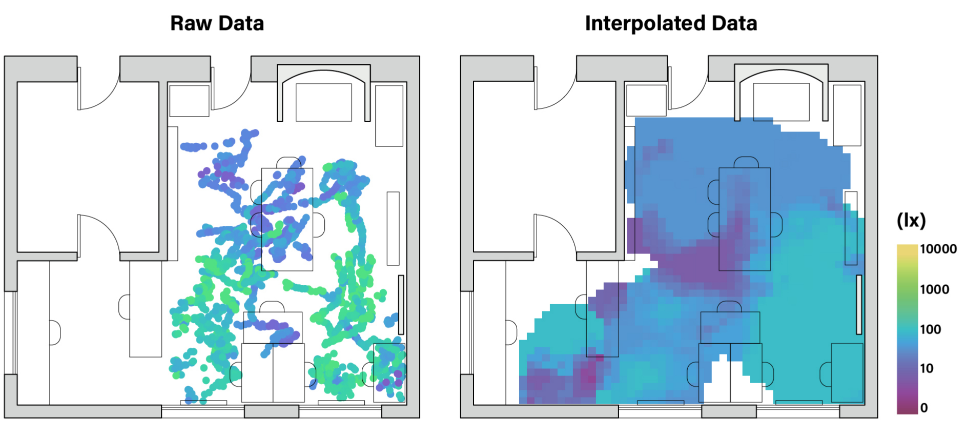

Figure 3: Left: illuminance as measured over space from a single collection period (12:00 — 12:20 local time, 7 January 2021). Raw measurements are plotted at their annotated location. Right: the result of interpolating with our regional interpolation method to 12:00. (log scaling used for lux.)

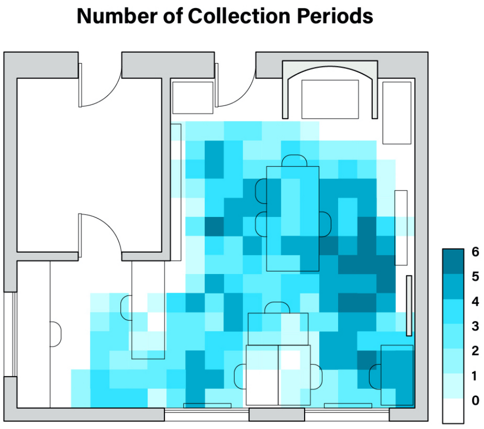

Figure 4: The frequency of measurements taken from $50cm^2$ regions of the space, over a single day (6 total collection periods on 8 January 2021.) Regions are counted if more than 10 measurements occurred in its bounding box.

at least 3 of 6 collection periods and that 29% of all navigable regions appeared in at least 4 of 6 collection periods; further, the regions that appeared in at least 5 collection periods are thoroughly distributed about the room, suggesting the vacuum is visiting several distant areas of the room within any given collection period. Variance in the number of measurements between collection periods is also notable but is likely to be a limitation of our location estimation network. The location estimation network depends upon inertial changes in the position of the device and performs best under line-of-sight. Instances where the vacuum visited areas with poor line-of-sight or when the vacuum maintained a steady velocity likely account for the variance in the number of measurements between collection periods.

The data shown on the left of Figure 3, while containing greater spatial resolution than conventional fixed sensors, remains relatively sparse. To produce a more consistent image we apply our regionally weighted interpolation method on the right of Figure 3. Figure 6 contains several time frames, interpolated with regional interpolation to produce spatially consistent images in sequence. This method produces a consistent visual representation of the measurements taken during a collection period. As our raw data occurs at several points in space, we also gain the ability to compare regions of different time frames containing nearby measurements. However, our regional method does capture the entirety of the dynamic range shown in Figure 3. The raw image in Figure 3 has several extremely bright lux measurements over 1000 lux that are likely due to daylight. Given that they are highly variable, they do not appear as significantly on the right of Figure 3. This is most likely due to our use of a local smoothing method which de-emphasizes outliers.

