# Integrating Contextual Data for Real-World Insights in Living Labs

by

**Alan Hwa Wang**
Rancho Palos Verdes, CA

A Dissertation submitted to the Graduate Faculty
of the University of Virginia in Candidacy for the Degree of
Doctor of Philosophy
in the

Department of Electrical and Computer Engineering
University of Virginia
Link Lab
August 2023

Doctoral Committee:
Brad Campbell, *chair*, Assistant Professor, Computer Science
Jack Stankovic, BP America Professor, Computer Science
Seongkook Heo, Assistant Professor, Computer Science
Mircea Stan, Professor, Computer Engineering
Arsalan Heydarian, *advisor*, Assistant Professor, Engineering Systems and Environment

# Dedication

*To my loving family.*

## Acknowledgement

First, I want to thank my advisor, Professor Heydarian, for taking a chance on a first-generation Taiwanese American and first-generation college student. This has been a journey since that sandwich on a lazy Thursday afternoon in the Southern California sunshine. I also want to thank Professor Campbell for graciously sharing his guidance on all aspects of academia, from how to set up a food shelf and how to write a proper thesis to how to network at conferences. To Professor Dugan, thank you for believing in me when I needed it the most. To Professor Heo, thank you for showing me that joy and learning can coexist unironically in graduate school. I want to thank the Taiwanese Graduate Student Association and the love of the 195 crew, especially Kevin, the Chen's, and the Lu's, for being my family and my home in a foreign town. I also want to thank all my mentors and friends at Cavalier Judo, Sensei Shih, and Sensei Rogusta for teaching me about *Seiryoku-Zenyo*, and *Jita-Kyoei*, flavors that pair well with dissertation work. Matt, Michael, Jack, and Faith, thank you for the lasting camaraderie you brought into my life, especially before and during the valleys of COVID. To Alexander, thank you for showing me the world of entrepreneurship and swing dancing, much-needed trips outside the ivory tower. To Angelique, thank you for bringing joy and responsibility into my life. And finally, I thank the original Link Lab gym crew, Yawen, Wengpeng, Jingyun, and Su, from whom I learned about cheer in adversity, gains, and the meaning of *Zha*.

# Contents

# List of Figures

# List of Tables

# Definitions

- **Cyber Physical System**: "Cyber-physical systems (CPS) are engineered systems that are built from, and depend upon, the seamless integration of computation and physical components." [2]

- **Living Labs**: A cyber-physical system that occupants normally inhabit that enables long-term and reproducible experiments

- **Contextual Data**: Data streams that are not directly measured by the sensors in a living lab (i.e., manually recorded medical events, building information, occupant information)

# Abstract

Research has shown that access to occupant behavior data in buildings can reduce energy consumption and improve occupants' productivity, comfort, and well-being. However, behaviors can vary across cultural, geographic, building, environmental, and contextual settings. Therefore, to increase our understanding of the long-term naturalistic behavior of occupants, more living labs are emerging across different countries, offering an opportunity to address existing research gaps. With the growth of IoT and ubiquitous computing, it has become easier to replicate and validate short and long-term data across different contexts. However, selecting the type, quantity, and position of sensors needs to be more cohesive with building information and activity simulation to avoid inaccurate, redundant, and privacy-intrusive sensing issues. This work tackles these critical challenges of living lab by demonstrating:

1. A methodology for integrating building simulation models to identify optimal sensor placements with privacy-preserving sensing considerations,
2. A longitudinal in-hospital case study that integrates medical events data and environmental sensor streams to predict momentary patient sleep disruptions, and in our remaining work,
3. A novel methodology for integrating information extracted from building plans to support fault detection of long-term energy harvesting sensor deployments

Overall, the three chapters in this dissertation proposal demonstrate contributions to three pillars of living labs, from *instrumentation* (respectful sensor installation in Chapter 1), to *utility* (support for occupant well-being using sensor and contextual data streams in Chapter 2), to *maintenance* (improving the reliability of long-term sensor deployments in Chapter 3).

# Chapter I

# Introduction

The average American spends more than 90% of their lives indoors [3, 4], and buildings account for 40% of the total energy consumption in America [5]. Together, it is unsurprising to find that if a building is properly designed and operated around the occupants' needs, preferences, and comfort levels, we can reduce consumption significantly [6]. In addition to reducing energy consumption, the study of the indoor environment has also been shown to affect the occupant's health and well-being significantly [7, 8, 9]. These studies indicate that improving of the indoor environment is a pressing health concern and also financially beneficial. However, a lack of standardization in the production of buildings compared to the automobile industry and poor information and communications technology (ICT) infrastructure in pre-existing buildings prevents building managers from achieving the 15%-50% energy-saving advanced control strategies have demonstrated [10]. Health studies have shown that proper management of the environment can lead to better physiological and psychological outcomes for occupants [11]. However, relying on employee self-reported surveys instead of quantitative measures such as Health Performance Indicators (HPI) limits the potential for buildings to support occupant health and well-being [12]. Researchers have created an approach called "Living Labs" to tackle these issues together. While many definitions for a living lab exist [13, 14, 15], a previous survey of existing living labs proposed a general definition [16]:

> *"A living lab... is a ... typical indoor environment where everyday tasks are performed by occupants over a significant period of time to experimentally characterize their activities and responses, and with a permanent setup that allows hosting scientific experiments ...*

*by monitoring and controlling the indoor conditions..."*

However, the definition assumes certain qualifiers that make it flexible for interpretation. A prior review of field implementations of occupant-centric building controls–a similar topic around the domain of Human-Building-Interactions (HBI)–showcases that a majority of studies were conducted for less than three months within ten zones and mainly covered zone-related controls. This means that studies do not observe potentially large seasonal effects of the environment, nor cover enough spatial diversity to be scalable. In the review, Park et al. further note a lack of standardization for measurement and verification and that occupants' privacy and data control are often overlooked. This and other works by the International Energy Agency Energy in Buildings and Communities Programme (IEA-EBC) [6] signaled to explore further ways to improve occupant privacy in buildings, ways to benefit occupant health and well-being in buildings, and ways to support maintenance and reliability of sensing for long periods. Amidst the many different and sometimes contradicting definitions for living labs [13, 14, 15, 17], this dissertation propose this definition for a living lab:

*"A living lab is a cyber-physical system that occupants normally inhabit that enables long-term and reproducible experiments."*

This definition is used to focus on the relationship between the buildings, sensors, and people and how interactions between the three entities can be improved for longer than three months. The following chapters explore these relationships among three different categorical environments. Experiments were conducted to map out the uses of residential space and all possible combinations of light states in an arguably more private and controllable indoor space in Chapter 1. Then, Chapter 2 focuses on experimentation in a university hospital, to demonstrate that these sensors can have practical applications in the real world. Finally, Chapter 3 explores experiments in a commercial research building, to explore issues in a large research test bed covering both public and private places.

**Thesis Statement** Installing environmental sensors in buildings has shown to be reliable and valuable, increasing health considerations, comfort, and energy efficiency. As technology improves, these sensors have grown smaller, more perceptive, and more ubiquitous. However, a lack of understanding exists of the information collected by different sensors at various quantities and distributions. Therefore, this dissertation introduces a novel method to contextualize building activities and utilize building simulation models to evaluate the information that can be collected with different quantities and distributions of sensors. By leveraging contextual data, applications can approximate the minimum number of sensor positions required for inferences, predict occupant behaviors, and classify faulty sensors that require maintenance.

## I.1 Problem Area

The area of contextual data integration spans multiple disciplines and different parts of a sensor deployment timeline. The closest match for this dissertation is in the realm of Human-Building Interactions (HBI), a subdomain of Human-Computer Interaction (HCI) [18]. For descriptions of the sensor deployment timeline, we can generally divide the problem into (1) before installation, (2) during installation, and (3) after installation. Before sensor installation, work is done to determine the sampling rate required to detect the intended behavior. Too low a sampling rate, according to the Nyquist-Shannon sampling theorem [19], and the behavior cannot be appropriately reversed. But too high a sampling rate, and *signal aliasing* occurs, and accurate inferences are sacrificed [20]. Once the correct sensor and sampling rate are chosen, we run into issues regarding the proper placement of the sensor. Prior works showcase this falling into a more computational realm, involving spatial statistics and autonomous sampling [21, 22]. The issue of optimal positions is further complicated by issues relating to the beneficiary and ownership of the data: a researcher interested in collecting the most amount of data from building occupants may not be as interested in the user's privacy, and the privacy-utility spectrum will sometimes be avoided for a straightforward solution, such as collecting data and then removing it [23]. Simulations and computational algorithms were combined not to find the most utility but rather to elucidate the bounds of inferences and how to safeguard user privacy before the first sensor is installed.

Following the issue of sensor installation, the issue of utility is explored. Sensors have been used for a vast myriad of applications and in an abundance of important and diverse areas, from predicting forest fires in California [24] to reducing daylight usage in buildings [25], to inferring noise using visuals for information retrieval. As such, we do not try to prove the usefulness of sensors but rather try to justify that long-term data can have a significant impact on the health of people in general. Specifically, we demonstrate a relationship between sleep and environmental disruptions, a known issue in hospitals [26]. In this instance, we again demonstrate the importance of *contextual data* and show how leveraging collaborations with our nursing partners can lead to valuable insights supporting better patient care. Finally, numerous works have showcased the difficulty of long-term deployments [27, 28]. We explored using energy harvesting devices because at face value they promise sustainable value without the need for constant retrofitting. However, in the maintenance chapter, we will demonstrate how modern energy harvesting sensors can come with a new set of challenges that can hinder reliability, namely issues surrounding diagnostics and indoor signal attenuation, and how we can use a technological solution to support them.

## I.2 Overview

Overall in this dissertation, we aim to demonstrate the overarching synergy that can exist when exploring multidisciplinary research involving the use of contextual data in long-term, naturalistic indoor environments. Specifically, we ask: 1) What are the challenges in translating inferences found in the real world to inferences found in the simulations? How can simulation be used to inform us about how informative a set of light sensors positions are given some assumptions about space use? 2) What are the important environmental patterns across different time scales that can negatively impact patient sleep? To what extent can environmental variables and medical events data predict patient sleep disruption? 3) What are the new modes of failure and consideration required to maintain and scale a network of energy harvesting sensors for long-term deployments? To answer these questions, we complete the following tasks:

**Task 1**: Light sensors were deployed, and lamps in a residential building were retrofitted with automated lamps. The total lighting state for all combinations of possible light states was sampled for all selected sensor positions by manually installing the sensor and automating the lights. A lighting simulation model using the floor plan of the unit was then developed. The lights in the simulation were parameterized to explore the same combination of lighting states to compare the digital twin results to the original twin. Afterward, we added a more nuanced but regularly occurring movement of doors, multiplying the state of possible inferences by a factor of nine. Given those states, we use the greedy set cover approximation algorithm to find an approximately optimal solution of required light sensor positions to cover all possible occupant behavioral states.

**Task 2**: Environmental sensors were deployed into five hospital rooms in a hospital setting and recruited participants to wear a smartwatch that tracks actigraphy over six months. The nursing team was consulted to acquire sleep-related medical events data to draw insights about the patient's sleep. The relationship between the frequency of the patient's arousal from sleep and environmental perturbations outside of the comfortable ranges listed in the literature was observed using a mixed-effect model.

**Task 3**: Over 125 energy harvesting sensors of varying types were deployed into a 17,000-square foot research infrastructure and collected data for six months, routinely adding more gateways to collect data. The building plans were then acquired from the architect and digitized to encapsulate distance and material information for tracing the medium between each sensor-gateway pair. Finally, the frequency of the data transmitted was analyzed. A representational encoder was applied to predict the signals of unseen sensor locations that were masked from our training set.

Following the completion of the tasks, this dissertation elucidates **Key 1**: The perfect sum algorithm combined with the greedy set cover approximation algorithm is a viable solution to make detailed inferences of indoor occupant behavior and find optimal sensor quantity and position, and is generally useful when run in simulation environments. **Key 2**: $CO_2$, medical events, lux, volatile organic compounds (VOC), and noise are significantly related to patient sleep disruption. Lower environmental perturbations were observed during the weekends and sometimes higher perturbations after midnight but before the start of the day shift when comparing interquartile ranges. **Key 3**: when digitized architectural plans are integrated, the silent rate of unseen sensor locations can be accurately predicted for different energy harvesting sensors. The predictions can inform future sensor placement or detect faulty downstream sensors.

This dissertation contributes by showing how contextual data integration can be helpful throughout various living lab scenarios. This dissertation demonstrates the quantity and placement of sensors to control the scope of inferences before installation, how the installed environmental sensors can elucidate and quantify the potential environmental disturbances, and how architectural layout can be leveraged to help maintain a network of energy harvesting sensors for scalable long-term deployments.

# Chapter 1

# Instrumentation

*One of the most fundamental ways contextual data relates to sensor data is about the placement of sensors. Building performance simulations enable researchers and sensor installers to evaluate the scope of observable behaviors prior to physically instrumenting the sensors. This chapter explores the distance between the digital and original twin and demonstrates with which algorithms and how much building performance-assisted placements can improve inferences and reduce the required sensors.*

## Integrating building simulation models to identify optimal light sensor placements with privacy-preserving sensing considerations

As IoT devices become cheaper, smaller, and more ubiquitously deployed, they can reveal more information than their intended design and threaten user privacy. Indoor Environmental Quality (IEQ) sensors previously installed for energy savings and indoor health monitoring have emerged as an avenue to infer sensitive occupant information. For example, light sensors are a known conduit for inspecting room occupancy status with motion-sensitive lights. Light signals can also infer sensitive data such as occupant identity and digital screen information. To limit sensor overreach, we explore the selection of sensor placements as a methodology. Specifically, in this proof-of-concept exploration, we demonstrate the potential of physics-based simulation models to quantify the minimal number of positions necessary to capture sensitive inferences. We show how a single well-placed sensor can be sufficient in specific building contexts to holistically capture its environmental states and how additional well-placed sensors can contribute to more granular inferences. We contribute a device-agnostic and building-adaptive workflow to respectfully capture inferable occupant activity and elaborate on the implications of incorporating building simulations into sensing schemes in the real world.

## 1.1 Introduction

Due to increased awareness of energy reduction measures in buildings over the past two decades, numerous technological advancements have been introduced to monitor the changes in indoor conditions. Sensors and actuators have become increasingly integrated into buildings to reduce overall energy consumption while improving occupant comfort [16, 6]. For example, building automation systems can reduce a building's energy consumption by dimming artificial lighting when sufficient daylight is sensed in the building [25]. The building can also utilize occupancy and air quality sensors to reduce energy demand by Heating Ventilation and Air Conditioning (HVAC) units in anticipation of occupants' presence or comfort [29]. The number of sensors installed in buildings will only grow with increasing energy prices and the known benefits of smart environments [30]. However, numerous challenges still exist to using sensor-collected data to improve the utility of occupants.

Firstly, and fundamentally, sensors have different data collection frequencies, so researchers cannot simply purchase any environmental sensor and install it to capture all the activity happening indoors. The frequency of data collection restricts the types of occupant behaviors that can be inferred. For instance, the Nyquist-Shannon sampling theorem demonstrates that you need to sample at more than twice the highest frequency component of the signal to correctly reverse it [19]. The time scale differences mean that researchers cannot use a one-image per thirty-second camera to capture the behavior of subjects that occur every 15 second without losing data. Increasing the frequency, on the other end, can cause other undesirable effects, such as signal aliasing. A researcher (designer or facility manager) is still required to decide the installed sensor's specifications, such as frequency, mode, and observed variable; the very act of purchasing the hardware restricts in perpetuity the downstream observable behavior.

Secondly, sensor placements have a large effect on the downstream utility of the sensor's collected data, but positioning is often overlooked or starts off randomly then iteratively improved [23]. Sensors deployed at incorrect positions can result in incorrect readings [31, 32], but it can feel more convenient to install sensors and start collecting data as soon as possible. Furthermore, prior sensors installed for evaluating a building can become insufficient or undesirable for new uses of space. For example, some residents can move out, rendering prior sensors in locations no longer occupied redundant. Similarly, additional residents can move in, making previous coverage inadequate for the new use of space. Optimizing the position and number of sensors can result in lower energy consumption and better readings, but it is challenging to incorporate and maintain considerations of changing contexts and sensing objectives for a scaling number of sensors manually.

Lastly, in the expediency of collecting ever more data, installed sensors can inadvertently expose more information than necessary for its intended use, leading to privacy concerns for the occupants. For instance, the "sensing by proxy" paradigm demonstrates how proxy measurements such as $CO_2$ can infer occupant count and activity [33, 34]. Similar granular occupant activity information has been observed by other sensors as well. For example, cooking activities can be observed via the fluctuation of PM 2.5 [35]. Granular appliance use can also be effectively disaggregated via non-intrusive load monitoring [36]. For instance, appliances such as coffee makers and hair dryers can have unique energy-use signatures relating to start-up processes and the physical makeup of the appliance. By installing a load monitor at the circuit level, the total energy use can be disaggregated from individual contributions based on the appliance's unique signatures, allowing for invasive inferences of occupant activity without needing to install any sensors inside the building itself. Other examples, such as using cameras and motion amplification, have allowed researchers to exaggerate the vibration of snack bags and reverse engineer decipherable sounds using visual data [37].

As more projects utilize machine learning and other computational methods to retrieve sensitive data from the indoor environment, we instead consider if similar computational methods exist to help reduce inferable information from sensors and protect the privacy of building occupants. One promising avenue that can help sensor installations navigate the potential overreach of sensing is at the intersection of simulations and sensor placements [38]. Simulations have traditionally been used to assess different building performance attributes during the design phase. The orientation of buildings and placement of windows, for example, can be explored and quantified using a score called daylight autonomy [39]. Simulations of the weather conditions, and movement of the sun, in conjunction with the location of the building, size, material, and orientation of the window and room, enable architects to uncover the total amount of time over a whole year when daylight can effectively stand in for artificial lighting. Similar simulations and metrics can be found for HVAC, where given the hours that the building will be occupied, the room size, the expected occupancy, the total energy use, and required airflow can be predicted and quantified [40]. The acoustic qualities of a building can also be designed and tailored to better match the intended use of space, such as longer reverberations for music halls and shorter reverberations for classrooms [41]. The timing in the pipeline during which these simulations are used represents a fundamental discourse in digital twins [42]. Specifically, a simulation model can help optimize the placement of sensors to inform on occupant activity, and it does not need to be run in real-time parallel to the physical twin for it to have a lasting impact on the smart environment.

Compared to manual or autonomous methods of sensor position optimization, where robots are used to sample the environment routinely and iteratively uncover the most optimal positions for the static environmental sensors [43], simulations enable a low-cost alternative for testing unlimited virtual sensors positions at the cost of computational power. Furthermore, instead of navigating protocols for institutional review boards (IRB) or logistics of the environment (e.g., to avoid a party or speaker event in the building), simulations enable researchers to avoid a broad range of complexity that can cause nontrivial disruptions for both study administrators and study participants. Researchers (designers or other decision makers) can run unlimited what-if scenarios to see how different environmental or user-related factors may impact the changes in the downstream signal before interacting with the physical environment. For example, the movement of occupants can be simulated using artificial agents to assess the ease of navigating the space in case of emergencies [44], the sun's movement can be simulated using historical sky information [45], and stochastic models can be used to simulate an occupants interaction with building controls [46].

In this chapter, we demonstrate the potential of physics-based simulation models to quantify the minimal number of positions necessary to capture sensitive inferences. We show how a single well-placed sensor can be sufficient in specific building contexts to holistically capture its environmental states and how additional well-placed sensors can contribute to more granular inferences. Specifically, we focus on lighting simulations as a test case because of its accessibility and geometrically-dependant nature. We answer two research questions:

- **RQ1**: What are the challenges in translating lighting inferences found in the real world to inferences found in the simulations? And,

- **RQ2**: How can lighting simulation be used to inform us about how informative a set of light sensors positions are given some assumptions about space-use?

We advocate for the use of simulation as a standard tool for 1) identifying the ideal location for sensors and minimizing the number of sensors distributed and 2) identifying potential information that sensors can collect when deployed in real life by showing the capabilities of simulations to calculate building states containing granular occupant activity exhaustively. In other words, we show how the simulated environment allows researchers to assess the effects that the position of sensors and the geometry of the building can have on occupant activity inferences. The workflow demonstrates an avenue for future researchers to verify the possible inferences of existing sensor positions or use the "adjustment of sensor positions" as a method to limit sensor inference overreach.