Another consequence of this method of interpolation

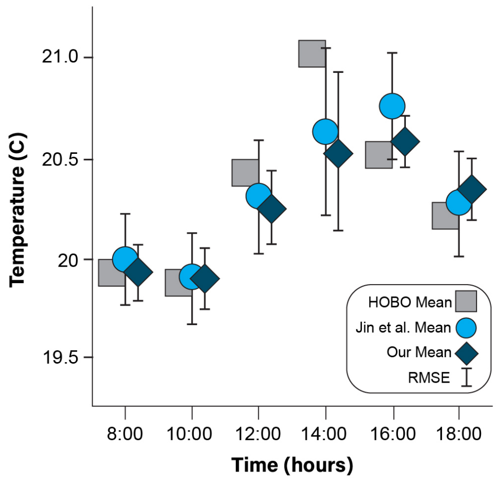

Figure 5: A comparison of our regional interpolation method compared to the method proposed by Jin et al. (2018). The dataset is comprised of sensor data collected 8 January 2021.)

is that the interpolation range does not respect the physical boundaries of the space, sometimes expanding the interpolated image into walls and solid objects where a measurement could not have been taken. In this pilot study we can simply remove these impossible values which cannot exist. In cases where the exact boundaries of the space are not known, it may not be possible to remove inaccurate interpolated values.

Evaluation & Baseline

We compare our method of interpolation with the method proposed by Jin et al. (2018) in Figure 5. They perform similarly, consistently interpolating within 0.1 C of each other. Our method infers with a lower root mean square error (RMSE), reliably 0.1–0.25 less than the global method. We obtain similar results without depending on an explicitly global trend. This may be because the global effects of environmental attributes (heat in this instance) are pervasive throughout the room. In other words, the propagation of heat over time is likely represented across space, making a temporally local point just as predictive as a spatially local point. We attribute our lower RSME score to the regions we derived from k-mean classifications. This preserves some of the regional features in the space, resulting in more accurate inferences. Both methods underfit the peak recorded by the HOBO data loggers at 14:00 local time.

The 6 HOBO data loggers fixed about the room we located approximately 15 cm above the floor to put them at the same height as the sensor system. In practice, this meant they were located on the lower faces of desks or table legs. As such, the fixed sensors did not gather light data that matched the profile of our device. The fixed sensors typically maintained a low lux level, with several incidental peaks. As our device traveled throughout the space, it was able to