## 1.2 Methods

We consider the scenario where a researcher is trying to ascertain the light state of on and off for individual luminaires. Given the additive nature of light, we approach decomposing the summed light contribution at a sensor point in the building by formulating it as a *Perfect Sum Problem* with a noise threshold $\epsilon$. Since light intensity diminishes equal to the inverse of the square distance from the source, and each luminaire has photometric data, moving a light sensor's final placement at a variable distance away from light sources can enable individual contributions to be disaggregated. For instance, we can utilize the geometry and drop-off of light intensity by distance to coordinate unique fingerprints for each luminaire in range. Furthermore, utilizing the entire building space as potential placement areas enables researchers to exhaustively explore the ability to utilize *sets* of potential sensor positions that traditionally require repeated trial-and-error to capture.

### 1.2.1 State Inferences

Let $L = [l_0, l_1, ...l_{n-1}]$, where $L$ is a configuration vector of $n$ individual light source states $l_i \in \{0, 1\}$. Then, for any location $s$, given a configuration $L$, the contribution vector can be seen as $X(s, L) = [x_0, x_1, ..., x_{n-1}]$, where the contributions from each light source $x_i$ correspond to each light state $l_i$ modified by distance and obstruction. The maximum number of possible configurations is then equal to the cardinality of the power set with each light source, or $2^n$.

Given a combination sum solver, we can take a target sum $\mathbf{K}$, threshold $\epsilon$, and list of individual contributions $[x_0, x_1, ...x_{n-1}]$ and return a list of lists that contain all possible combination of contributions that add up to $\mathbf{K} \pm \epsilon$. Each list corresponds to a possible configuration that fits the constraints, but since ultimately only one configuration is correct, we calculate the accuracy for each inference using Accuracy $= \frac{|L \cap L_{\text{infer}}|}{|L \cup L_{\text{infer}}|}$ and return the mean.

Sensor readings from multiple points do not always corroborate to the same lighting configurations. To overcome this, we disambiguate misaligned light configuration inferences by using a voting vector $V(s, L) = \left[ v_0, v_1, ...v_n \right]$, from each sensor, where $v_i$ is 1 if the luminaire is determined to be on, -1 if the light is determined to be off, and 0 if the light cannot be detected (i.e., the sensor is out of range of the luminaire). The final inference is chosen based on the summed votes. If the value is greater than zero, the luminaire is on. If the final value is less than zero, we consider the luminaire off.

### 1.2.2 Distinctness Score

To generalize the real-world sensing to building simulations, we introduce an error threshold $\tau$, and define a distinctness vector $D_\tau$ as:

$$D_\tau(X(s, L)) = e_0, e_1, ..., e_{n-1}, \text{where } e_i = \begin{cases} 1, & \text{if }_j |x_i - x_j| > \tau, i \neq j \\ 0, & \text{otherwise.} \end{cases} \tag{1.1}$$

For example, given a contribution, $X(s, L) = [1, 2, 4]$, if the error threshold $\tau = 1$, then $D_1(X(s, L)) = [0, 0, 1]$ and the distinctness score $\mathbf{D}_1 = \sum D_1 = 1$. However, given a $\tau = 0.5$ for the same contribution, $D_{0.5}(X(s, L)) = [1, 1, 1]$, so $\mathbf{D}_{0.5} = 3$. The sum of the distinctness vector helps us decipher the total number of detectable lighting states at a given position and accounts for the sensor's resolution when assigning credit. We utilize this score in our simulations, where the virtual light sensors in simulation do not have the added lumen degradation and other measurement noise terms.

For our activity inferences, we account for $j = 3$ door angles (i.e., $\{0°, 45°, 90°\}$), for $m = 2$ doors. To find the distinctness score for each given application state $a$ is then:

$$\mathbf{D} = \sum_{p=0}^{2^n-1} \sum_{q=0}^{jm-1} D(X(s, L_p, a_q)) \tag{1.2}$$

Because choosing the minimum combination of sensor locations that can detect all possible applications is a known NP-Hard problem called the Minimum Set Cover (MSC) Problem, for simplicity, we visualize the inferable states for only the single-sensor scenario. Finally, we show an example set of light sensor locations that can capture our latent variable and the building light states using a known algorithm for the MSC problem: the Greedy Set Cover Approximation (GSCA) [47]. To see if a set of light sensors capture the dynamic building information, we can then use the entire application states as the *universe* of states to cover a threshold $< \tau$ against all other sensor locations rows for the same application state column to inform *membership*, to find the minimum subset of sensor locations required to cover the entirety of the application states. We target door usage because of its fixed nature. The opening and closing of doors have a large effect on perceived lighting by the sensors, with the potential to block off sections of lighting altogether consistently. Further, door usage is a deeper insight into space usage not expected by light sensors. Using the bathroom, bedroom, or living room naturally involves opening and closing lights and doors. In this case, where the lighting does not automatically turn off, being able to discern door movement even while all lights are on allows more granular inferences about space used to be made.

### 1.2.3 Assumptions

To simplify the analysis, external lighting sources outside the building were ignored. Possible external contributions of the outside can include but are not limited to direct sunlight, ambient environmental bounce light, and outdoor artificial lighting from other building units and street lights. Building materials are also simplified. We elaborate on these limitations in Section 1.5.3). Occupant behaviors of interest are limited to light switch behavior and door interactions. Specifically, in the former, the state space of 64 possible states is defined by which light fixtures are on. This is referred to as the “simple” example. The action space is defined by which light fixtures are turned on or off, but we are primarily focused on how signals sampled can be used to infer the correct state space since it is trivial to infer the actions between two states for light switch behavior. In the complex scenario, with a state space 576, the effects of opening and closing doors are paired with the possible light states. This is referred to as the “complex” example. For instance, there are nine possible configurations of door states if the possible angles of two doors (closed, half-open, and open), each with a different interaction with light, are discretized. For the complex scenario, the purpose is to demonstrate the possibility of conducting experiments in a simulation space that is not practical to conduct in the physical space. Another assumption here is that each light fixture's contribution is known. This is assumed because scaling by each light fixture's on-off state is polynomial time, compared to scaling to the cardinality of the power set of the light fixtures, which is exponential. Finally, we also assume in this chapter that the differences and similarities demonstrated in the simple scenario are sufficient to enable exploration in the complex scenario without exhaustively evaluating the complex scenario in the physical space. We do not intend these experiments to demonstrate, for example, how robust a sensing system built with only doors in mind could be resilient against a set of unseen occupant behavior patterns.

### 1.2.4 Experiments

As shown in Figure 1.1, even for a single light source, the angle of a door's rest state can significantly impact the final detected light signal. Using these methods, we conducted two experiments: 1) we deployed light sensors into the modeled residential setting to explore real-world challenges with detecting light states and fusing inferences from multiple sensors together, and 2) we modeled the residential building and simulated a set of dynamic building elements to analyze the number of light states and latent states that can be observed. The real-world study was conducted first as a sanity check; if we can arbitrarily place sensors using only human intelligence, it is unnecessary to utilize simulations to support the sensor placements. The simulation study was conducted to see whether the simulation space would corroborate with the findings found from our real-world study and also to gauge the minimum number of sensors that can capture a set of luminously disruptive behavior.

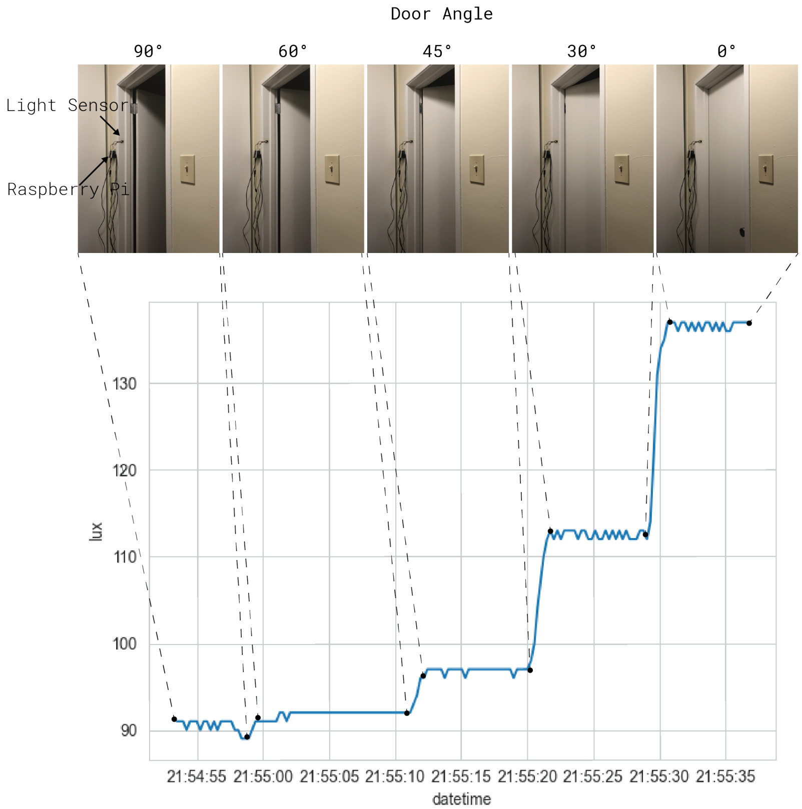


Figure 1.1: Relationships between building elements and lighting can be used to inform about changes in the physical environment.

### 1.2.5 Real-world Experiment

In the real world, we retrofitted all lamps in a residential apartment building with 800 lumens 10 watts Philips Hue A19 Lamps and grouped each set of lights to their corresponding luminaire. For example, two to three lamps can be associated with each light switch. We controlled each Philips Hue lamp using the Hue API and Python to limit the need to alter the light switches manually. We then used Raspberry Pi 4s connected with CQRobot TSL2591X light sensors with an effective sensing range of 0 to 88,000 Lux as our sensor, communicating using the I2C interface with the Pi. To account for the jitters, we used the mean lux values for three seconds as the baselines, collected at 4.7 Hz after our code changed the lighting configurations configuration $L$ for three seconds because of the changes in light intensity during state transitions. We used the proximity of the light sources as a guide to installing the light sensors in each location and orientation, as shown in Figure 1.2. Specifically, we visually looked for sensor positions on different walls, enabling the sensors to capture light from different light sources. Then, we moved our sensor system across each position and automated the lighting transitions using the Raspberry Pi, exporting the final data into a CSV file for post-processing. The final accuracy we report is explained in the results, where the ground truth is the input command we used to automate the light states. To install an initial set of sensors, we identified walls in the testbed that all the luminaries can reach and then placed seven sensors on those walls, seven feet above the floor shown in Figure 1.2. The idea is to see if it is possible to install sensors in locations that avoid the noise in the data caused by shadows in human traffic, the reflection of furniture, and other LEDs from appliances and objects. This allows us to accurately detect the lighting state of the building for the static open door scenario before we dive into permutations of doors in the simulation experiment.

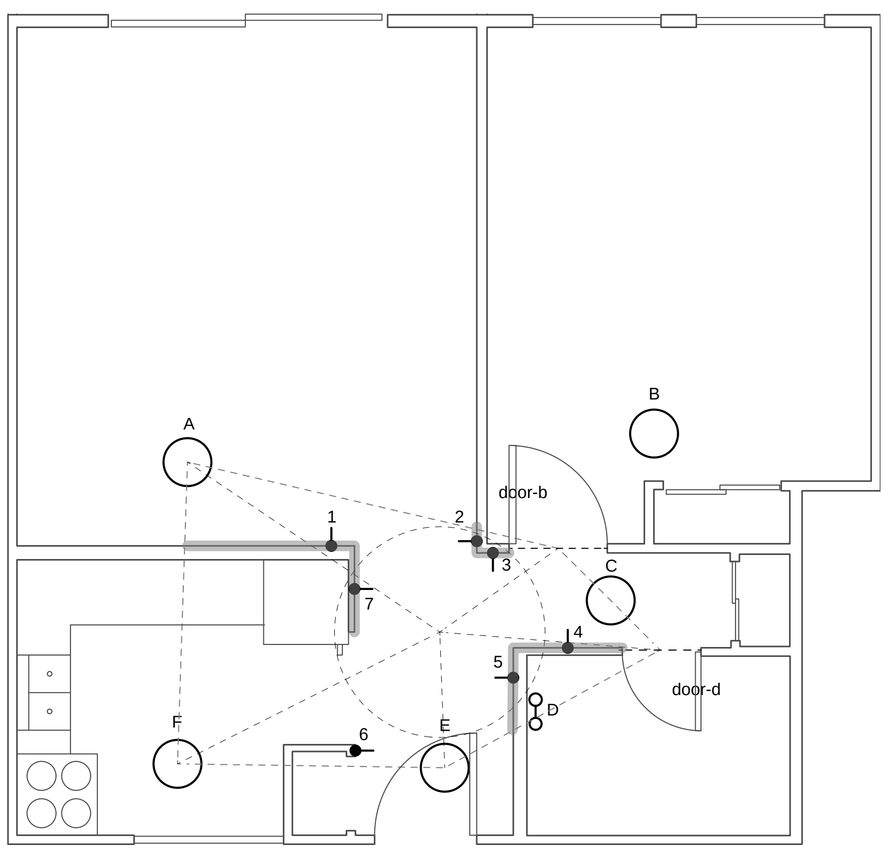


Figure 1.2: Building and Lighting Layout. The numbers (1,2,3..., 7) denote the light sensor positions, and the alphabets (A,B,C,...,F) denote the light sources. The dashed lines denote geometric boundaries we used to search for our candidate walls to install the sensors.

### 1.2.6 Simulation Experiment

In a simulation, we propose to utilize grasshopper to parameterize multiple door movements and simulated each of the allotted door angle combinations using Rhino [48]. We can then use the grasshopper plug-in honeybee [49] as we have in our previous work [50] to extract the lighting contributions at each sensor point. In essence, the plug-in acts as a middleware that takes the building geometry and photometric lighting files (IES files) that describe the geometric intensity distribution of light and passes them into the lighting render engine radiance [51]. The experiment will be assumed to be undertaken at night with covered windows to avoid external light sources. We can then use the simulation default material for all wall, floor, ceiling, door, and window objects. For the photometric lighting files, we propose to use generic flush-mounted dome lights for fixtures A, B, C, E, and F and a generic wall-scone fixture for D. We can then automate the inputs to the grasshopper workflow and export text files we convert for post-processing using Python modules inside grasshopper. The process will then return a ray-traced light rendering of the building for each virtual sensor position defined. We can then extract each lighting value for post-processing by passing it through the Perfect Sum solvers and the GSCA algorithm. This workflow will enable us to explore potential lighting differences in the overall environment without needing to physically place new sensors, change which luminaires are on and off, and move the doors in the real world.

## 1.3 Results

We first conducted a feasibility study to explore the potential challenges of collecting sensor data manually. We then used the experiences we have learned to inform our modeling of the building in simulation space. We summarize the real-world experiment and the simulation experiment below.

### 1.3.1 Real-world Experiment

Figure 1.4 shows the spread of individual accuracy for those locations. The median accuracy of inferences was generally above 80%, with location 6 having the lowest median accuracy. We suspect this due to the direction the sensor is facing being directly opposite to luminaire F and reflected lights off the wall being less potent than the drop off of intensity due to distance. Overall, we still found it possible to disaggregate all possible light states using a single sensor (i.e., a light sensor located at position 4), but many factors can contribute to the imperfect inferences. For example, when an occupant closes the door to any room, the information about the light state in that room is lost to sensors outside. We also experience situations where the inferred light state of the signals does not align with each other, because the noise in the environment and sensor was larger than the resolution required to differentiate the states. For example, the lighting contribution from two light sources can be the same (e.g., 2+5 and 3+4 equal 7), leading to ambiguous readings. This led us to utilize a voting mechanism to reduce the overall error of the system, which is further described in the methods section.

When we automated the lighting states, we also realized that, unlike what may happen in simulations, the switching on and off of lights in the real world is not instantaneous. Specifically, when turning on the lights, there is a distinguishable start-up time when the light starts dimmer after the switch is flipped and approaches its final brightness after a delay. Furthermore, when retrofitting the luminaires, we noticed that not all of the lights were using the same bulbs and that likely more frequented areas had bulbs that were more frequently replaced. This suggests that keeping track of the light usage in a building can also be helpful to track which lights might need to be replaced and account for the lamps' brightness degradation over time. Finally, we found that the number of available sensors, microcontrollers, and outlets also limits how many positions can be tested simultaneously. In addition to purchasing a long extension cord to move our sensing apparatus across the building, we also adhered to moving sensors, installing sensors, running through all the light states, and uninstalling sensors to sense each position.

### 1.3.2 Lighting Simulation Results

Utilizing the information we've learned from the real-world experiment, we developed a lighting simulation that deviated from traditional lighting simulations to explore locations where static sensor installations can give us the most information. Specifically, in our lighting simulation, we included *walls* instead of using the traditional work plane–an imaginary plane set at the level of a desk where work is done–because we are not interested in the utilization of the space but rather in the ability to install real static sensors and detect different behaviors in the space (e.g., light switch behavior). The results from the simulation are shown in Figure 1.5. Figures 1.5a through 1.5i represent the different states of dynamic building elements, with the heatmap showing the

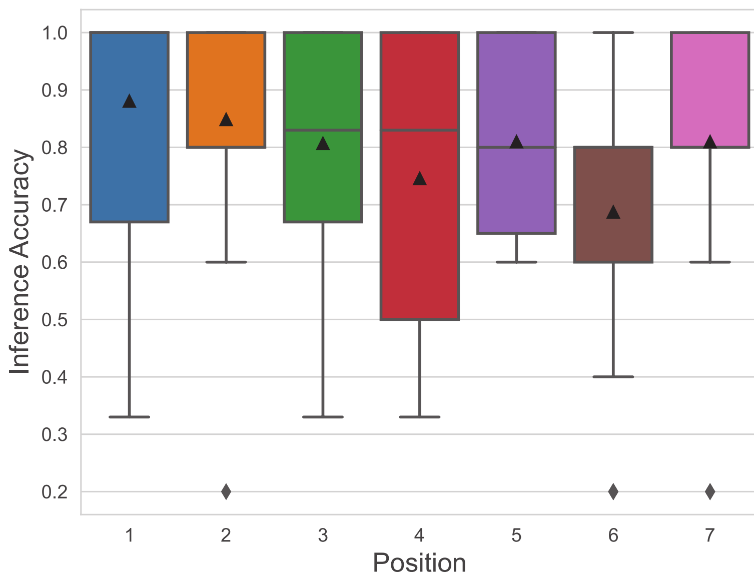


Figure 1.3: Inference accuracy of physically sensed values

Figure 1.4: In Figure 1.3, we observe state inference accuracy for the physical experiment. Each box plot shows the seven positions for 64 possible lighting states for manually selected positions. The triangles mark the mean accuracy for each sensor location, while the line across the middle of the box marks the median accuracy.