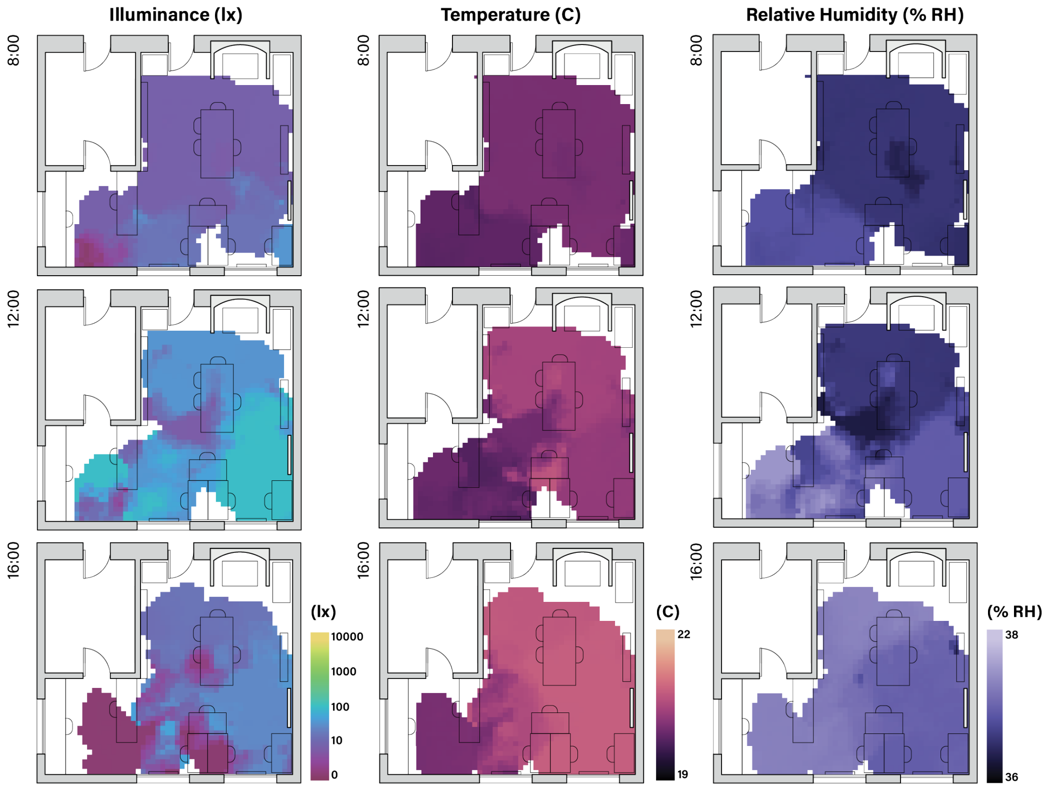

Figure 6: Illuminance, temperature, and relative humidity measured over space, interpolated with our regional interpolation method to 3 points in time on 8 January 2021. (log scaling used for lux.)

collect a robust profile of light as it behaves, not just on walls and surfaces of the room, but throughout it. Figure 7 plots average illuminance at 8:00–18:00 over the 5 days we collected data. Notably, Figure 7 follows a typical daylight decay pattern with the most intense lux levels read between 10:00–14:00. The only anomaly is 6 January where we note the weather was particularly saturated in cloud coverage.

In Figure 8 we plot a similar graph of the average temperature over the same time interval as Figure 7, plotted against the average fixed sensor measurements. From this we see a similar trend between the temperatures observed by both sensing systems. We attribute the peaks in our device's data to particularly warm heat sources discovered by the device while roaming. The fixed sensors only capture the heat after it has propagated from the source.

Discussion

The model we presented conveys a uniquely thorough spatial resolution when compared to conventional sensing devices. The results of our pilot study demonstrate that existing, low-cost infrastructure in the form of a robotic vacuum cleaner can be leveraged to produce highly accurate isotopic maps of environmental variables over space and time. This is accomplished with inexpensive parts and deployed with relative ease. The environment in which we deployed the pilot device is analogous to a residential space or a small commercial space. The frequency at which occupants visit the space poses a threat to our validity because we cannot be sure if they interrupted the device. In fact, we recorded 2 instances where the device was interrupted over the 5 days when we collected data. In one case, we cannot be certain that an occupant did not interfere, because no physical obstructions were found, but in either case these are challenges that would face a device in a real-world environment.

The results demonstrate that it is the relationship between measurements and even the relationship between multiple collections of measurement that convey the true value of this method. In practice we saw that no single vacuum cycle covered the entire room, yet we are able to derive form out of somewhat disparate sets of data. Similarly insightful value could be derived from the evaluating the relationship between multiple rooms, or in ensemble with multiple sensors.

Cost & Ownership

The hardware purchased to perform the pilot study included 4 UWB radios, a Create 2 Roomba, a Raspberry Pi 3, a Bosch BME280, an Adafruit TSL2561, 2 cell phone batteries, and several connectors to assemble the system. The total cost of these components is below \$500 USD. The total cost of the Onset HOBO data loggers used in the pilot study exceeds this easily at over \$600 USD. As the size of the space grows, up to the area that can be covered by a single robotic vacuum, our system will only require more UWB ra-

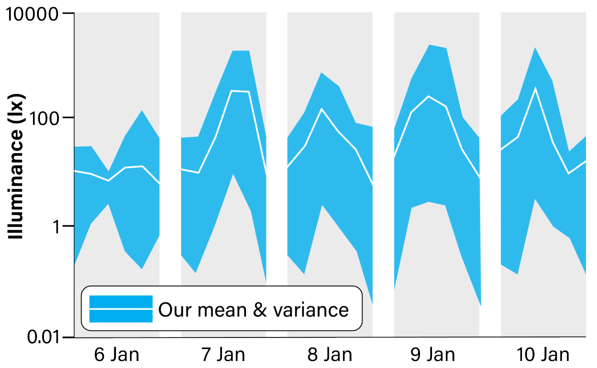

Figure 7: Mean illuminance and variance taken over 5 days (8:00 — 18:00) from 6 January 6 to 10 January 2021 and calculated over each collection period (log scaling used for lux.)