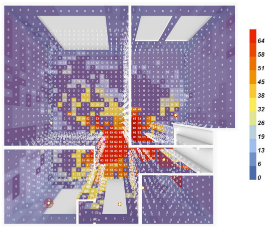

(a) Bed Door 90°, Bath door 90° (open-door scenario)

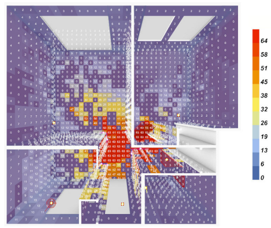

(b) Bed Door 45°, Bath door 90°

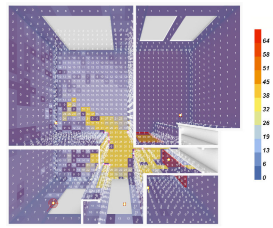

(c) Bed Door 0°, Bath door 90°

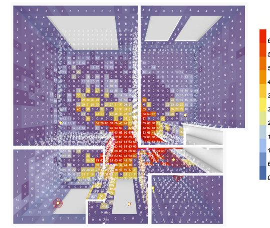

(d) Bed Door 90°, Bath door 45°

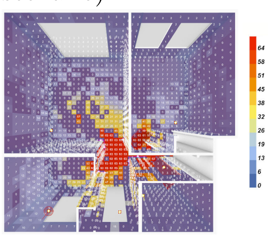

(e) Bed Door 45°, Bath door 45°

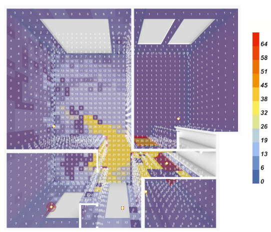

(f) Bed Door 0°, Bath door 45°

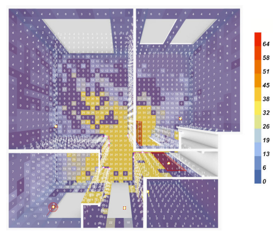

(g) Bed Door 90°, Bath door 0°

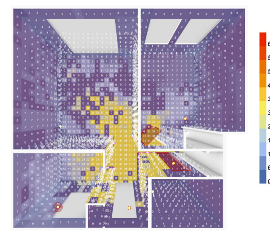

(h) Bed Door 45°, Bath door 0°

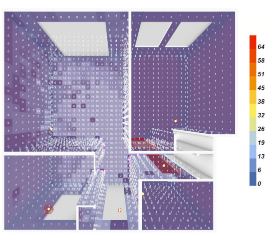

(i) Bed Door 0°, Bath door 0°

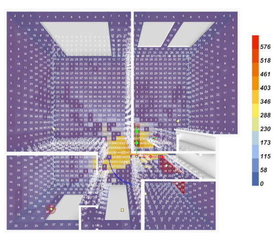

(j) Aggregate latent variable privacy $D_{0.01}$

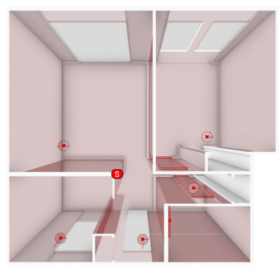

(k) One light sensor can sense all 64 light states for the open-door scenario using GSCA

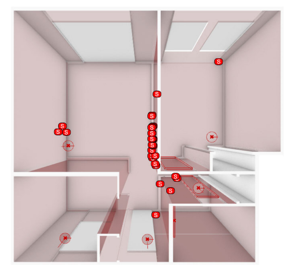

(l) 31 sensors are needed to sense a total of 567 states using GSCA for the dynamic-door scenario

Figure 1.5: Compilation of the distinctness score $\mathbf{D}_{0.01}$ heat maps for each building configuration. In addition to the resolution of the sensor, the building's physical configuration can also systemically alter the inferable states of the sensor set. Note that the scale for Figures 1.5a through 1.5i is on a scale out of 64, while the aggregated scale on Figure 1.5j is out of 567.

distinctness score $\mathbf{D}_{0.01}$ value from 0 to 64, representing the number of collective states a sensor placed at the location can detect. Figure 1.5a represents a typical light sensing simulation, where the movement of additional building elements is not considered. In this scenario, many positions in the middle of the room where all the lights can reach can be used to infer the light state of all luminaires. From Figure 1.5b to Figure 1.5i, we show that these informative middle positions diminish as the doors close and block out lighting contributions from different sources. In Figures 1.5d and 1.5e, when all doors are still partially open, we see a slight reduction in areas that can still make all targeted inferences. However, in Figures 1.5c, 1.5f, 1.5g, 1.5h, and 1.5i, we see the total possible inference visibly diminish by half to three quarters.

Figure 1.5j represents one of the most informative single locations accounting for all possible door states. As much as there are informative locations in the middle, we also observe there to be spots of lost information, dark zones that have lower inference potential. This is a result of the clashing of light contribution combinations that lead to ambiguous readings. Comparably, while Figure 1.5i has lower number of possible inferences in the middle, there are also less informative "dark spots" as a result of collisions. Figure 1.5k shows how Greedy Set Cover Approximation (GSCA) found a single location that can detect all light states, but Figure 1.5l shows that the previous best location is no longer valid when we account for the opening and closing of doors. More specifically, we can see a set of sensors being placed deep into rooms away from noisy areas in the center, which help to disambiguate the readings when the lighting signals are muddied by reflection and attenuation in the center area.

## 1.4 Evaluation

To further evaluate the validity of the proposed method, three different value functions for sensor selection were explored, as well as improvements on how the voting scheme can apply to states larger than light states. Specifically, three different value functions were explored for sensor placement: 1) random selection, 2) value based on maximal increase in observable states, and 3) manual expert selection. An updated merge strategy was used for evaluating the number of total states covered, where the value assigned to each state for a sensor position corresponds to the $v_s = \frac{1}{n}$, where $v_s$ is the weighted vote variable, and $n$ is equal to the total number of ambiguous states (states where the lighting values clash). Compared to the previously merge strategy proposed in Section 1.2.1, this updated method properly represents the confidence of assertions for each state. It simplifies ways to merge information when further states are to be observed. For example, if a sensor at a location can detect each possible building state with a unique lux value, and we are searching for 576 states, for 2817 possible sensor positions, we create a 576 by 1,622,592 matrix, representing each unique (location, state) combination, and their confidence for the knowability of each unique state. If, for a certain building state configuration (i.e., for a row), two configurations result in the same value, the row will have all zeros and two 0.5 values for the readings. If three states results in the same value, three values in the row will be 0.33, and so on. Using this construction, to compute the final coverage of each state given a set of sensor locations, we sum the confidence vectors for a given state and observe if the index of the largest confidence cell in the voting vector corresponds with the index of the 1 in the ground truth vector. Figure 1.6 compares the average of 100 random selections with the random selection policy with the sensor coverage and sensed state value function using GSCA, showing how GSCA can reliably outperform the two other placement strategies.

## 1.5 Discussion

### 1.5.1 Key Takeaways

The key takeaway from our study is that: *sensor position is important, and simulations can be used to quantify just how important position is.* By quantifying inferable information in simulation, building operators can adjust the privacy-utility spectrum for where sensor installation should occur before deployment. The quantity and location of sensors can be altered to purposefully remove possible inferences based on the physical attributes of the environment. Even after deployment, simulation elucidates where the current installation is on this spectrum and ways to navigate this trade-off. In our scenario, this trade-off between privacy and utility is exactly the distinctness score, a quantifiable value between zero and the total number of states we are considering. Towards answering **RQ1**: we found that the informativeness of the sensor location also hinges on the sensor's resolution. Paradoxically, the more contributions from different light sources sensed by a sensor-location pair, the more chance there can be ambiguous readings in the perfect sum solver because of the number of possible combinations and jitters in the sensor signals. These jitters, shown in

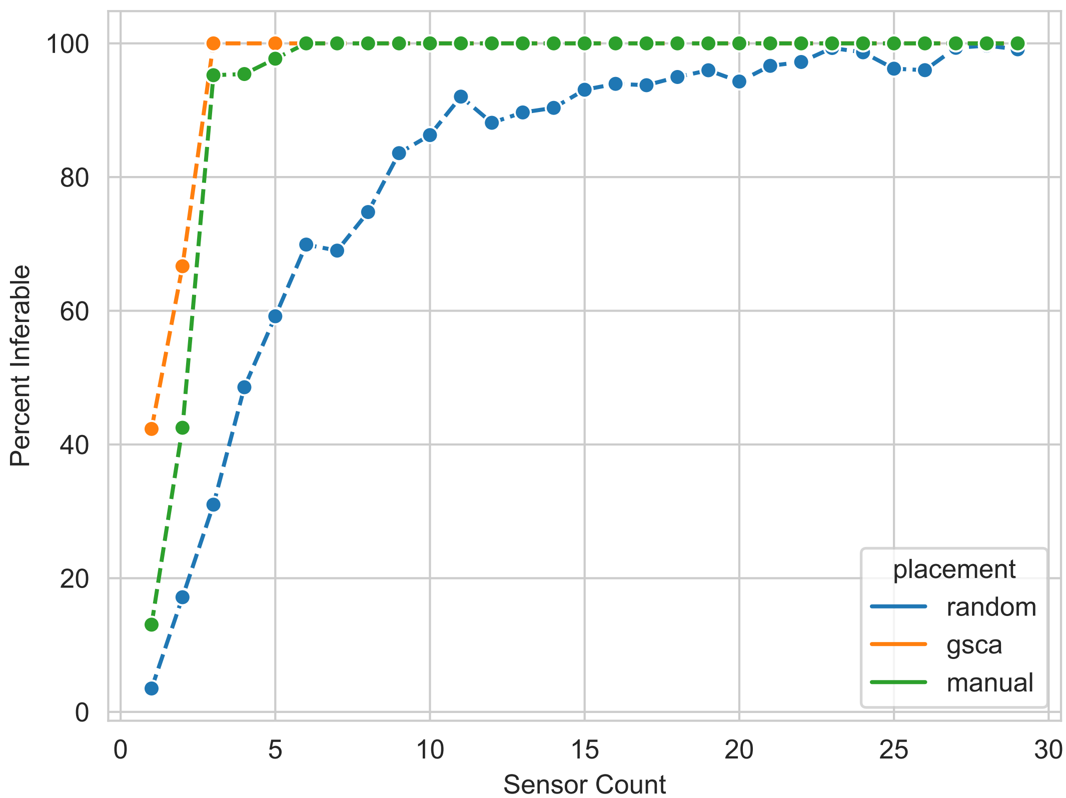


Figure 1.6: A comparison of the number of sensors required to infer all 567 possible building states for the complex scenario and the weighted probability strategy (the nine states in the dark remain ambiguous). Using GSCA with the weighted-vote merging, only three sensors are found to be necessary to infer all states in the complex scenario. The algorithm can perform marginally better than how a human agent would place the sensors. An example of the manually selected location for a three-sensor scenario is shown in the Appendix, Figure A.1).

Figure 1.1, signify a vital distinction between real-world and virtual sensors in simulations. Adding more sensors in the physical world does not, by default, increase the accuracy of the final light state inference. The requirement for accurate sensing relies on the majority of the votes cast being accurate. To know the lowest level of resolution permissible for accurate sensing before purchasing a sensor, simulations can be a useful tool to assist with planning. In the simulation, the resolution of the virtual sensors is deterministic under the same parameters, and adjustments can be made to simulate different sensor resolutions by introducing additional noise terms. Further, by adding more information about the activity, simulations can be improved to account for different sampling frequencies by discretizing sample points based on a continuous response function. With sufficient computational resources, permutations of different sampling policies and sensor descriptions can be used to optimize the inference accuracy, redundancy, and efficiency. With simulations, researchers can achieve more intrusive inferences with fewer sensors, fewer samples, and less energy compared to without simulations.

Towards answering **RQ2**: we found that utilizing lighting simulations with a formulation of *Perfect Sum Problem* with the *Set Cover Problem* allowed us to quantify the minimal number of light sensors that are required to capture the light state of the building, including modification of the doors. We found that as long as the sensor is placed in a location where all light sources can reach and result in a different contribution, a single sensor is theoretically enough to infer all of the possible light configurations in the building if the resolution equivalent $\tau$ is sufficiently small. However, as doors are introduced that can block off lighting contributions from other luminaires, the minimum number of sensors required to sense the lighting state of the building becomes equal to the total number of independent zones. For example, three separate rooms require at least three sensors to detect the lighting state, regardless of the number of luminaires in each room. The minimum number of sensors required increases from one to 31 when we accounted for the movement of the doors as seen by Figure 1.5k compared to Figure 1.5l. This indicates that even simple residential buildings with no dimmers can result in complex luminous environments if commonplace building elements such as doors are considered. Unless researchers had thousands of sensors, placed at every inch of the space, they wouldn't be able to practically test the entirety of the space at once.

Finally, simulations can be much faster at the search for optimal positions. In our experiment, it took us roughly 30 minutes to test one position in the real-world, where as in simulation space, we cover roughly one sensor location every 0.11 minutes of simulation (about 2,800 points can be calculated every 5 minutes for the 64 different light states). This approximates a 270-times increase in efficiency when using simulations to test positions compared to real-world testing, not accounting for set-up time in either scenario. Because simulations also do not require the researchers to be physically present in space, simulations hold a distinct advantage over manual testing as an important step to enhance physical sensor deployments. With increased simulation scenarios that incorporate human movement and other modalities such as noise and HVAC, the digital space will increasingly become more critical not just for sensor positioning but also for a large myriad of selection tasks. From what type of sensors to deploy, at what frequency to sample data, what information they provide at what different times of day, and seasons, with what different levels of occupant activity to expect, and with which soft sensors[52] to combine and make inferences with, simulations will take an increasingly important role in controlling and testing the scope of inferences in buildings. Methods like this demonstrate that simulations have the potential to serve as a stand-in for domain experts. If experts can digitize the knowledge of specific sensor placements for building commissioning, for example, they can enable accessible and code-compliant occupant privacy protection designs while further providing availability to interact with other simulations using the same building model.

### 1.5.2 Broader Impacts

The work we have completed represents both predictive model tasks, where we anticipate the use-case of the occupant before installing the sensors, but also a step towards reducing the gap between the digital twin and the original twin. Sensor installations can take advantage of more than just their placement concerning common building elements such as floors, walls, and doors. Sensor placements can also benefit from being aware of other sensors in the context. The work-flow we demonstrate allows for explorations in designing buildings that can be more effectively commissioned with fewer sensors. Developing metrics to quantify possible inferences also provides an additional avenue for designers and researchers to consider user privacy. For instance, there could be dedicated "silent zones" where sensors cannot detect any occupant activity as protected by the laws of physics. Simulations can be an effective tool to compete with the scale of sensor developments because they can protect their users from scenarios that have yet to happen and inform and adjust models using real data. While the digital and original twins divide is shortening, we consider their distinct identities to carry certain benefits. For example, digital twins can be operated "offline" to explore reactive and predictive scenarios that inform on the optimal corrective action without interfering with the operations in the real system. However, this is not to say that simulations will not also be an increasingly important part of the operations in real-time systems. Decisions to navigate the potential split incentives between the building operator and the occupant will likely depend on circumstance and might require routine updates to support new management and new tenants. A building operator might install light sensors in all rooms to avoid the complexity of inference but expose information about the occupant's kitchen, living room, and bathroom use they consider private. Similarly, an occupant might install one sensor to understand lighting in the living room but accidentally leak the lighting states of other rooms to the building operator. Decisions regarding which data should be hidden for privacy or which data should be available for utility would require coalescing of ideologies regarding ownership of space, ownership of data, ethics, among other considerations. Regardless of the perspective, the first step is showing in a data-driven and reproducible manner where a sensor installation theoretically lies on the privacy-utility spectrum.

### 1.5.3 Limitations

One limitation of this study was that it was conducted on an older residential unit, where buildings might not reflect a more modern understanding of the efficient usage of light fixtures. The privacy implication of sensor use in public and semi-public situations such as offices and libraries could have a broader impact on the number of people affected. Another limitation is our assumption of snapshot views in simulations. We did not incorporate time (analysis of signals instead of values) such as through bulb response functions [53]. This can limit the ability for simulations to reflect the time in-between snapshots and the additional inferences that can be drawn due to more realistic sample rates. Another limitation is that we only collected light intensity levels but did not look further into other properties of light, such as lighting colors. We suspect colors can be an important avenue to disambiguate lighting signals further. For instance, the individual lights might be able to be first filtered by color, reducing the total number of possible combinations and collisions that could happen. Another limitation is that the involvement of some furniture has the potential to alter the indoor environment. For example, having a mirror on the wall can drastically alter the luminous environment, similar to having light-absorptive materials on the floors. Finally, our work does not address the difficulty of constructing a representative building model nor the potential diminishing returns of modeling the environment in more realistic detail. While numerous benefits can be achieved with an informative building model, the cost of building a representative model can eventually outweigh the demand to protect an occupant's privacy. The cost of the building model can be further exacerbated when more computation time is required to calculate physical interactions in the space, such as increases in the number of bounces for lighting simulations or the number of particles in Computational Fluid Dynamics (CFD) simulations.

## 1.6 Conclusion

We demonstrate a theoretical framework for indoor activity inference selection through simulation experiments and real-world sensor placements. We show how simulations can quantify inferable occupant activities using a distinctness score and how to find a mathematically minimum set of sensor positions required to detect them by applying the concept of set cover. The resulting metrics to quantify distinguishable activities enable future sensor deployments to consider building geometry better and limit potential sensor data overreach. We anticipate using sensor positioning paired with building simulations to grow as an essential technique for researchers to navigate the privacy-utility trade-off for the smart buildings of tomorrow.

# Chapter 2

# Utility

*The relationship between environment and health outcomes are hard to quantify because of the inherent complexity of humans, the time duration to establish significance, and the large variety of environmental variables that have an effect. This chapter demonstrates how contextual data recorded by the nurses can support automated tracking of medical events behavior and narrow down the required environmental variables and time windows. Further, this chapter uses statistical methods and signal processing to show how environmental variables are related to patient sleep disruption.*

## Exploring Environmental Signals to Analyze Hospital Patient Sleep Disruptions and Medical Events

Environmental factors, such as lighting and noise, have a history of disrupting patient sleep in hospitals. However, until the recent advent of affordable ubiquitous environmental sensing techniques, it was not feasible to conduct the long-term recording of both widely prevalent and less available (such as volatile organic compounds (VOC)) disruptive environmental factors. Quantifying the impact and timing of sleep disruptions owing to hospital environmental conditions has greater significance than merely knowing its existence. This information can help nursing teams better administer care, improving patient sleep and recovery. Further, environmental data streams have been shown to provide additional insights about the medical environment, such as occupant presence or patient fall events. In this study, we deployed commercially available environmental and actigraphy sensors in five hospital rooms for general medicine and geriatric care at a University Hospital to detect sleep-disruptive environmental factors (lighting, temperature, humidity, $CO_2$, VOC, PM2.5, and noise levels) and classify five classes of medical events using of a subset of the high-resolution environmental data as the input features. We analyzed this data alongside the recorded medical data over 169 days and 38 patients to understand the relationship between medical events, environmental factors, and patient sleep disruptions. We utilize visual analysis, generalized Linear Mixed Models, and Decision Tree models to identify environmental patterns. Furthermore, we establish meaningful connections between environmental signals and patient sleep disturbances while effectively classifying medical events through the use of environmental signals in a hospital environment.

## 2.1 Introduction

Sleep is essential for patients in the hospital, as it promotes healing, improves mental health, and reduces the length of hospital stays [54, 55]. Without enough sleep, patients may experience increased pain sensitivity, delayed healing, and a greater risk of infection. Therefore, healthcare providers are incentivized to promote good sleep hygiene and appropriate sleep support to improve positive outcomes for their patients. However, despite the positive benefits that adequate sleep provides, patients typically sleep on average less than 1.5 hours in hospitals compared to at home [56]. While environmental factors are now considered just as crucial as a patient's physiological changes as a potential cause for poor patient sleep in hospitals [57], descriptions of these disruptions vary. Kulpatcharapong et al., for example, report how pain, light, and sound are the most significant predictors of sleep quality for hospitalized patients [58] while Tan et al. divide the interruptions based on patient characteristics, hospital routines, and the hospital environment [59]. While the principality of disruption varies across the studies, these works highlight the essential and interconnected relationships between patients, hospital routines, and environmental factors. Similarly, the possible interventions to improve patient sleep in hospitals are multi-faceted. Ricio et al. divide sleep-disrupting factors into environmental, illness-related, and sleep-promoting factors into pharmacological aids and non-pharmacological aids [60]. Environmental factors disrupting sleep can be noise, lighting, or interruptions from hospital staff, whereas illness-related can be pain from operations, such as surgery. Pharmacological aids correspond with sleep medicine, whereas non-pharmacological aids correspond with ear plugs and sleep masks. The persistent, spontaneous, and numerous ways a hospital environment can negatively impact sleep indicate that we can no longer rely solely on the observation of caregivers to ensure a sleep-promoting environment for the patients. In addition to carefully monitoring caregivers, hospitals can benefit from deploying IoTs and integrating the following data with existing collection pipelines to provide continuous and objective observations. Prior works have demonstrated how IoT deployments can successfully inform about environmental variables that perturb sleep, but many of these studies were conducted in residential areas [61, 62], which embodies a different context than a hospital setting. By testing environmental IoT in a hospital environment, researchers can better inform the maintenance of a comfortable sleep environment in a hospital and, ultimately, better patient care. We further summarize prior IoT solutions exploring sleep in Related Works (Section 2.2).

To track physiological components of sleep continuously, researchers have utilized actigraphy in place of polysomnography (PSG), or electroencephalography (EEG) for more scalable, long-term assessment [63, 64]. Actigraphs are devices typically worn on the wrist that can record a patient's sleep state via movement and are based on the idea that people tend to move less while asleep [65]. Continuous tracking of sleep and environmental variables using IoTs, has also been carried out via deployed static environmental sensors in various residential, commercial, and outdoor settings [1, 27]. Besides the increased awareness of tracking indoor environmental quality attributes in general because of its significant relationships with occupant health [12], the parallel tracking and fusion of different but equally significant data streams have also received renewed attention and interest [66]. Like temperature, humidity, and lighting, researchers have also found relationships between sleep and other indoor air environmental quality metrics, such as $CO_2$ and VOC [67]. With the influx of evermore data to consider, we look to merge the different data streams and identify the principal contributors to poor sleep.

In addition to relationships with poor sleep, environmental variables have been successfully used in past literature to detect granular occupant activity. A well-known example, the concept of "Sensing-by-proxy," demonstrates how $CO_2$ data can be utilized to detect occupancy [33]. Beyond occupancy, VOC has been demonstrated to be a proxy for cooking activity [68]. Environmental sensor streams have been successfully paired with machine learning in hospital environments to support numerous applications, from fall detection [69] to irregular environmental event detection for pregnant females [70]. Since environmental variables have been showcased to detect granular occupant behavior, we explore the possibility that environmental patterns in a hospital room can be utilized to classify medical events activity to see if environmental sensors can bring additional utility besides the commissioning of the environment.

In this work, we pursue an exploratory study of patient sleep quality in general medicine and geriatric care unit hospital to observe long-term contextual challenges and opportunities to tackle the issue of combining the different types of data to help improve the sleep quality of patients. Further, we analyze to explore the possibility of using environmental signals to provide insights about the medical procedures to showcase the possibility for automated checking of manual data entries. This work will focus on methods to quantify sleep-disruptive environmental factors with input from physiological and medical events data. We utilize actigraphy as our baseline and environmental and medical events data as our predictor to answer the following research questions (RQ):

- **RQ1**: What are notable environmental patterns across different time scales that can negatively impact patient sleep in hospitals?
- **RQ2**: To what extent can environmental variables and medical events data be predictors for patient sleep disruption?
- **RQ3**: Which combination of time-windows and environmental signals is the most accurate in predicting sleep-related medical events?

We structure the rest of the paper as follows: In the related works (Section 2.2), we review the literature investigating the different environmental fields that affect sleep and prior methods that captured and related sleep data with the environment. In methods (section 2.3), we detail our participant recruitment, data acquisition, and how we evaluated our study using a Generalized Linear Mixed Model (GLMM), Devision Trees, and Random Forest models. In the results (Section 1.3), we showcase the variance, including our random effects, across percentiles of patients and rooms, and scores of the machine learning model with various inputs. In the discussion (Section 2.5), we discuss the potential causes for the observed differences, the importance of window length, the highest performing combination of environmental signals, the work's limitations and future directions, and summarize the contributions and broader impacts of the work.

## 2.2 Related Work

The relationship between environmental factors, nursing intervention, and sleep has been explored in various prior studies in different combinations. This section summarizes some important relationships, technologies, known comfortable ranges and considerations.

### 2.2.1 Known Environmental Set Points

In this section, we summarize a list of indoor environmental quality metrics that we track and prior works that help describe comfortable ranges that are conducive to sleep.

**Lighting** Measured in lux (one lumen per square meter), lighting has been found to have a direct relationship with select ganglion cells influencing a person's natural sleep-wake cycle, or circadian rhythm [71]. Specifically, our eyes act like a blue-sky indicator, and exposure can cause delayed melatonin secretion, increasing alertness during the daytime but inhibiting sleep when exposed at night [72]. As the number of electronic tablets, television, or computer becomes more ubiquitous [73], nighttime lighting exposure also grows as an important environmental variable to scrutinize for better sleep. Crucially, proper exposure to light during the day is equally important to darkness at night [74]. Furthermore, in addition to the timing, the amplitude of the lighting also amplifies this delay, up to about 1,000 lux [75]. In summary, past literature mark the amplitude and timing of light (and its absence) as an important environmental factor to track across time anywhere people sleep, not just in hospitals. We mirrored a previous study and marked 10 lux as a cutoff point for tolerable lighting during sleep [76].

**Noise** Noise, or unwanted sound, is widely known as an issue in hospitals and has become more of an issue in modern hospitals because of the increasing number of devices and monitors with audio cues [77]. It can be measured in decibels (dB) or A-weighted decibels (dBA) – which is dB weight adjusted as perceived by the human ear. Noise is well known to negatively impact sleep. A prior review by Hume et al. demonstrates the difficulty of survey-based approaches for noise disruption research because noise events can be too short to be consciously perceived by the subjects [78]. Further, Hume's review also notes potential homeostatic mechanisms for internal monitoring and control of waking arousal, where most (90%) of the noise-induced awakenings merely replaced awakenings that would have occurred spontaneously. These findings suggest that even with continuous sensing, only 10% of the disruptive noise observed would significantly impact the overall sleep quality. Regardless, we follow the World Health Organization (WHO) guidelines that recommend less than 40 dBA at night for good sleep. In comparison, the Environmental Protection Agency (EPA) expresses a 24-hr exposure limit of 55 dBA to protect the public from all adverse effects on health and welfare in residential areas [79]. The EPA further notes that it is not the average noise level that is important and that a bedroom with an average noise level of 35 dB with no instantaneous peak levels substantially higher would be more conducive to sleep than a room with an average noise level of only 25 dB but which stillness is disturbed by occasional shrieks. These findings advocate for tracking noise in the hospital environment and motivate us to incorporate methodologies such as change point detection to monitor large changes in the distribution of environmental variables in conjunction.

**Temperature and Humidity** While temperature and humidity are measured separately, they are often included in the same measure regarding thermal sensation because sweat evaporates more slowly when the air is saturated with water. Researchers in environmental health utilize a combined measure such as a heat index to calculate perceived temperature by adjusting device-sensed temperature values based on humidity readings to more closely align with human thermal sensations [80]. Separately, exposure to low humidity can cause sensory irritation in the eyes and upper airways, while exposure to high humidity carries with it risks of fungal dispersion [81]. Exposure to higher temperatures can be conducive to higher quality sleep but also cause more sleep fragmentation during sleep [82]. Manzar et al. note an increase in longest wake episodes, a decrease in total sleep time, and a decrease in sleep efficiency with the seasonal increase in bedroom temperature and relative humidity [83]. For temperature, ranges between 18 and 28 degrees Celsius have been reported to be optimal for sleep; for humidity, ranges between 40% to 60% have been reported as optimal for sleep [84, 85]. We adopt these reported ranges for our analysis, using individual ranges for visualizations and apparent temperature for our modeling.

**Carbon Dioxide** Measured in parts per million (ppm), $CO_2$ has been positively correlated with increased sleep awakenings, and poor subjective sleep quality [86]. During wake time, increased $CO_2$

concentrations are also associated with impaired cognitive functions [87], making it an important environmental variable to track in hospitals for patients and caregivers. Additionally, because humans naturally exhale $CO_2$, it has also been used as an effective proxy for counting the number of occupants in a room [33]. This indicates that negative sleep quality observed during elevated $CO_2$ ranges can also be attributed to other movements in the room, such as visitor activity or nursing events. Compared to humidity, light, or noise, $CO_2$ and other air quality metrics can be an especially important environmental variable to track because they cannot be as readily detected by humans but can still affect sleep. Acceptable ranges of $CO_2$ have been reported between 400 and 900 ppm [88].

**Total Volatile Organic Compounds (TVOC)** As defined by the EPA, Volatile Organic Compounds are "compounds that have a high vapor pressure and low water solubility" that can be 2 to 5 times greater concentrations indoors compared to outdoors [89]. Total volatile organic compounds, or TVOC, are groupings of VOC to simplify reporting. In the pages following, we refer between TVOC and VOC interchangeably. VOCs are in many chemicals in paint manufacturing, pharmaceuticals, and refrigerants. They can cause eye, nose, and throat irritation, headaches, loss of coordination, and nausea when exposed to them. Prior works have demonstrated people can be routinely exposed to VOC via inhalation and trans-dermal pathways related to off-gassing by bedding [90]. Fritz et al., in a study of occupant residential homes, found a relationship between VOC and increased sleep time [91]. VOC has also been shown in studies to be a result of personal care products [92], demonstrating its potential as a proxy to occupancy detection similar to $CO_2$. VOCs are also measured in parts per billion (ppb). An upper bound of 500 ppb was suggested as an upper limit for comfortable sleep [93].

**Particulate Matter 2.5 ($PM_{2.5}$)** $PM_{2.5}$, measured in micrograms per meter cubed, are tiny particles in the air that are 2.5 microns or less in width. Short-term exposure of $PM_{2.5}$ has been linked with "with premature mortality, increased hospital admissions for heart or lung causes, acute and chronic bronchitis, asthma attacks, emergency room visits, respiratory symptoms, and restricted activity days", while long-term exposure to $PM_{2.5}$ has been linked to other adverse health effects such as "premature death, particularly in people who have chronic heart or lung diseases, and reduced lung function growth in children" [94]. Besides health effects, prior work has also showed a link between increased $PM_{2.5}$ exposure and reduced total sleep time [95]. The EPA suggests keeping an annual standard of 12 µg/m3, and a $PM_{2.5}$ values under a 24-hour fine particle standard of 35 µg/m$^3$ [96]. We adopted the 12 µg/m3 standard in our traction of deviations for any epoch. For convenience, we summarize all the above-mentioned environmental ranges we selected in Table 2.2 of the Methodology section.

### 2.2.2 Prior IoT Solutions to investigate sleep and Medical events

With the rapid advancement of sensing technologies, we now have access to devices that allow us to collect more data using smaller devices, allowing the capture of both user and environmental data on the same device. Prior work such as SleepGaurd demonstrates that it is possible to capture rich information, including body posture and movements, acoustic events, and illumination conditions, by using just a smartwatch [97]. Similar detection schemes using smartphones have been demonstrated by Toss 'N' Turn, which found that features of noise and movement were useful to infer sleep quality [61]. Wahl et al. also showcase a power-efficient Expert Model-based smartphone app for realistic everyday sleep monitoring [98]. With ever more effective and efficient sensing and analysis schemes, why haven't these technologies seen more pervasive use in hospitals? One of the main challenges with smartwatches that capture high-resolution data from multiple modalities is that they consume more energy and thus require more frequent upkeep regarding charging and maintenance. This maintenance imposes additional burdens on the caregivers, not to mention they can become too frequent and disrupt the very comfort the devices were hoping to protect. Another challenge with the patient population is due to the short duration of stay, resulting in insufficient data to build a useful model. In a study of users in a residential area, Toss N' Turn notes that at least three days of ground truth data to train an individual model and three weeks of data to train a general model is required [61]. However, during our observational period, we found that patients often do not sleep more than two nights in the hospital. We thus explored the potential of statically placed environmental sensors that do not have to be reinstalled every time there is a new patient. These additional burdens can sometimes become unsustainable for the nursing staff. We explore the practicality of deploying static environmental sensors into a resource-constrained hospital environment to observe the potential benefits of the observed signals and the perceived reduction in necessary maintenance efforts.

### 2.2.3 Research Gaps

While prior works have found the different environmental factors to be sleep-promoting or sleep-disrupting, we purvey that sudden environmental changes–even within known comfortable ranges–can still be disruptive for sleep or indicative of a medical event or nurse visit. Furthermore, the patient population admitted to a hospital bed can be unpredictable to adjust to and draw inferences from. Compared to previous studies where users do not sleep in the same rooms, we can explore the context of shared hospital space and scale the utility of installed devices for patients in the most need. Further, we can capture information in a hospital setting without patients to observe the *potential* that the semi-public room can be hospitable to patients not yet admitted. Finally, deploying sensors with the support of hospital staff and nurses enables us to leverage recorded medical data to arrive at a more holistic assessment of contextual variables that can affect a patient's sleep. We did not find studies that explored the search-window space necessary for environmental variables to help classify medical events, nor the minimal combination of environmental variables that can lead to an accurate classification, As such, we explore the practical potential of longitudinal medical, environmental, momentary environmental, and actigraphy changes in a hospital setting.

## 2.3 Methods

This section is divided into multiple subsections, including participant recruitment (Section 2.3), data acquisition (Section 2.3.1), data preparation (Section 2.3.2), and evaluation methods (section 2.3.3).

**Participant Recruitment** The study was conducted at a University hospital's General Medicine and Geriatric care unit. The inclusion criteria for the study were: Subjects recruited 1) must have been assigned to one of the five single-occupant study rooms, 2) must have been admitted to a study room within the last 72 hours, 3) must be at least 18 years old, and 4) must be able to provide verbal consent. We admitted 39 participants (16 females, 23 males) with a mean age of 56 and a standard deviation of 16 years. The participants stayed at the hospital anywhere from 2-69 days, with a mean of 7.5 days and a standard deviation of 9.2 days. Patients sleep anywhere from 2 to 22 hours per sleep episode, with a mean of 7 hours and 56 minutes and a standard deviation of 4 hours and 40 minutes.

### 2.3.1 Data Acquisition

To capture the environmental and patient sleep-related data, we used a set of sensors, devices, and other data streams as described in Table 2.1:

Table 2.1: Data Sources and Details

| **Source** | **Sampling Frequency** | **Fields and Units** | **Error/Sensitivity** |
|---|---|---|---|
| Actiwatch Spectrum | Every 30 seconds | Interval Status (Active, REST, REST-S), Sleep/Wake (S/W), Activity Count (ACT) | Accelerometer: 0.025 G (a 2 count level), light sensor: 10% at 1500 Lux (typical) |
| Awair Omni | Every 10 Seconds | Light Intensity (lux), Temperature (°C), Humidity(%), $CO_2$(ppm), VOC(ppm), $PM_{2.5}$($\mu$g/m$^3$), Noise (dB) | Temperature (±0.2°C), Humidity(±2%), $CO_2$(±75ppm or 10% of reading), VOC(±10%), PM2.5(±15% or ±15$\mu$g/m$^3$), Noise (-26dBFS) |

Data Sources and Details

| **Source** | **Sampling Frequency** | **Fields and Units** | **Error/Sensitivity** |
|---|---|---|---|
| EnOcean ELLSU (OEM) Light Level Sensor | Samples every minute. Transmit if change > 50 lux versus last transmission. Heartbeat every 20 ... 30 minutes (affected at random) | Light Intensity (Lux) | ± 5 % at full scale / 68°F |
| Pressac Mini Temperature and Humidity Sensor | Transmit if temperature changes is greater than 0.6 degrees or humidity change of 2%, else Heartbeats every 15 minutes. | Temperature (°C), Humidity (%) | Humidity: ± 5 %, Temperature ± 1 °C |
| EPIC Systems Medical Data | Recorded manually by caregivers | The columns we look at specifically are chosen by the nursing team (see Figure A.2 for more details). | ± 15 minutes between care administered and care recorded |

The study was conducted in two phases; phase 1, accounting for 320 days worth of Energy Harvesting (EH) sensor data, and phase 2, accounting for an additional 93 days of Awair data.

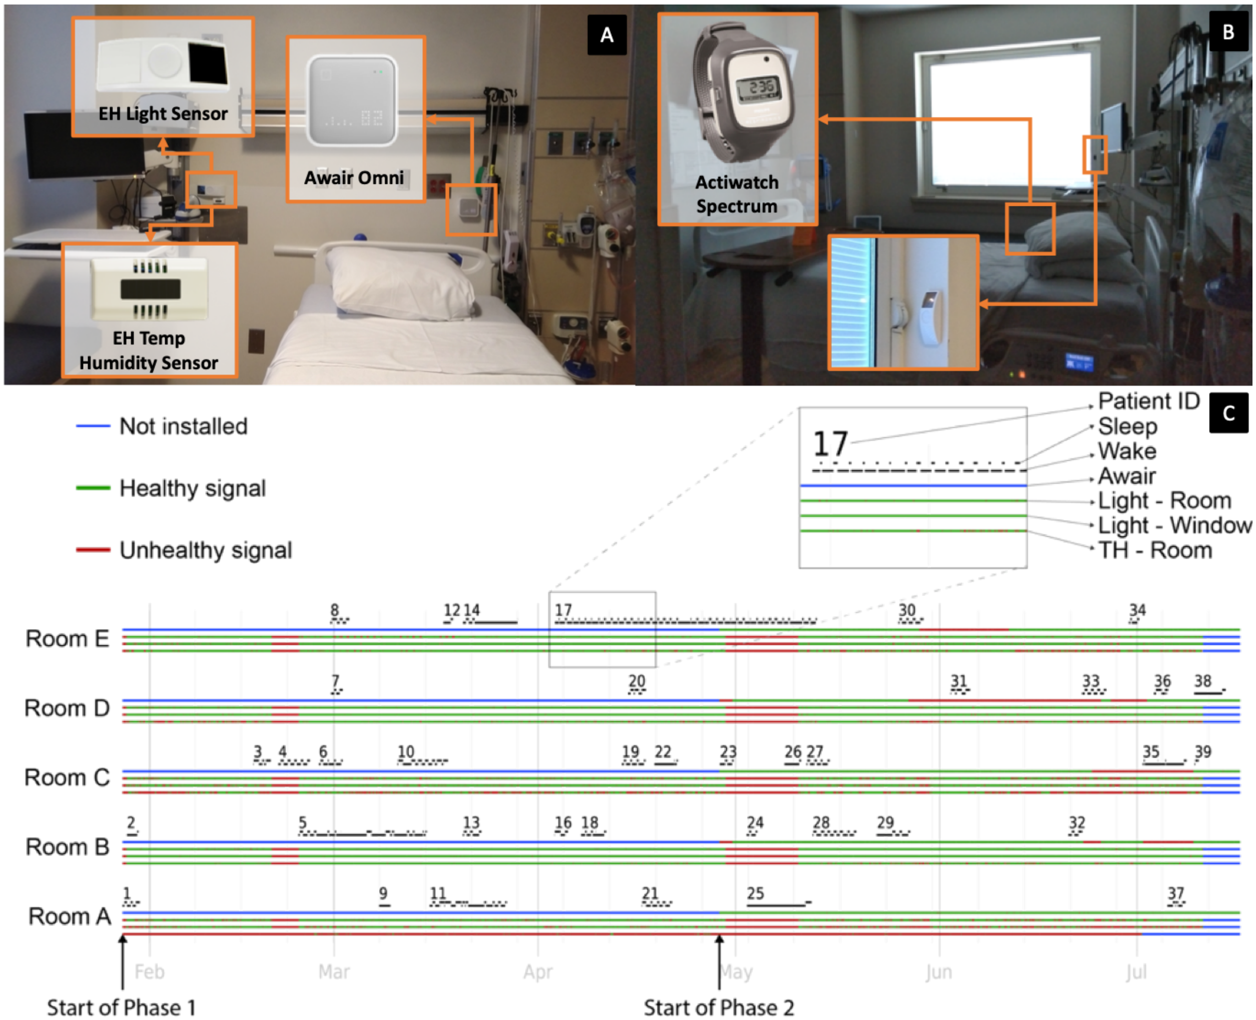


Figure 2.1: Example installation of EH and Awair Omni sensors (A), at the bed as well as the window (B), and Overall signal health of our hospital environmental sensors (C). A set of environmental sensors were installed at the end of August 2020, before the patient began to be admitted at the end of January of the following year. Healthy signals are determined using the signal health described in [1]. The unhealthy signals allow us to evaluate periods where data can be lower than expected or missing, which helps describe the representativeness and quality of our collection data. However, since the value collected during lower-than-expected periods are still valid, we included all data as part of our analysis. The sleep status of the patient is overlaid on the top of the health signals for reference.

**Data Inclusion and Data Exclusion** Figure 2.1 shows the variety of stay duration and sleep-wake patterns exhibited by the different patients throughout our study period. For example, Patient 9 is observed not to sleep at all, Patient 17 has the most extended stay duration, and Patient 38 is not detected by the Actiwatch to have slept until after a few days in the hospital.

### 2.3.2 Data Preparation

#### Data Processing for the mixed effect model

The treatment of the different data streams warranted additional processing before we combined them, and this section details that processing.

**Processing Actiwatch Data** In this work, sleep arousal is characterized by higher movement during a patient's sleep as tracked by wrist-worn actiwatch. The Actiwatch records data at 30 second intervals and characterizes patients' sleep during those intervals using multiple metrics. The activity count metric is essentially the activity level calculated based on accelerometer readings. Based on the activity count, a binary "sleep-wake status" variable is pre-calculated by the actiwatch software based on a threshold value of 40. If we detect more movement during a labelled sleep time, we consider it not conducive for restful sleep. The watch then uses a black-box algorithm to "sleep-wake status" as part of a state machine to finally identify episodes of sleep and differentiate a patient between resting and sleeping, or specifically "REST-S" (asleep), "REST", or "ACTIVE". In other words, a patient can be asleep (REST-S), but have a sleep-wake status of disrupted sleep (wake) in order for the Actiwatch to calculate sleep metrics such as the number of awakenings and sleep efficiency. We build upon the insight that Actiwatch provides where and utilize the combination of low activity epochs during REST-S as our normal sleep, and high activity epochs during sleep as our disrupted sleep.