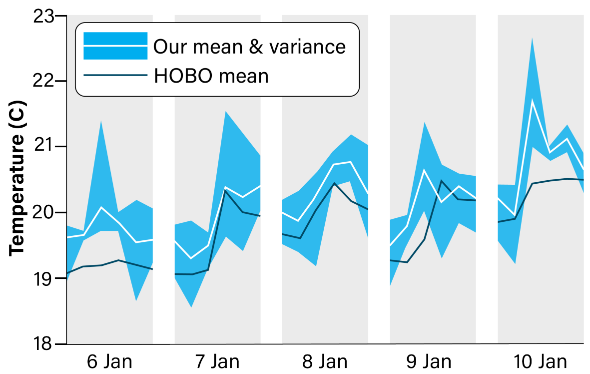

Figure 8: Mean temperature and variance taken over 5 days (8:00 — 18:00) from 6 January 6 to 10 January 2021 and calculated over each collection period. Plotted against HOBO fixed sensor baseline.

dios to maintain the strength of the positioning system. This however should only ever be necessary in spaces without line-of-sight or where the distance between the radios approaches 100m Decawave (2020). In contrast, to extend the range of a fixed sensor network, sensors should be added at regular intervals. As the size of the space grows the cost of the HOBO network grows proportionally.

Selecting locations for fixed sensors is also not inexpensive. Even if fixed sensors were preferred for long-term use, our method could be deployed temporarily to inform where in the space fixed sensors would provide the most relevant information.

Typical licensing agreements with commercial sensing systems do not place direct ownership of the system with the building operator. However, in this arrangement the system is likely to remain installed for decades. A mobile and easily deployable system not only broadens the types of spaces in which environmental sensing can be deployed, but it also the occupants *who* can deploy them. An inexpensive solution would favor renters or even prospective homebuyers.

Limitations

There are several limitations to our proposed sensing method.

- **Height.** In the pilot study we fixed the array of sensors at 15 cm above the floor. We chose this height because it integrated into the autonomous vacuum platform with relative ease. For some applications, measurements at typical waist height (300 cm), head height (600 cm), or ceiling height (800 cm) may be more desirable than floor height. While our implementation cannot sense at these heights, future work could explore the engineering required to perform measurements at different heights. Even so, the operational advantages of acting parasitically upon an existing means of locomotion may outweigh this limitation. Further, even though we cannot directly sense at greater heights, the variance that we detect between locations, and the accuracy at which we can fix location are their own benefits.

- **Location estimation.** The pilot study depends on a positioning system requiring a properly configured network of UWB radios for accurate location estimation. This system requires marginally more infrastructure than some other solutions, and not all applications may require the same degree of accuracy we achieved. Other methods, such as SLAM or WIFI ToF for example, may prove to be more appropriate to the specific needs of the application or of the space. To construct our system we first programmed into each of the UWB radios a relative position in space which it will occupy. To find a point in space of each device measured each point from a local origin point. This is a small inconvenience as we only did this once, but it does represent an opportunity to introduce human error. However, anchor self-localization methods such as Shi et al. (2019) address this issue.

- **Energy usage.** Our sensor system was powered by a single 10000 mAh battery. Daily inspection of the device included replacing the battery. We chose hardware for low cost and convenience rather than the lowest energy. Many consumer IoT devices share this characteristic. In practice, the cost of repeated sensor readouts and network usage has a significant energy cost which is only compounded by the rapid consumer adoption of IoT devices. We mitigate this in some regard by acting as a parasite upon existing infrastructure. The additional energy cost of sensing is insignificant in relation to the energy cost of operating the vacuum.