**Adjusting for different time scales** One challenge with using multi-modal data was asynchronous and different time scales at which the data was recorded across different sources. To address this, we binned the data to the time scale with the least resolution and omitted from consideration bins that do not have at least one data point. For example, the EH sensors transmitted both on a periodic and event-driven schedule at around once every 30 minutes, but the Awair transmits routinely once every 10 seconds. Thus, we binned the data at a longer time scale of one hour to allow for comparison of our plots across different streams.

**Monotonic environmental disruptions based on deviation from known ranges** Unlike lighting, which disrupts sleep only above and not below a threshold, environmental factors such as temperature and humidity can be uncomfortable for sleep by being too high or too low. We, therefore, transformed the environmental values by calculating their distance from known optimal ranges based on literature before utilizing them as input to our models. Specifically, to get the adjusted environmental deviation value $e'$, we use Equation 2.1:

$$e' = \begin{cases} e - \theta_{\text{upper}} & e > \theta_{\text{upper}} \\ 0 & \theta_{\text{lower}} \leq e \leq \theta_{\text{upper}} \\ \theta_{\text{lower}} - e & e < \theta_{\text{lower}} \end{cases} \tag{2.1}$$

where $\theta_{\text{upper}} > \theta_{\text{lower}}$, and represent the selected tolerable upper and lower environmental set points respectively optimal for sleep, and $e$ represents the observed environmental readings. The optimal values and actual observed values are summarized in Table 2.2 along with a description of what those values actually represent. Finally, note that we combined temperature and humidity into a single signal called the heat index of the apparent temperature using MetPy for our models. [99].

Table 2.2: Selected non-disruptive environmental ranges for sleep and observed values by the Awair sensors. Observed values outside of optimal ranges are marked in bold.

| **IEQ Factor** | **Selected Ranges** | **Notes** | **Units** | **Sensed Ranges** | **Notes** | **Ref.** |
|---|---|---|---|---|---|---|
| Light | $\theta_{\text{lower}} = 0$ | No light | Lux | $e_{\text{min}} = 0$ | No light | [85, 75, 76] |

| | | | | | | |
|---|---|---|---|---|---|---|
| | $\theta_{\text{upper}} = 10$ | "Horizon, clear sky after sunset" | | $e_{\text{max}} = 7770.9$ | Indoors between scattered clouds and complete overcast daytime sky | |
| Temperature | $\theta_{\text{lower}} = 18$ | Cold | °C | $e_{\text{min}} = 16.59$ | Cold | [84, 85] |
| | $\theta_{\text{upper}} = 28$ | Warm | | $e_{\text{max}} = \mathbf{30.75}$ | Warm | |
| Relative humidity | $\theta_{\text{lower}} = 40$ | Dry | % | $e_{\text{min}} = 20.3$ | Dry | [85] |
| | $\theta_{\text{upper}} = 60$ | Humid | | $e_{\text{max}} = \mathbf{67.8}$ | Humid | |
| Noise | $\theta_{\text{lower}} = 0$ | "Threshold for normal human hearing" | dBA | $e_{\text{min}} = 44.1$ | "Normal living, talking, or radio in the background" | [85, 77] |

| | | | | |
|---|---|---|---|---|
| | $\theta_{\text{upper}} = 60$ "Noisy lawn mower at 10 meters" | | $e_{\max} =$ **84.3** "Heavy traffic at 10 meters, door closure" | |
| $PM_{2.5}$ | $\theta_{\text{lower}} = 0$ $\theta_{\text{upper}} = 12$ | $\mu g/m^3$ | $e_{\min} = 0$ $e_{\max} =$ **1,000** | [96] |
| $CO_2$ | $\theta_{\text{lower}} = 400$ $\theta_{\text{upper}} = 500$ | ppm | $e_{\min} = 400$ $e_{\max} =$ **1,008.0** | [88] |
| TVOC | $\theta_{\text{lower}} = 0$ $\theta_{\text{upper}} = 500$ | ppb | $e_{\min} = 20$ $e_{\max} =$ **38,769** | [93] |

**Change point detection** We applied change point detection to investigate momentary disruptions in environmental conditions and their association with sleep disruption during phase 2 (after the heat index was calculated for temperature and humidity). Change-point detectors have been used in the literature in different fields such as health [100], transportation engineering [101], and behavioral science [102] for a variety of applications. Note that we are not interested in random peaks in the environmental changes, which might be due to sensor reading error; rather, we are looking for substantial changes in the distribution of the environmental attributes data. We specifically use the Gaussian Kernel change point detector as implemented in the *ruptures* package [103], assuming an unknown number of changes ahead of time. The Gaussian Kernel solves a penalized optimization problem for the change points, which takes the form:

$$k_{\text{Gaussian}}(u, v) = \exp(-\gamma \|u - v\|^2) \tag{2.2}$$

Where $u$ and $v$ are two-dimensional vectors and $||\cdot||$ the Euclidean norm, and $\gamma > 0$ is a user-defined parameter we chose manually. The literature shows that human discomfort during sleep arises from environmental conditions falling out of the optimal ranges and the rates of change. For instance, regular hospital activities from nursing staff, medication carts, and roommates can cause sudden changes in the environment around the patient [104]. We run the change point analysis on contiguous portions of environmental data from the beginning till the end of each patient sleep episode to identify events/change points in the signals. Figure 2.2 shows an example of patient waveform and found change points. As we can see in Figure 2.2, several change points in environmental attributes overlap with disruptions in sleep during both sleep episodes. Specifically, change points in temperature, humidity, $CO_2$, and VOC overlap with the moments of sleep disruption during both sleep episodes. Noise and lighting change points precede the awakenings during the first sleep episode and overlap during the second. Similarly, *Other Night time medical events* overlap with disrupted sleep during the second sleep episode. Considering the medical events and environmental disruptions together, we witness some co-occurrences implying that any nursing or other activities in the room result in a change in the room environment happening. For instance, *Other Night time medical events* overlap with noise disruptions.

**Processing Medical Events Data for the mixed-effect model** The hospital uses an EPIC electronic health record (EHR) system [1], a combination of automatic and manual data entries. The EHR system assists with the care of the patients and is maintained by the nursing staff. However, based on the expertise of nursing staff, there can be a margin of error of around 15 minutes between the actual occurrence of an event versus its recording in the system. Furthermore, the information tracked on the system to provide better care might differ across hospitals, even with the same system. As a result, we relied on the nursing staff to provide a set of potential variables most relevant to patients' sleep. Due to the large size of this list, we document these columns in Figure A.2. Finally, we converted all medical events to a binary format based on time for any epoch within 15 minutes of a recorded time event. Thus, if an event happened at time $t$, the window of an event is considered as $t \pm 15$. For any epoch that falls within the range of a recorded event, we label our medical event row as true. Otherwise, we label the epoch as false.

[1] https://www.epic.com/

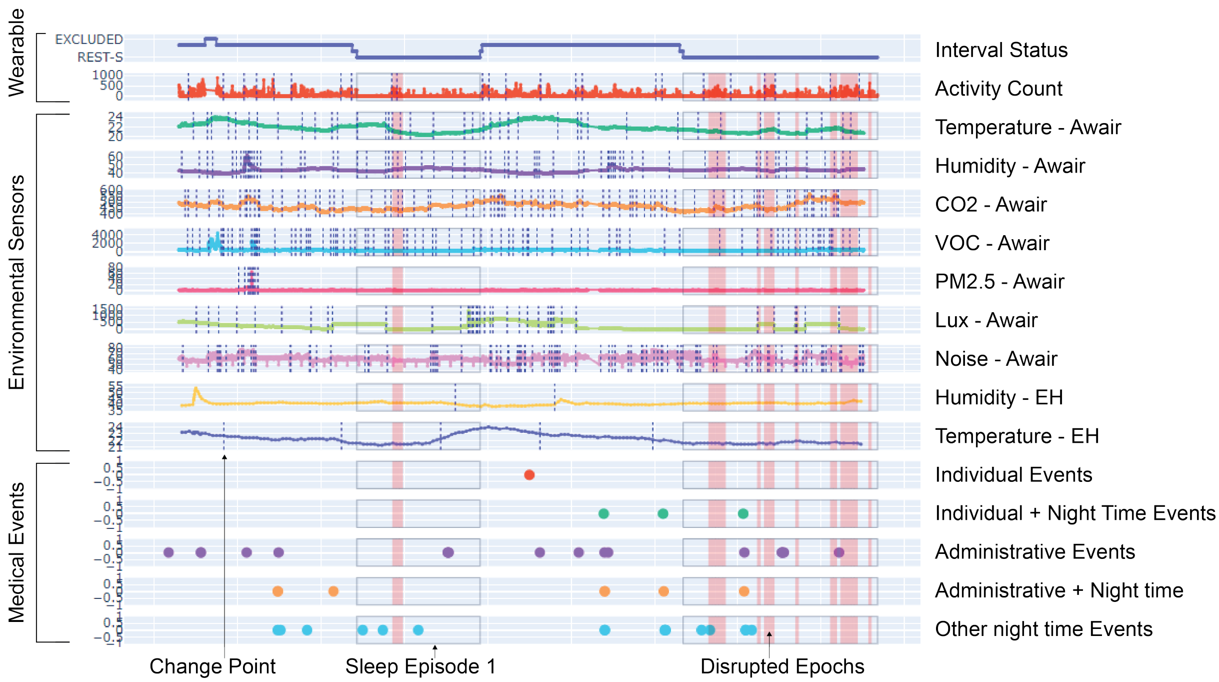


Figure 2.2: Example patient waveform including two sleep episodes (boxed in gray). The interval status refers to the Actiwatch built-in state machine, and "REST-S" refers to the state where the patient is considered sleeping. The vertical dash lines refer to the calculated change points described in Section 2.3.2, the grey surrounding boxes refer to each sleep episode (REST-S intervals), and the highlighted red areas correspond to disrupted epochs.

**Processing for multi-class prediction of medical events using Environmental Signals**

**Feature Preparation** The processing step for multi-class prediction uses the same data but was prepared differently. To account for missing data, we first standardized data from the different Awair environmental sensors to one-minute aggregates. Then, we varied the size of the time window per classification, up to the nurse-specified 15-minute error. This means, for example, that the time window for classifying a medical event can go fifteen minutes before and after the time recorded. When multiple medical events are observed, the closest event to the center is chosen as the event class. After splitting the testing and training set 20-80, we convert the values to standard scores before feeding them into a classifier. We converted data within training and testing sets separately into standard scores for each environmental category. The following equation calculates the standard score [105]:

$$z = \frac{x - \mu}{\rho} \tag{2.3}$$

where: $z$ is the $z$ standard score for the population, $x$ the raw value, $\mu$ the mean of the population, and $\rho$ the standard deviation of the population. We calculated the $z$ score after splitting data into training and testing sets to avoid information leakage. Figure 2.3 showcases an example heat map of the final processed feature set. We further describe the event class label below.

**Event class and data augmentation** We are interested in predicting five classes using the environmental features. Specifically, we predict: No events, Individual Events, Other Night Time Events, Medication administration Events, and Combination of Events. *No events* represent time windows where no medical events are recorded. Individual events are recorded for Falls, MET calls (emergency response call to other nurses when a patient's vitals fall outside healthy ranges, for example), and $O_2$ Delivery Devices. *Other nighttime events* group together events such as measuring blood glucose levels, EKGs (measures a Heart's electrical signal), and vital signs. *Medical Administration Events* represent instances where non-narcotic pain medications, opioids, and other sedative medications are distributed. *Combiantion of Events* represents instances where combinations of previously described classes co-occur (as part of the manual input, not as a function of the time-window). We include in the Appendix Figures A.3 to A.5 a detailed mapping of classes to columns for reference. A significant class imbalance was observed for classifying medical events data, as shown in Figure A.6. To address this imbalance, a Synthetic Minority Over-sampling Technique (SMOTE) was used [106, 107]. Specifically, we utilized the implementation in *imblearn* [108]. Compared to regular over-sampling, where samples are drawn with replacement, SMOTE prevents over-fitting by generating new minority samples by creating synthetic samples close to the feature space. Using SMOTE, we generated synthetic samples for all of the minority classes after we converted values in the training set into z-scores.

### 2.3.3 Evaluation Methods

#### Visual Analysis

We approach the exploration of environmental factors by plotting our observations in different time scales, from seasonal (i.e., spring, summer) to hourly. We first approach the plotting of box plots to understand the different temporal scales and when certain environmental factors fall outside of known sleep optimal ranges. We plot the entirety of our observations, including times when no patient is present, to observe the room's potential to serve a quality sleep environment. Secondly, we plot split violin charts of 1) environmental values per hour with patient sleep disruption and 2) environmental change points with patient sleep disruption. The split violin plots will give us insight into how the different categories (i.e., normal sleep versus disrupted sleep) can be identified via individual environmental variables. However, multiple relationships can be overlooked by plotting the environmental factors separately. Therefore, we describe the use of a statistical model in the next subsection to see if we can gather further insights.

#### Statistical Modeling

To evaluate our observations, we utilized the GLMM described above with binary outcome variable as sleep disruption/awakening every 30 seconds and predictor variables as light, apparent temperature, noise, $CO_2$, VOC, $PM_{2.5}$, their change points, and medical events. The environmental variables from the Awair sensor record data every 10 seconds but are aggregated every 30 seconds. Regarding the medical events data, if the sleep epoch falls within ± 15 minutes of any medical event occurrence, it is considered True; otherwise, it is False. GLMMs enable us to make predictions with continuous and count variables and allow us to account for inter-patient variability [109]. GLMM follows the general form of:

$$y = X\beta + Zu + \epsilon \tag{2.4}$$

Where $y$ is a $N \times 1$ outcome column vector; $X$ is a $N \times p$ matrix of the $p$ predictor variables; $\beta$ is a $p \times 1$ column vector of the fixed-effects regression coefficient; $Z$ is a $N \times q$ design matrix for the $q$ random effects; $u$ is a $q \times 1$ vector of the random effects; and $\epsilon$ is a $N \times 1$ column vector of the residuals. GLMM was chosen to account for repeated observations from the same individuals over time and to evaluate the many predictor variables together. Specifically, because of the wide variety of samples in our observed patient population in terms of length of stay, we decided to model the patients as random effects and the environmental variables and perturbations as fixed effects. We removed highly correlated variables to reduce their overall influence over the model, such as using apparent temperature instead of humidity (see Appendix A.7).

### Decision Trees and Random Forest Models

The use of decision trees and random forest models to explore multi-class classification has been successfully utilized in many in-hospital applications using environmental data [70, 110]. Decision Trees (DTs) are a non-parametric supervised learning method used for classification and regression that predicts the value of a target variable by learning simple decision rules inferred from the data features [111]. A random forest (RFs) model is an ensemble method that merges the results from multiple decision trees [112]. While sometimes more accurate, a Random Forest model in practice can take more time to run because of the number of trees the model needs to generate and merge. In our case, we accounted for the additional time commitment by first running Decision across all possible environmental combinations, then running Random Forest Classifiers on a subset of hyperparameters used for the top Decision Tree Classifiers. We use the Python package *sklearn* [111] module's implementation of DTs and RFs to conduct our experiments. We display the $F_1$-score, calculated as:

$$F_1 = \frac{2}{\text{recall}^{-1} + \text{precision}^{-1}} \tag{2.5}$$

To balance between computation time and avoid over-fitting for our large scope of hyperparam-

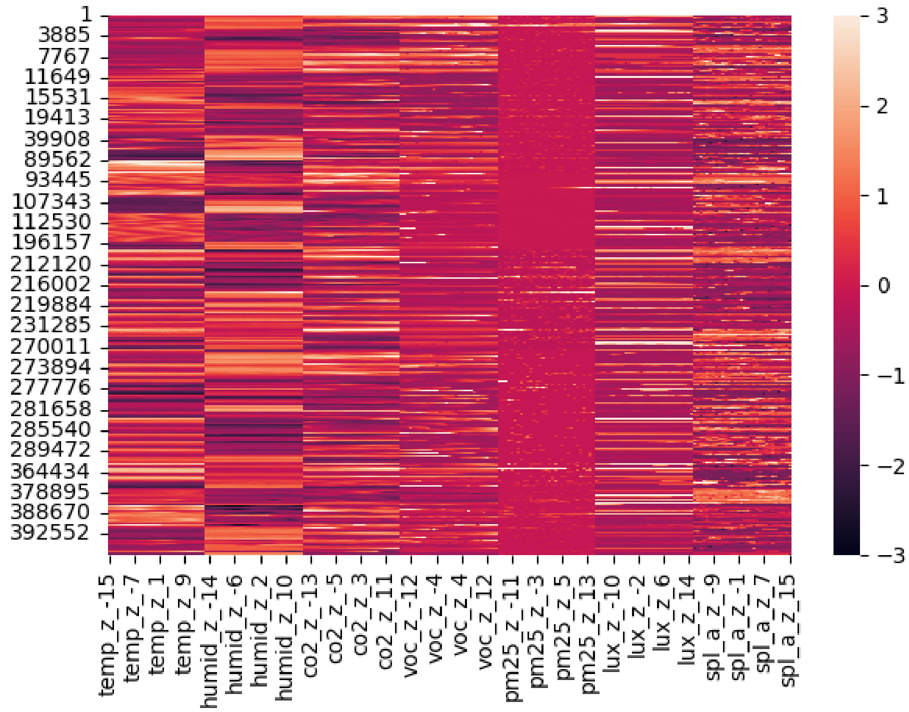


Figure 2.3: Example feature set as a heat map. The standard $z$ score on the right is capped at -3 to 3 for visibility. The y-axis shows the index of the signal window we process, and the x-axis shows the environmental variables, with -15 representing 15 minutes before the recording of the event and two being two minutes after the event.

eters, we cross validate our models by randomly splitting the data between training and test set (20% to 80%) ten times and return the aggregated $F_1$-score. Specifically, we use the macro-weighted $F_1$-score, giving equal weight to the importance of predicting all classes.

### 2.3.4 Assumptions

In our statistical modeling task, we assume that all patients are perturbed to the same degree with regard to environmental deviations. This generalized patient model is likely not representative of individual patient biases, but we chose this model to help account for the large variety of different sleep times across patients. For our machine learning task, we assume that the manual inputs by the nursing team into the EPIC system are accurate. Because we did not install cameras in the room, we rely on their input as ground truth. Our study mostly focuses on the ability to predict these labels, with the assumption that utility can be further provided in reduced cognitive load for computer-assisted inputs.

## 2.4 Results

### Environmental Trends in the hospital

**Seasonal and hourly environmental factors in the hospital** In Figure 2.4, we show lighting, temperature, and humidity values change by season. Specifically, we see that the temperature and humidity fluctuations do not corroborate with higher temperatures in the region during the summer. This indicates that HVAC regulates indoor temperatures and that the range of allowable deviations is smaller during the summer than in the other seasons. The smaller variance in observed summer temperature also indicates potential benefits in altering. On the other hand, the humidity trends fall in line with known historical values, being more humid closer to the summer and less humid around the winter. This showcases the potential value of utilizing humidity control to improve the overall comfort of patients in the hospital, especially since ranges routinely fall out of known comfortable zones.

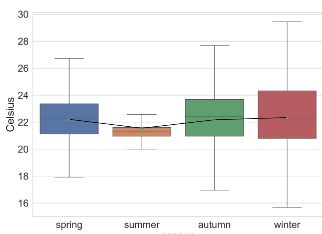


(a) Seasonal temperature trends

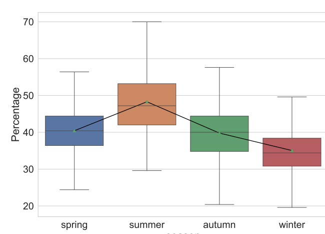


(b) Seasonal humidity trends

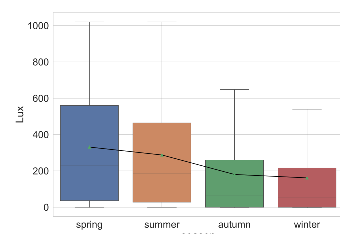


(c) Seasonal lighting trends

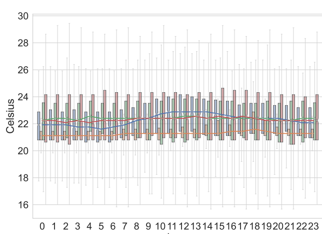


(d) Seasonal, hourly temperature trends

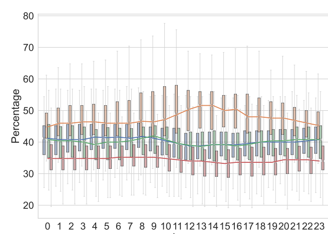


(e) Seasonal, hourly humidity trends

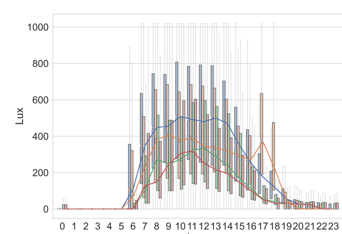


(f) Seasonal, hourly lighting trends

Figure 2.4: We notice annual seasonal fluctuations in temperature, humidity, and lighting from the EH sensors. Surprisingly, the temperature variability was smallest during the summer and largest during the winter. Correspondingly, we saw, on average, higher humidity values during summer and lower humidity levels during.

Figure 2.5 demonstrates a "washing out" effect, where the sensor can observe a higher range of lighting values because it detects more lighting contribution from sunlight outdoors than artificial

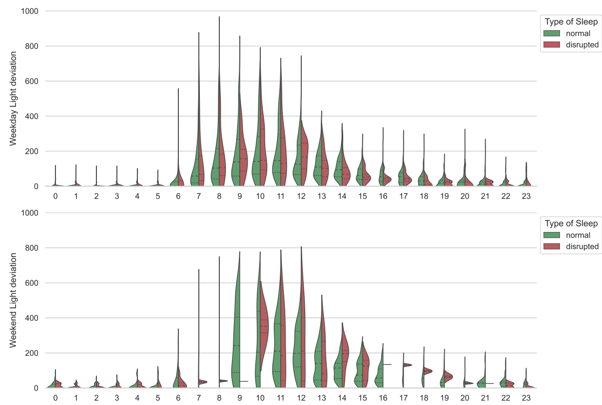

(a) Weekend Versus Weekday Room Level Lighting

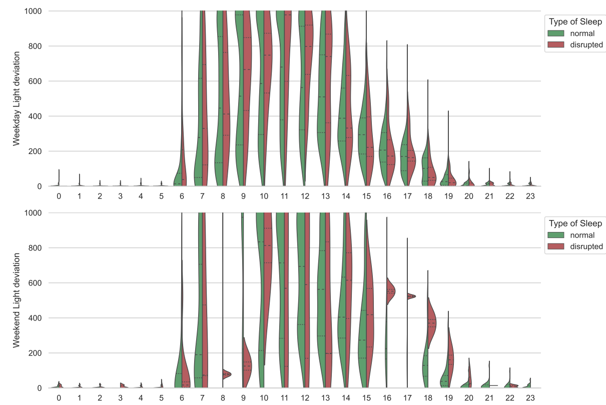

(b) Weekend Versus Weekday Window Level Lighting

Figure 2.5: Hourly split violin plots of room-area and window-area EH sensors. Sensors placed at the window area show higher variance during daylight hours. The left side of a violin shows distributions during normal sleep epochs, while the right side of a violin shows distributions during disrupted sleep epochs.

light indoors. We generally observe higher lighting values during late afternoon hours (4 pm to 12 am) having a noticeable relationship with sleep disruption. Further, we observe anticipated daylight fluctuations, being lower at night and higher during the day. This indicates the necessity to either incorporate black-out blinds to further protect a patient's luminous environment or the importance of utilizing sunlight or other methodologies [113] to help entrain a patient's circadian rhythm towards a regularized cycle that rises and falls with the sun. Moving on to Phase 2 Awair sensors, in Figure 2.6, we show CO2, VOC, and Noise values for Spring and summer. Specifically, we observe a peak in environmental perturbation across VOC and CO2 around noon hours, indicating perturbations based on dining activities. Further, we observe summer-time VOC median values to be higher than spring-time VOC values and spring-time CO2 values to be higher than summer-time CO2 values. The similar signals between VOC and $CO_2$ indicates that VOC values can also correspond with higher occupancy and movement. However, VOC being higher than $CO_2$ during the summer indicates that additional mechanisms can result in more elevated VOC values, not just elevated occupancy, such as higher decomposition rates [114].

**Pattern of higher environmental variability during the weekdays compared to the weekends, especially during the day shift** Split violin charts, which show the distribution differences between two categories of data to help demonstrate the relationship between different

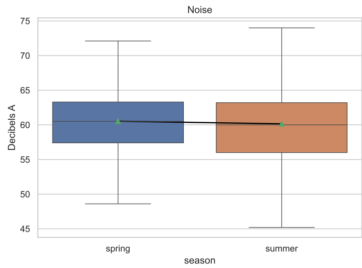


(a) Seasonal noise trends

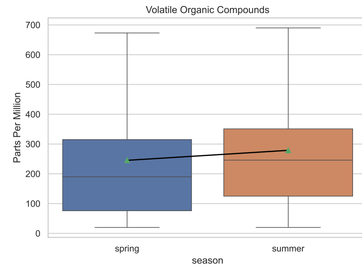


(b) Seasonal VOC trends

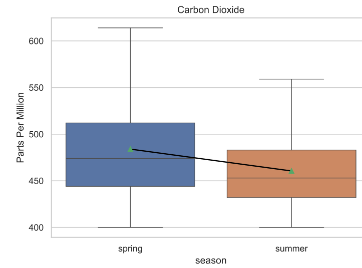


(c) Seasonal CO2 trends

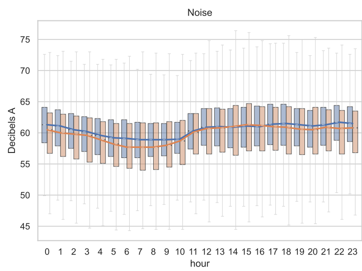


(d) Seasonal, hourly noise trends

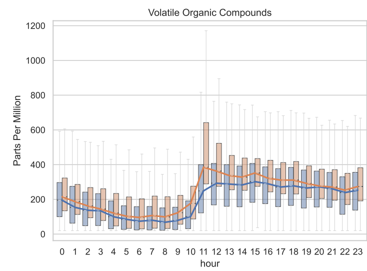


(e) Seasonal, hourly VOC trends

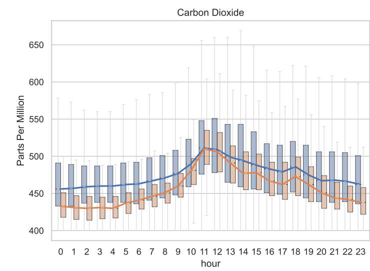


(f) Seasonal, hourly CO2 trends

Figure 2.6: During Phase 2 of our study, including the Awair sensors, we observed a noticeable difference in overall VOC across all hours during the summer and increased $CO_2$ during the spring season. VOC fluctuations are noticeable around noon. We moved plots of $PM_{2.5}$ to the appendix because we found minimal perturbations.

distributions of environmental factors and the categories in question. For example, by drawing a split violin plot per hour, we can see not only the min, max, median, 25th, and 75th quartile but also see how the same environmental factor distribution measure compared to the category in question (in our case, normal versus disrupted sleep). Looking at Figure 2.7c, we observe the anticipated rise in lux during the daytime from around 7 am to 7 pm and lighting distributions from midnight to about 6 am. This demonstrates the value of the environmental sensors in providing additional information about the hospital's luminous environment, specifically undesirable nighttime lighting. Furthermore, the fluctuations in lighting have much lower ranges over the weekends compared to weekdays. Similar undesirable can be observed in temperature (Figure 2.7a), and VOC (Figure 2.7e), where values can be lower or higher than known optimal ranges. For space considerations, we moved the $PM_{2.5}$ plot into the appendix Figure A.8, but here we observe low amplitude spikes during lunchtime and dinner time, which are also smaller during the weekends.

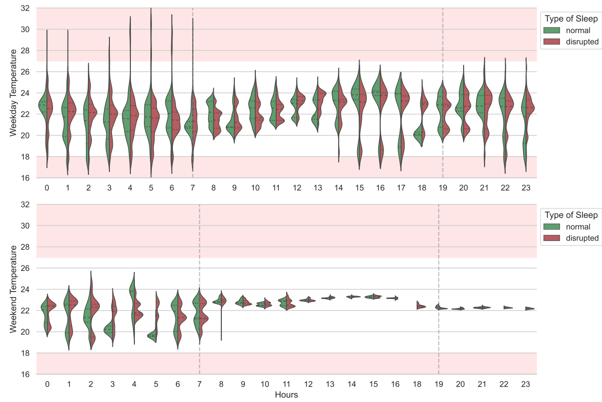


(a) Weekend Versus Weekday Temperature and Sleep Disruption

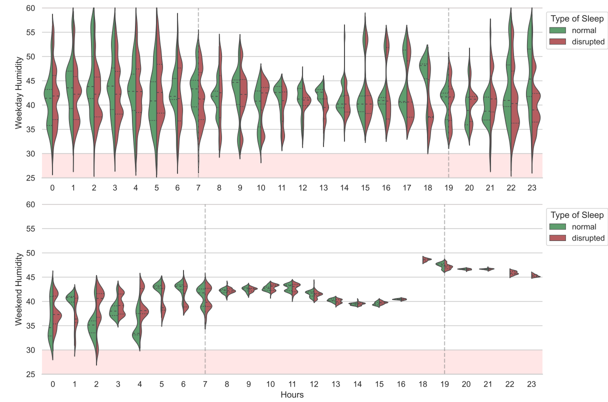


(b) Weekend Versus Weekday Humidity and Sleep Disruption

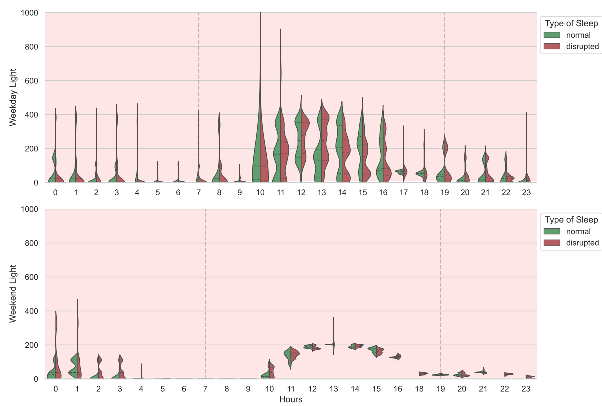


(c) Weekend Versus Weekday Lighting and Sleep Disruption

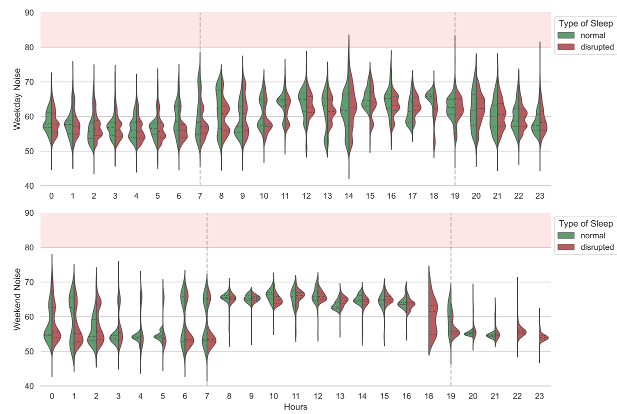


(d) Weekend Versus Weekday Noise and Sleep Disruption

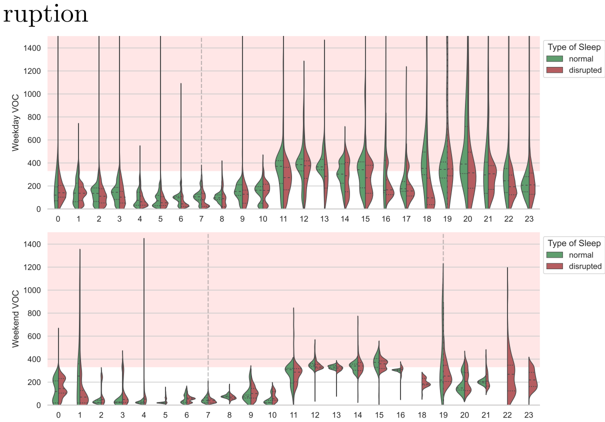


(e) Weekend Versus Weekday VOC and Sleep Disruption

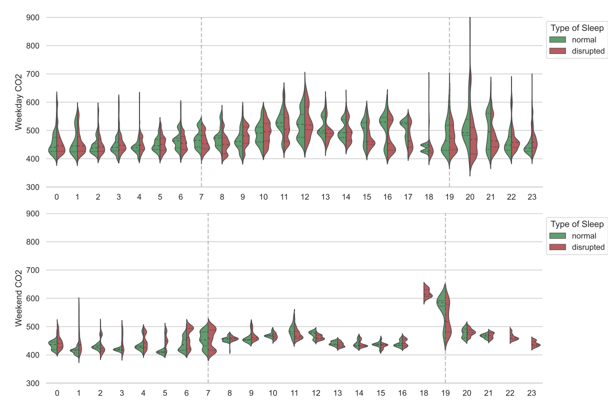


(f) Weekend Versus Weekday CO2 and Sleep Disruption

Figure 2.7: Split violin plots of weekday (top) versus weekend (bottom) environmental factors and their relationship with patient sleep disruption. The dotted vertical lines around 7 am and 7 pm indicate the start of the day and night shifts, respectively. We observe a variety of different types of disruption, from lighting (Figure 2.7c), where for the majority of times, the hospital is above the recommended threshold of light (10 lux), to humidity (Figure 2.7b), where patients can be either too hot or too cold; to $CO_2$ (Figure 2.7f), where the ranges are entirely within known thresholds that disrupt sleep.

### 2.4.1 Generalized Mixed-Effect Modeling of relationships between environmental, medical, and sleep disruption

Using a GLMM described in the methodology section, we found noise deviations, $CO_2$ deviation, $CO_2$ change points, VOC deviation, Lux deviation, and medical events to be significant predictors for sleep disruption as shown in Table 2.3 and Figure 2.8. We also found VOC and $CO_2$ change points to be significant. In this table, the intercept describes the residual unaccounted-for variability as a value. The estimate describes the positive or negative correlation the predictor variables have with the outcome variable. For example, in Table 2.3, we see a positive relationship between lux deviations and the likelihood of a disrupted sleep epoch. The $z$ value is the estimate divided by the standard error and describes the deviation from the mean. The $p$ value ($Pr(>|z|)$) is calculated based on the $z$ value and represents the likelihood that the predictor variable is an outcome of chance. The lower the value, the more likely the predictor variable is significant. These results suggest that the tracking of lighting, noise, medical events, and change points for $CO_2$ to be relevant in analyzing patient sleep characteristics, even for a diverse set of hospital patients. In our discussion section, we elaborate on our speculation as to the mechanism of disruption and the relationship between the predictor variables and outcome variables.

Table 2.3: Fixed effects results from the binomial family generalized linear mixed model fit by maximum likelihood, with patient id as random effects.

| | **Estimate** | **Standard Error** | **$z$ value** | **Pr($>|z|$)** |
|---|---|---|---|---|
| (Intercept) | -2.437e+00 | 1.834e-01 | -13.288 | <2e-16 *** |
| Noise deviation | 1.984e-02 | 3.944e-03 | 5.029 | 4.94e-07 *** |
| Noise Change Points | 5.339e-02 | 1.665e-01 | 0.321 | 0.74852 |
| Heat Deviations | 7.465e-02 | 8.065e-02 | 0.926 | 0.35469 |
| Heat Change Points | 2.389e-01 | 1.516e-01 | 1.575 | 0.11514 |
| $CO_2$ Deviation | 2.808e-03 | 6.139e-04 | 4.574 | 4.79e-06 *** |
| $CO_2$ Change Points | 3.410e-01 | 1.368e-01 | 2.492 | 0.01269 * |
| Lux Deviation | 3.233e-04 | 1.037e-04 | 3.117 | 0.00182 ** |
| Lux Change Points | -5.523e-02 | 2.199e-01 | -0.251 | 0.80169 |
| PM 2.5 Deviation | 5.747e-03 | 6.697e-03 | 0.858 | 0.39077 |

| | | | | |
|---|---|---|---|---|
| PM 2.5 Change Points | 2.018e-01 | 4.579e-01 | 0.441 | 0.65949 |
| VOC Deviation | 1.641e-04 | 7.364e-05 | 2.229 | 0.02582 * |
| VOC Change Points | -3.921e-02 | 1.651e-01 | -0.237 | 0.81230 |
| Medical Events | 7.600e-01 | 2.751e-02 | 27.628 | <2e-16 *** |

Signif. codes: 0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

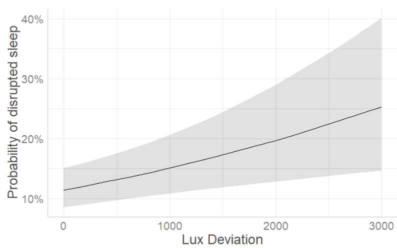


(a) Probability of disrupted sleep and lux deviation.

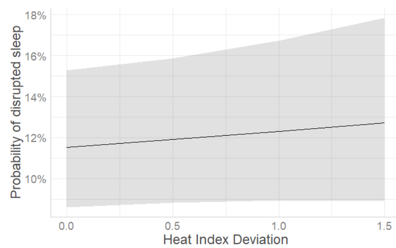


(b) Probability of disrupted sleep and heat index deviation

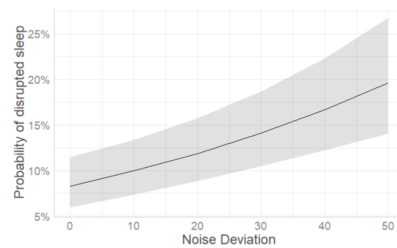


(c) Probability of disrupted sleep and noise deviation

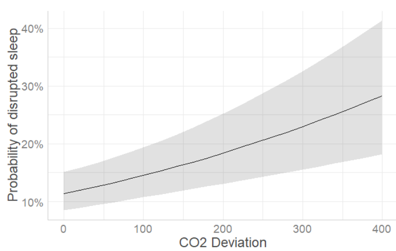


(d) Probability of disrupted sleep and $CO_2$ deviation

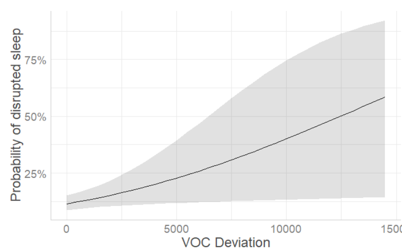


(e) Probability of disrupted sleep and VOC deviation

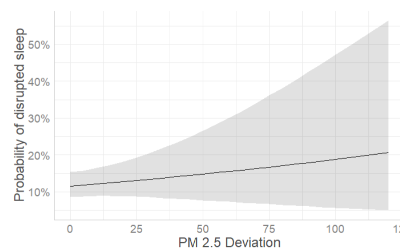


(f) Probability of disrupted sleep and $PM_{2.5}$ deviation

Figure 2.8: Predicted Probability of Sleep Disruption Per Environmental Factor, showing a positive relationship between lux, $CO_2$, noise, and VOC deviations and patient sleep disruption.