Conclusion

We presented a method of measuring the ambient attributes of the built environment with a unique spatial resolution. We demonstrated an implementation of the method and conducted a pilot study. The results contain approximately 100 000 measurements over 10 hours. The details of the results show how data can be accurately measured over space and interpolated to convey the dynamics of the room. The benefits and limitations of this method are discussed in context with current conventional devices. Future work and applications are explored.

References

Amundadottir, M. L., S. Rockcastle, M. S. Khanie, and M. Andersen (2017). A human-centric approach to assess daylight in buildings for non-visual health potential, visual interest and gaze behavior. *Building and Environment 113*, 5–21.

Bluyssen, P. M. (2013). *The healthy indoor environment: How to assess occupants' wellbeing in buildings.* Routledge.

Chen, Y., H. Cai, Z. Chen, and Q. Feng (2017). Using multi-robot active olfaction method to locate time-varying contaminant source in indoor environment. *Building and Environment 118*, 101–112.

Cleveland, W. S. (1979). Robust locally weighted regression and smoothing scatterplots. *Journal of the American statistical association 74*(368), 829–836.

Pacific Northwest National Lab.(PNNL), Richland, WA (United States) (2013). *Sensor Characteristics Reference Guide*.

Danell, M., A. M. L. and S. F. Rockcastle (2020). Evaluating temporal and spatial light exposure profiles for typical building occupants. Vienna Virtual Conference, May 25 – 27, 2020.

Decawave (2020). Real time location systems: An introduction. `https://www.decawave.com/sites/default/files/resources/aps003_dw1000_rtls_introduction.pdf` Accessed 2020-01-10.

Pacific Northwest National Lab.(PNNL), Richland, WA (United States) (2017). *Impacts of commercial building controls on energy savings and peak load reduction*.

Jin, M., S. Liu, S. Schiavon, and C. Spanos (2018). Automated mobile sensing: Towards high-granularity agile indoor environmental quality monitoring. *Building and Environment 127*, 268–276.

Kim, J., F. Bauman, P. Raftery, E. Arens, H. Zhang, G. Fierro, M. Andersen, and D. Culler (2019). Occupant comfort and behavior: High-resolution data from a 6-month field study of personal comfort systems with 37 real office workers. *Building and Environment 148*, 348–360.

Lan, L., Z. Lian, and L. Pan (2010). The effects of air temperature on office workers' well-being, workload and productivity-evaluated with subjective ratings. *Applied ergonomics 42*(1), 29–36.

Reggente, M., A. Mondini, G. Ferri, B. Mazzolai, A. Manzi, M. Gabelletti, P. Dario, and A. J. Lilienthal (2010). The dustbot system: Using mobile robots to monitor pollution in pedestrian area. *Proc. of NOSE*.

Rousseeuw, P. J. (1987). Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics 20*, 53–65.

Shi, Q., S. Zhao, X. Cui, M. Lu, and M. Jia (2019). Anchor self-localization algorithm based on uwb ranging and inertial measurements. *Tsinghua Science and Technology 24*(6), 728–737.

National Renewable Energy Lab.(NREL), Golden, CO (United States) (2020). *Innovations in sensors and controls for building energy management: Research and development opportunities report for emerging technologies*.

Ulaganathan, S., S. A. Read, M. J. Collins, and S. J. Vincent (2017). Measurement duration and frequency impact objective light exposure measures. *Optometry and Vision Science 94*(5), 588–597.

Ulaganathan, S., S. A. Read, M. J. Collins, and S. J. Vincent (2019). Influence of seasons upon personal light exposure and longitudinal axial length changes in young adults. *Acta ophthalmologica 97*(2), e256–e265.

U.S. Department of Energy, Washington, U.S. (2011). *EnergySavers: Tips on Saving Money Energy at Home, prepared by the Office of Energy Efficiency Renewable Energy*.

U.S. Energy Information Administration, Washington, U.S. (2020). *Annual Energy Outlook*.