### 2.4.2 Predicting medical events using environmental data

To explore the combination of values without running all the computations, we first ran DT models across all possible combinations of environmental variables across all possible window lengths. We then selected a set of the highest-performing environment and window size combinations to run the random forest model to generate Table 2.4, comparing two of the highest-performing models. We omitted other classifiers such as SVM and Naive Bayes because of how poorly they performed during our testing (both achieved roughly 30% accuracy). With the same environmental combination and window size, the Random Forest Classifier achieved generally higher scores than the Decision Tree classifier. When comparing the highest-scoring models, the Random Forest model was able to achieve a 93.2% Macro-weighted $F_1$-score, with a 31-minute window combining temperature, humidity, $CO_2$, VOC, light, and noise, while the Decision Tree classifier was able to achieve an 85.1% $F_1$-score with a 1-minute window combining temperature, humidity, light, and noise.

| **Model** | **Combination** | **Iterations** | **Window** | **$F_1$-score** |
|---|---|---|---|---|
| Decision Tree | temp,humid,light,noise | 10 | 1 minute | 85.1% |
| Random Forest | temp,humid,co2,voc,light,noise | 10 | 31 minutes | 93.2% |

Table 2.4: Comparison of highest scoring (macro-weighted) models between Decision Tree and Random Forest Classifiers

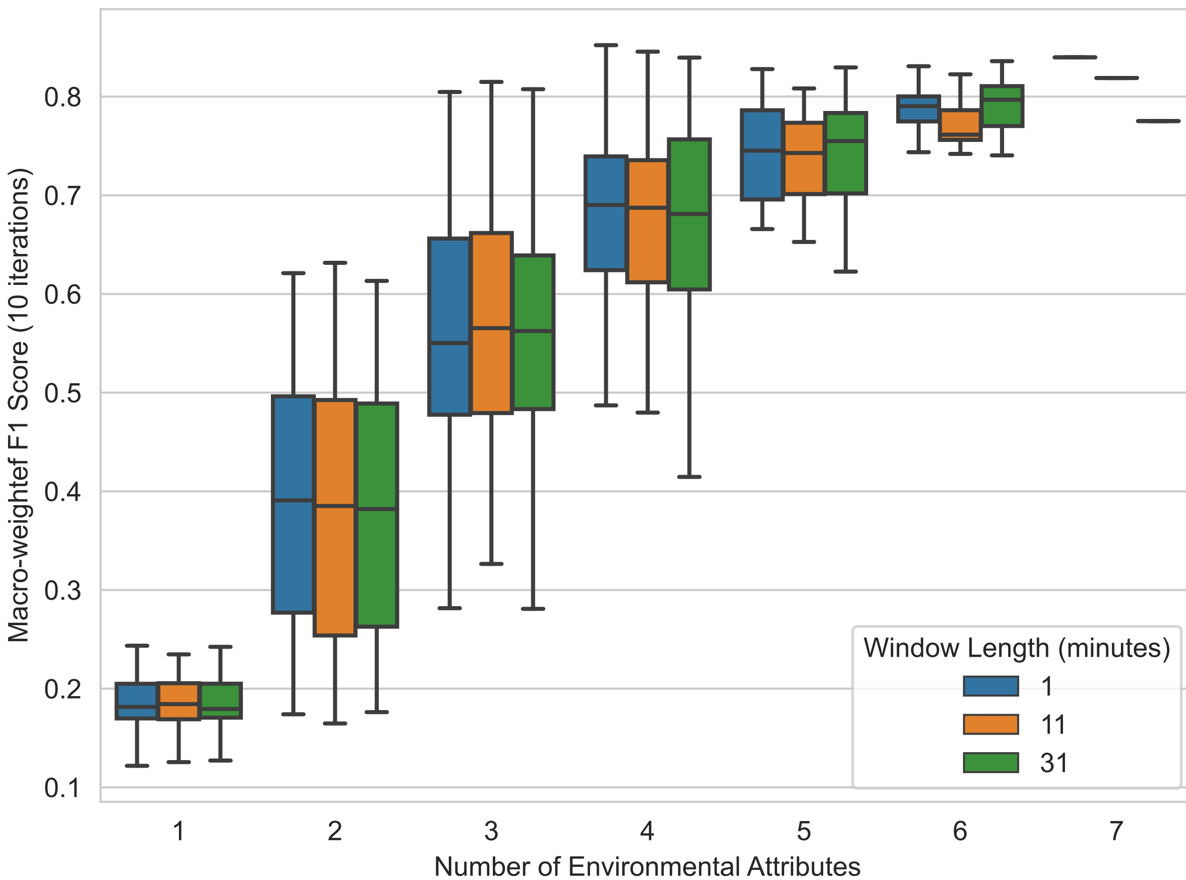


Figure 2.9: Window size, environmental combinations, and $F_1$-score using a Decision Tree Classifier. The figure demonstrates how additional accuracy can be gained by combining environmental measures and how the window length in our study was inconsequential in the large picture to achieve accurate classifications. Important to note is that while more environmental attributes *generally* lead to higher median $F_1$-scores, the highest-scored model was achieved only using four environmental streams (see Table 2.4).

## 2.5 Discussions

In this section, we elaborate on the findings for each results section and then dive into broader impacts, lessons learned, and limitations.

### 2.5.1 Visualization Results

To answer **RQ1**, while we did observe a trend of more minor variances (via interquartile range) for each environmental factor across the weekend, we also observed more significant variances during weekend night shifts than weekend day shifts for temperature and noise. Furthermore, we observe higher median values across the different environmental variables during work hours. Together, this indicates that the nature of environmental disruptions might not be able to be fundamentally separated from medical routines and human traffic. We observe that the hourly fluctuation trends of VOC visually match the hourly CO2 fluctuations and can have the potential to be used as a proxy for occupancy measures as well. However, compared to $CO_2$ trends, VOC is higher during the summer than during the spring, the opposite of $CO_2$. The elevation VOC observed during summer can be due to thermal plumes due to elevated temperatures [90]. This, together with elevated VOC values around noon, suggests that using VOC as an occupancy proxy would benefit from dining and seasonal adjustments. Interestingly, while humidity becomes uncomfortable below and above thresholds, only lower humidity values were observed during a patient's sleep time. For instance, we observe seasonal humidity in all seasons except the summer humidity is a problem on the lower end instead of the upper end. This suggests a humidifier might be of more use in larger percentages of the year. Tracking just patient sleep timing in conjunction with uncomfortable environments trends can mislead environmental deviations and the potential for the room to be a conducive environment with high sleep quality because a patient's room may be either unoccupied or not sleeping during an otherwise uncomfortable environment.

### 2.5.2 Generalized Linear Mixed Model Results

To answer **RQ2**, the shape of the models shown in Figures 2.8 suggests a significant positive relationship between environmental variables and patient sleep disruption, and also to what degree of increased probability of disruptions is related with each increase in environmental deviation. Lux, heat, noise, $CO_2$, VOC, and $PM_{2.5}$ correspond with patient sleep disruption. We also observed scarce $PM_{2.5}$ perturbations and found it potentially useful in tracking meals because the peaks coincided with lunch and dinner times (See Figure A.8). This suggests that more granular modeling of patient activity could be relevant for future studies, but since we did not install video cameras out of privacy concerns, we cannot verify with a ground-truth set. For change point signals found not to be significant, such as $PM_{2.5}$, lighting, or noise change points, one possibility can be due to the inherent nature of this particular hospital environment. For example, gradually changing environmental variables can be continuously disruptive without abrupt changes in their sampled distribution. In the case of $PM_{2.5}$, because the values we've detected are mostly small, with most distributions found around 5 $\mu g/m^3$ (see Figure A.8), even sudden changes may be imperceptible to the patients. However, we should note that it can not be verified without the introduction of video cameras or other human activity tracking mechanisms whether the detected environmental perturbations are proxies of human interactions or the ambient environment.

### 2.5.3 Decision Tree and Random Forest Results

The random forest and decision tree models demonstrate that environmental signals adequately tailored for medical events detection can accurately classify from among various classes and events. The results further demonstrate how environmental values, under the sensing by proxy paradigm, can provide utility as an objective observer. The inconsequential nature of the window range likely indicates that the frequency of data, even when aggregated at one-minute intervals, can contain information necessary to identify patterns of environmental perturbations that can be linked to a medical event. Figure 2.9 demonstrates that adding more environmental streams, in general, can improve the median score for the machine learning model but can lead to less efficient and less effective models than the optimal model. Additionally, while adding environmental measures is helpful as a trend, increasing the window size for consideration did not improve our model. In fact, at the maximum number of combinations, the most significant observation window of 31 minutes performed the worst compared to the other time windows at 11 minutes and 1 minute. This trend suggests that it is sometimes not the quantity of data collected that matters after a particular frequency but rather the diversity of data that improves our model. To answer **RQ3**, surprisingly, in both of our highest-performing models, temperature and humidity were crucial features to retain beyond the well-known environmental proxies of light and noise, as shown in Table 2.4. Lastly, adding more information does not seem to, by default, improve inference accuracy. Additional information can inadvertently lower the score depending on the model and hyperparameters.

### 2.5.4 Broader Impacts

Improving the environment and sleep of people in a critical shared space such as the hospital carries with it immense societal if it can lead to better care and health outcomes. By integrating environmental, medical, and physiological data streams and demonstrating the key predictors of sleep disruption, researchers and caregivers can identify factors to which to dedicate more resources in education and awareness. In this work, we demonstrate that all proposed streams of data can have significance with regards to the patient's sleep when considered together (i.e., noise deviation $p < 0.05$), $CO_2$ deviation $p < 0.05$), medical events ($p < 0.05$)). Furthermore, even the ones shown not to be directly tied to patient sleep disruption could be routinely monitored to improve the understanding of the hospital environment and routines. For example, heat deviations can help inform thermal comfort in the hospital, and $PM_{2.5}$ deviations might help inform dining activities. Our work demonstrates methodologies to combine and assess different facets of the environment towards a quantifiable health metric, in addition to enabling the identification of abrupt environmental changes ($CO_2$ change points ($p < 0.05$). The results indicate the potential benefits that broader adoption of environmental tracking in hospitals can provide. Environmental sensors have a real potential to provide substantial utility beyond environmental commissioning; they can be considered an invaluable addition to support privacy-friendly hospital data products, automation, and computer-assisted decision-making.

### 2.5.5 Limitations

Sleep is a complicated phenomenon of human existence, with which exercise, age, gender, sleep regularity, and medication have all been shown to have significant relationships [65]. Additionally, past literature has shown that actigraphy's sleep-wake measure decreases whenever sleep is disturbed or distorted [115, 116]. These demonstrate potential confounding variables for the complexity of sleep and technology limitations. While we did not find deviations about apparent temperature to be significant in our observation, that does not mean that temperature was not important to track, nor temperature not affecting the patient. One possible disconnect can be between the sensor's observed temperature and the actual perceived temperature accounting for clothing and bedding [117]. Similar disconnects can also exist between the sensor and patient for the other environmental variables. Generally, we placed sensors by locations specified by the nurse manager so as not to interfere with the hospital's operations. However, the more accurate the sensed values can be to the perceived values, will likely benefit from further research. Our study contains patients from only one hospital, so we note it is not a representative sample of hospital environments or hospital patients elsewhere. Primarily, we anticipate our exploration to find potential avenues to improve hospital environments and quantify the environment's impact on patients. The literature demonstrates that the disruptive environmental fields are found to be different across different hospitals [58, 118], suggesting that a further study is required to make more substantial claims about the need for more hospitals to incorporate environmental sensors and how the sensors can help improve patient sleep and by extension patient recovery.

### 2.5.6 Future Work

We recommend that future researchers carefully consider what can be considered "ground truth" when running future experiments. While recording patient rooms with cameras can lead to higher data quality labels (e.g., construction lights outside the hospital bed), and patient family visits can then be more accurately recorded), we understand the difficulty of acquiring consent from the caregivers and the patients for video recording, as well as the difficulty in limiting the patient population to the sensor-installed room since hospital beds are often crowded and in great demand. We anticipate exploring more devices or methodologies that enable more convenient, private, yet still precise measures of activities occurring in hospital rooms to explore interventions that can help inform decision-making positive patient outcomes. We also look forward to studying more of the mechanics that correspond to the relationship between environmental fingerprints and medical events class. Further study is needed to verify that the environmental fingerprints are generalizable and not a product of the study hospital's specific machines, caretakers, or environment.

## 2.6 Conclusions

Our work explores the relationships between patient sleep, the environment, and medical events and provides hospital-specific lessons learned and environmental attributes to consider for future researchers to explore. While it is unsurprising that the environment dramatically impacts our sleep quality, our work demonstrates quantifiable significance in including medical and environmental information in analyzing momentary disruptive factors affecting patient sleep. It demonstrates the potential benefits of integrating environmental sensors and medical data to improve the patient's sleep environment.

# Chapter 3

# Maintenance

*The placement of sensors dictates not only the observable behavior but also the silent rate of signal transmissions. Modeling this silent rate enables applications that can identify locations in the building with lower silent rates and diagnose sensor transmitting sub-optimally. This chapter shows how contextual data about the walls and materials of a building paired with the location of gateways and sensors help us determine whether transmission rates are expectedly low due to the attenuation of the context.*

## Integrating building plan information to support fault detection of long-term energy harvesting sensor deployments

As the number of Internet of Things (IoT) devices continues to increase, energy-harvesting (EH) devices eliminate the need to replace batteries or find outlets for sensors in indoor environments. This comes at a cost, however, as these energy-harvesting devices introduce new failure modes not present in traditional IoT devices: extended periods of no harvestable energy cause them to go dormant, their often simple wireless protocols are unreliable, and their limited energy reserves prohibit many diagnostic features. While energy-harvesting sensors promise easy-to-setup and maintenance-free deployments, their limitations hinder robust, long-term data collection.

To continuously monitor and maintain a network of energy-harvesting devices in buildings, we propose the *EH-HouseKeeper*. *EH-HouseKeeper* is a data-driven system that monitors EH device compliance and predicts healthy signal zones in a building based on the existing gateway location(s) and building profile for easier device maintenance. *EH-HouseKeeper* does this by first filtering excess event-triggered data points and applying representation learning on building features that describe the path between the gateways and the device.

We assessed *EH-HouseKeeper* by deploying 125 energy-harvesting sensors of varying types in a 17,000-square-foot research infrastructure, randomly masking a quarter of the sensors as the test set for validation. The results of our 6-month data-collection period demonstrate an average greater than 80% accuracy in predicting the health status of the subset. Our results validate techniques for assessing sensor health status across device types, for inferring gateway status, and approaches to assist in identifying between gateway, transmission, and sensor faults.

## 3.1 Introduction

As buildings strive to be green and healthy [12], so too increases the need for indoor sensing of environmental conditions and occupant activity. Studies focusing on energy consumption and occupant comfort, performance, and well-being have demonstrated that continuous environmental sensing can aid building automation systems in adjusting the environmental settings to suit users' needs [119, 120, 121, 122, 123]. Since the needs of occupants are complex and multi-faceted, there is an increasing need for richer and more comprehensive sensors to provide multiple modalities of information about the user and for this data to be accurate and consistent.

Increasing the number of sensors while ensuring reliability presents competing challenges. Increasing the density and quantity of sensors suggests they should be smaller, cheaper, and easier to deploy. But ensuring reliable data suggests that devices should be sophisticated and hard-wired. The low-power embedded sensing community has largely focused on the first set of challenges, namely developing small, wireless sensors capable of instrumenting existing buildings.

As devices continue to reduce in size, they have started to swap larger batteries for smaller energy-harvesting power supplies [124]. Not only can harvesting outperform batteries when devices are smaller than a sugar cube [125], energy-harvesting increases the range of location for sensing versus wall-powered devices, and eliminates the periodic battery swaps needed for battery-powered devices. These traits make them attractive for dense but aesthetically pleasing retrofits in existing buildings.

As energy-harvesting devices become more accessible [126], and as such more used in buildings [127, 128], the set of challenges related to reliability and robustness become more pressing. While a small, "stick-on", and photovoltaic-powered sensor [129] is easy to deploy and quickly generates useful data, these types of sensors have three characteristics that are significant regressions from the mains-powered and BACNET capable sensors commonly found in buildings. First, they are dependent on the availability of harvestable energy. If their energy source disappears, for example a room is dark for an extended period of time, they will enter a hibernating state and stop transmitting data. Second, to enable low-energy operation, they typically use simple, unreliable wireless protocols. This means data may not be received even if the sensor successfully samples and sends its data. Third, intermittent energy availability and low-cost hardware can result in less consistent operation. For example, the sensor may have poor timekeeping and not sample at precise intervals. Each of these hinders the reliability of the overall sensing deployment, but together they present a significant challenge for long-term monitoring, and worse, they all tend to manifest with the same symptom: no data packets are received from the sensors.

To realize the upside of ubiquitous energy-harvesting sensors while managing the uncertainties they present, we propose a comprehensive monitoring system specifically for networks of energy-harvesting sensors and the unique challenges they present. Our system, *EH-HouseKeeper*, is a diagnostic system for energy-harvesting sensors that identifies faulty devices that require manual intervention, and supports planning for more effective future device placements to increase reliability.

To enable the monitoring, *EH-HouseKeeper* collects data from every energy-harvesting sensor and automatically creates a unique data-driven profile of expected behavior for each sensor. This is necessary because devices can vary widely. First, some devices transmit periodic readings, others only respond to events, and some are event-based but also transmit periodically if no event has occurred recently. Second, devices experience different harvesting conditions and will have differing amounts of available energy. Third, devices experience different RF environments and will successfully deliver packets at different rates. And fourth, slight differences in sensor hardware will cause otherwise identical sensors to behave slightly differently. By using the device's actual behavior, *EH-HouseKeeper* can compensate for these variabilities.

With the profile created, *EH-HouseKeeper* then provides a health score for each sensor based on how well the sensor is performing with respect to its expected behavior. This health score is then used to identify sensors that have failed and need to be either repaired or replaced, and not just devices that have been unable to harvest or have had a few lost packets.

Because *EH-HouseKeeper* has profiles of devices in the sensing deployment with a range of health scores, *EH-HouseKeeper* can also be used to predict the health score of future energy-harvesting devices installed in different locations in the same environment. *EH-HouseKeeper* uses a predictive model to estimate where sensors will perform well in the future. This can guide deployment managers on where to place devices to optimize performance, or on what level of redundancy or overprovisioning is required to obtain a certain level of sensing performance.

To demonstrate the efficacy of *EH-HouseKeeper*, we test it using an in-building testbed with more than one hundred energy-harvesting sensors of various operating modes and sensing modalities. Due to the size of the testbed, there are several gateway devices distributed throughout the space that collect the wireless packets from the sensors, and each sensor may transmit to one or more gateways. *EH-HouseKeeper* must consider this gateway deployment as well, and must account for gateway failures when assigning health scores to individual sensors.

We run *EH-HouseKeeper* during a six-month study and observe its performance. We find that:

- Significant data loss can occur even when both the gateway and the EH sensor are working correctly.

- It is possible to calculate a comparable signal health score for a mixed periodic and event-triggered sensor using the device's Largest Heartbeat Interval (LHI).

- It is possible to automatically and accurately predict data packet loss due to signal attenuation given the building plan.

- The prediction method can accommodate a variety of different device types with varying intervals of heartbeat and event-trigger conditions.

- The average prediction accuracy for healthy signal zones is greater than 80% across all investigated device types.

## 3.2 Related Work

Because of its lower cost during upfront installation and better scalability of maintenance compared to battery and mains-powered devices [27], an energy-harvesing (EH) sensor based architecture was proposed as an ideal infrastructure for building monitoring and event detection [130]. *EH-HouseKeeper* is built on top of this architecture, exploring new challenges on dependability for a network of EH sensors.

Laprie provides a framework that we adapt for describing dependable computing, including a nomenclature to help distinguishing between fault, error, and failure [131]. Kavulya et al. extends on this nomenclature and describes different diagnosis techniques, limitations, and examples [132]. Notably, she explains how rule-based techniques are human-interpretable and extensible but difficult to maintain at scale; statistical techniques require little expert knowledge but might not distinguish legitimate changes in behavior; and machine-learning techniques automatically learn profiles of system behavior but can suffer from the curse of dimensionality when the feature set is too large.

To further analyze the reliability of our sensors, we explored works that modeled the effects of radio signal attenuation in an indoor environment [133, 134, 135, 136]. The literature points out a clear relationship between radio signals and indoor factors such as distance, number of walls, wall-depth, and wall material. However, the studies were mostly conducted in a static setting and for a short time-period. This makes it difficult to adopt the findings to a naturalistic setting, where the movement of people and furniture could add noise into the system. The difficulty is further increased when the status of the receivers are variable. Thus, we designed *EH-HouseKeeper* so that it can detect receiver status changes and account for them for future predictions.

We found a longitudinal study of radio signal attenuation for an experiment recording moisture content in an outdoor environment using battery-powered devices [137]. The study demonstrates a clear relationship between the modeled signal attenuation and reduction in periodic device transmission probability via silent rates. However, the devices used in the study are outdoor periodic sensors transmitting every 10 minutes while the sensors used in our study are a mixed event-triggered and periodic sensor. Additionally, the experiment considers a scenario with only one receiving antenna. Therefore, we propose a slight modification to the silent rate (signal health score) to account for the hybrid sampling of our sensors. Furthermore, we build on top of this signal health score to model data loss in a multiple receiver scenario in an indoor setting.

## 3.3 Methodology

In order to investigate the feasibility to model data loss through signal attenuation and extend the prediction to similar building spaces, the methodology section is divided into subsections of sequential order. Section 3.3.1 describes the geometry of our testbed, documenting the location and related specification of each of the deployed sensors. Then, because our study deals with multiple divisions of time, Section 3.3.1 defines the different divisions of time we use, and details when and what gateways were installed on the timeline. In Section 3.3.2, we build upon the defined time definitions to derive a signal health score calculation that resists bias caused by event-triggered sensing. Section 3.3.3 provides background information on how signal attenuation is modeled for an indoor environment, which is demonstrated to result in measurable data packet loss [137], which we can now detect using the derived signal health score. Lastly, Section 3.3.4 describe the details of how we use feature-representation learning to predict future healthy signal locations for the different aforementioned spatial and temporal arrangements.

### 3.3.1 Testbed Overview

The testbed is embedded within a laboratory and office space complex of approximately 17,000 square feet at a university and includes occupant-based wearables, interactive mobile robots, and comprehensive environmental sensors. The testbed is designed to support research on occupant behavior and new occupant-focused building control techniques through the capture of data associated with several dimensions of variability in human-building interactions. While more than 250 different types sensors (wired, battery-powered, and EH) have been deployed in the space to date, only the EH sensors with location information are considered in this paper. The testbed is supported by a generic gateway platform in a one-hop network that stores the received data in a cloud-hosted time-series database. Figure 3.1 documents the gateway location, device location, and relevant device specifications. The gray circles drawn around the gateways mark a 25 meter radius.

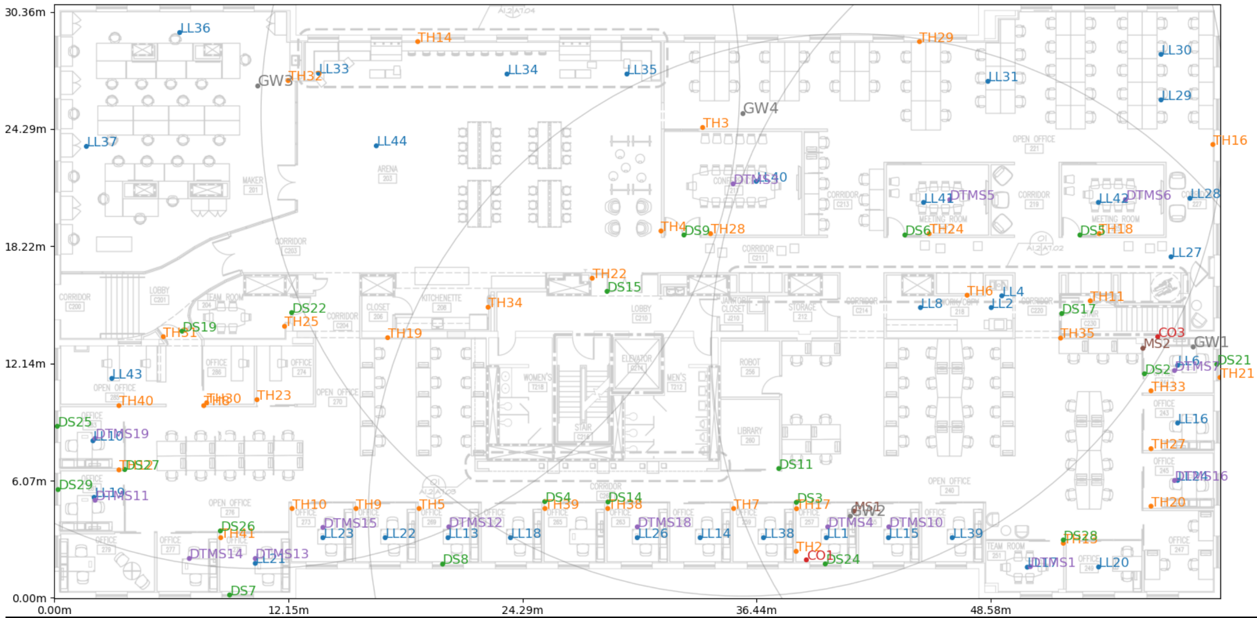


| Device | Prefix | Units | Transmission Radius | Operation Time | LHI |
|---|---|---|---|---|---|
| Swarm Gateway [138] | GW | 4 | N/A | N/A | None |
| EnOcean Light Level Sensor ELLSU-W-EO | LL | 40 | 25 meters | 80 hours | 30 minutes |
| Pressac Mini Temp Humidity Sensor (Discontinued) | TH | 37 | 30 meters | 4 days | 15 minutes |
| EnOcean Wireless Door/Window Sensor ExT-MDCCP | DS | 27 | 20 meters | 5 days | 25 minutes |
| Echoflex Dual Tech Ceiling Mount Sensor MOS-DT | DTMS | 17 | 24 meters | 7 days | 100 seconds |
| Pressac Wireless CO2 60.CO2 SLR TMP HUM | CO | 2 | 30 meters | 5 hours | 15 minutes |
| Illumra Motion Sensor E9T-OSW | MS | 2 | 25 meters | 80 hours | 30 minutes |

Figure 3.1: Projected Device Plan and Descriptions

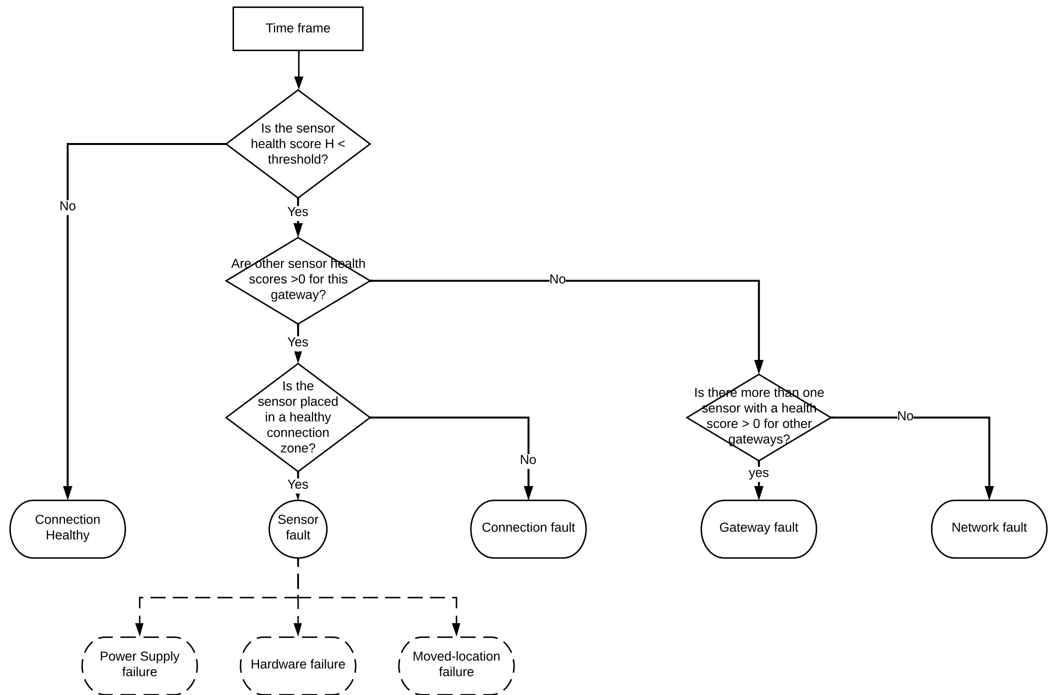


Figure 3.2: Fault Identification Flowchart

### Time Definitions

Because there are three different subdivisions of time used in this paper, we will clarify them here, from longest to shortest:

- **Time range**, where we describe the encompassing datetimes for a specific gateway configuration. In our study this is a variable, and further described in Table 3.1. For example, during time range $T_0$, only one gateway was installed.
- **Time period**, where we describe the division within a time range, used for signal health calculations. In our study this is a constant set to 24 hours.
- **Time frame**, where we describe the subdivisions within a time period. In our case we use a constant equal to the device's corresponding Largest Heartbeat Interval (LHI), as described in Figure 3.1. For example, the time frame used for light level sensor health score calculations is a constant equal to 30 minutes.

Table 3.1 details the divisions of time as well as which gateways were on during which time range.

After $T_2$, the space had close to zero occupants due to COVID-19 restrictions.

| Name | Start and end Date | Gateways online |
|---|---|---|
| T_0 | [2020-01-01 , 2020-02-25] | $G_1 = \{GW2\}$ |
| T_1 | [2020-02-26 , 2020-03-18] | $G_3 = \{GW1,GW2,GW4\}$ |
| T_2 | [2020-03-19 , 2020-04-28] | $G_4 = \{GW1,GW2,GW3,GW4\}$ |
| T_3 | [2020-04-29 , 2020-07-01] | $G_4 = \{GW1,GW2,GW3,GW4\}$ |

Table 3.1: Time ranges and gateways configuration

### 3.3.2 Calculating Device Signal Health Scores

To account for variable heartbeat intervals when doing health score calculations, we elect to use a device's Largest Heartbeat Interval (LHI), the largest interval of time after which a data point is expected. For instance, for the EnOcean Light Level sensor, which heartbeats randomly between 20 to 30 minutes, we elect to use 30 minutes. We use the heartbeat interval as defined by each device's corresponding datasheet as their baseline. For our study, we did not customize any configuration on our devices to sample at different intervals.

To calculate whether or not a periodic sensor is transmitting correctly for a time period, our basic approach is to divide the total number of received data points by the total number of expected data points for every time frame in the time period to arrive at a health score:

$$H = \frac{1}{N} \sum_t \frac{r_t}{e_t}$$

Where $H$ is the overall health score for the time period, $t$ is the index of the time frame within that time period, $N$ the number of total time frames for the time period, $r_t$ the number of received data points for that time frame, and $e_t$ the expected number of received data points for the time frame. This basic method is straightforward, but if $e_t$ is lower than the LHI and therefore zero, the score is undefined. Similarly, if $e$ is not a multiple of the LHI, the subsequent rounding results in loss of information.

To solve this, we subdivide the time period into time frames that are equal to the device's LHI.

The expected number of received data points $e$ is then always one, giving us:

$$H = \frac{1}{N} \sum_{t} r_t$$

However, in the mixed sensing scenario where the sensor is both periodic and event-triggered, and the event-triggered data point resets the heartbeat interval, doing so could allow event-triggered data points in one time frame to bias the entire time period (i.e. $r$ could be greater than one). As such, it is important to also cap the transmission count for each LHI frame to arrive at:

$$H = \frac{1}{N} \sum_{t} min(1, r_t) \tag{3.1}$$

The silent rate of the time period for the device as described in [137] is then just $1 - H$. It might be helpful to note that it is impossible to completely disambiguate between heartbeat and event-triggered data points for health score calculations in this scenario since it is theoretically possible for a sensor to be event-triggered at the start of every heartbeat interval.

### 3.3.3 Radio Signal Attenuation

The Keegnan-Motley model of logarithmic signal loss $L$, as described by [135], is:

$$L(d) = L_{FS}(d) + n_w L_w + n_f L_f \tag{3.2}$$

With $L_{FS}(d)$ the theoretical loss in free space for an isotropically[1] radiating antenna, $d$ the distance between transmitter and receiver, $L_w$ attenuation per wall, $L_f$ attenuation per floor, $n_w$ number of traversed walls, and $n_f$ number of traversed floors. This model has further been shown to be adjustable to account for the thickness of the wall [139]. While signal attenuation is then generally calculated using a constant attenuation per unit path length $\alpha$, in our feature preparation section (Section 3.4) we detail our process in tracing the discrete partitions and free space for each device-to-gateway path.

[1] not varying in magnitude according to the direction of measurement

### 3.3.4 Feature Representation Learning

Feature representation learning is commonly used for applications such as natural language processing and one-shot image recognition [140]. Once discriminative features have been learned, the predictive power of the network can be applied to new data. In this paper, we resort to a feedforward siamese network [141] to learn discriminative features from the raw sensor-to-gateway path features in order to predict the signal health status of future sensor locations. A feedforward siamese model consists of $L$ feedforward layers each with $N_l$ units. For the first $L-1$ layers, each is followed by a ReLU activation layer. For the remaining layers, each is followed by a sigmoid layer. Our model takes a pair of sensor data as inputs. Let $h_{1,l}$ represents the hidden vector in the $l$-th layer for the first twin and $h_{2,l}$ denotes the same for the second twin. A non-negative function is deployed after each activation layer to restrain the learned hidden vectors are non-negative. Hence, the operation at the $l$-th layer takes the following form:

$$a^k_{1,m} = \max(0, \sigma(W_l h_{1,l} + b_l))$$

$$a^k_{2,m} = \max(0, \sigma(W_l h_{2,l} + b_l))$$

, where $\sigma$ denotes the ReLU activation function, $W_l$ and $b_l$ represent weights and bias in the $l$-th layer respectively, $l \in \{1, \ldots, L-1\}$.

Once siamese twins, $h_{1,L-1}$ and $h_{2,L-2}$ are outputted, the induced distance metric is computed by the final layer. More specifically, the prediction vector is given by:

$$P = \sigma(\sum_{j=1}^{N_L} \alpha_j |h^j_{1,L-1} - h^j_{2,L-1}|) \tag{3.3}$$

Here, $\alpha$ denotes weights in the final layer,and $\sigma$ represents the sigmoid activation function.

Let $y(x_1, x_2)$ be the vector that contains the label for a pair of data sample, where $y(x_1, x_2) = 1$ if $x_1$ and $x_2$ are from the same class and $y(x_1, x_2) = 0$ otherwise. The network is optimized by minibatch gradient descent to minimize the following loss:

$$L(x_1, x_2) = y(x_1, x_2) \log P(x_1, x_2) + \\ (1 - y(x_1, x_2)) \log(1 - P(x_1, x_2))$$

In this study, we aim to learn discriminative features to distinguish healthy sensors from unhealthy ones based on their geometric information. The two classes in our case are healthy and unhealthy sensors. The siamese network takes raw features, which will be explained in the following section, of a sensor $i$ as inputs and outputs a feature vector $h_{x_i,L-1}$. Once the siamese network is trained, we can apply it to generate features for sensors that are currently not installed, which assists us in predicting the quality of newly proposed sensor locations.

### 3.3.5 Assumptions

This chapter assumes that the buildings where sensors will be installed will have an available building plan that can be digitized with minimal effort. This assumption might not be generalizable to older buildings with manual drawings. Further, the study assumes that a stable harvesting condition is satisfied in the training set. This assumption is further elaborated on in Section 3.8. It is supposed that the installers of sensors will first collect a profile for a sensor's typical interaction with building materials when energy is not constrained and then use the ideal profile to model the silent transmission rate of other sensors. Finally, this study presupposes that attenuation by building material is relevant (i.e., sensors are transmitting in the ultra-high-frequency range or higher [2]). Since attenuation generally scales with frequency, the work in this chapter is likely not generalizable to transmissions at lower-frequency ranges.

## 3.4 Feature Preparation

**Gateway Status** Using the health score calculations, we defined the time period to be uniformly 24 hours for each device to represent a realistic response time for us to investigate a device failure. We used the LHI of each corresponding device as their time frame. Figure 3.2 shows a flowchart for our fault identification process.

[2] See https://www.enocean.com/en/technology/radio-technology/

Because we did not explicitly store data points of the gateway status, we assume that the gateway is down if for that time period, no values were transmitted from that gateway by any sensor. By extension, we assume that the network is down for the time period if all gateways did not transmit data. While this could be sufficient in a one gateway scenario, there is a minute possibility of incorrectly classifying gateway down status in a multiple gateway scenario if the data point was pushed to the database by another gateway for the time frame for all of the devices in range. Figure 3.3 shows our identification of a gateway powered by an occupancy-controlled outlet through the health score of all the devices.

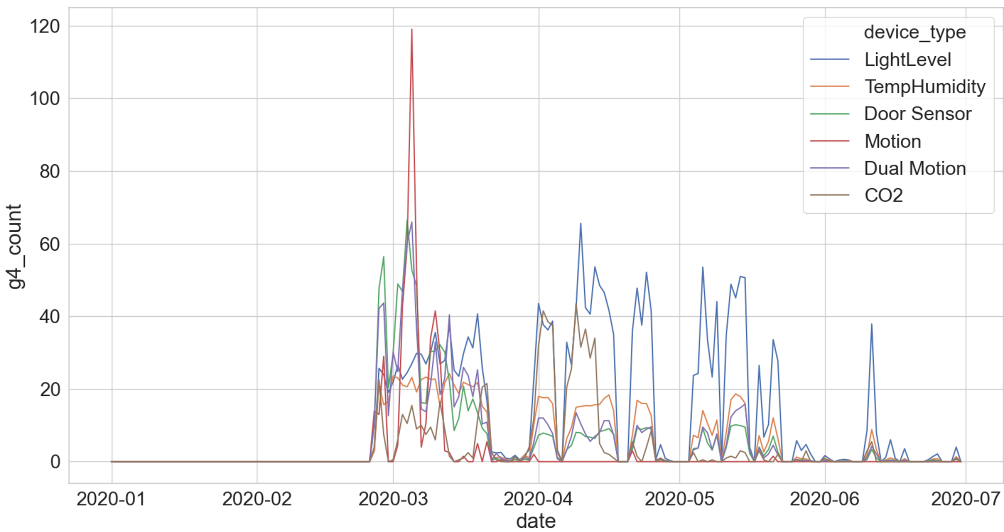


Figure 3.3: Occupancy-controlled-outlet plugged gateway 4 related transmission count per device type over time

**Sensor to Gateway Path** Using the health score calculations, we sought out to explore a relationship between sensor health score, sensor-to-gateway distance, and the wall profile between the sensor to wall.

Plotting the distance to health score relationships for the light level sensors, shown in Figure 3.4, we observed a decline in signal health score over distance. We observe a similar trend for the other device types as well. This indicates that there is indeed a measurable reduction of data transmissions across all types during the one gateway scenario $G_1$. We observed a similar trend when counting for number of walls traversed, since generally the longer the distance between the sensor to the gateway the more walls were traversed.

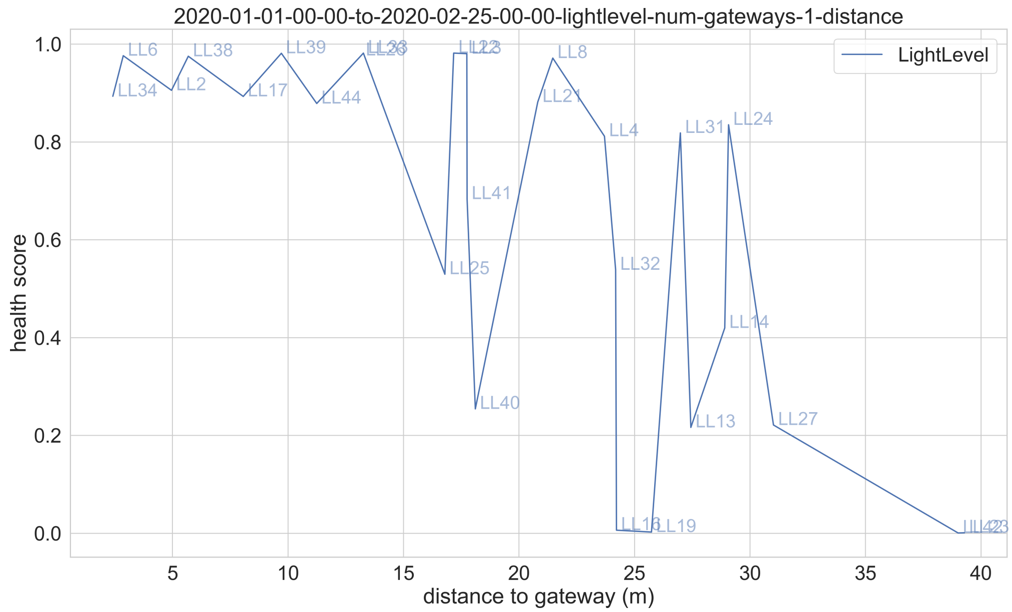


Figure 3.4: Relationship between signal health score and distance to gateway for light level sensors during $T_0$

Having observed that data loss does occur, we further employed ray-tracing on a plan from the sensor to the gateway. Codifying the wall depth, material and air space using the traversed pixel colors, Figure 3.5 demonstrates the example wall waveform from point A to point B. We counted only signals that exceed a threshold of 0.5 (i.e. when the trace hits the corner of two materials). Once the features are ready, we can represent them using a siamese network.

In the next section, we detail how we adopted this information into our machine learning model that accounts for all the different device to gateway raytraces.

## 3.5 Experiment and Data Preparation

For this study, additional gateways were installed over time to explore whether or not we could improve the health scores of our devices, starting with a single gateway. We used our assumptions (see Section 3.4) of the gateway status to evaluate whether or not the gateway was on or off during that day, and when it was first installed. In addition, during our light sensor installation, we installed LL4 and LL8 in a room with little access to daylight and their harvesting surface pointed away from the artificial light source as a test case for our detection system. Table 3.2 summarizes the columns of the data frame with an example row.

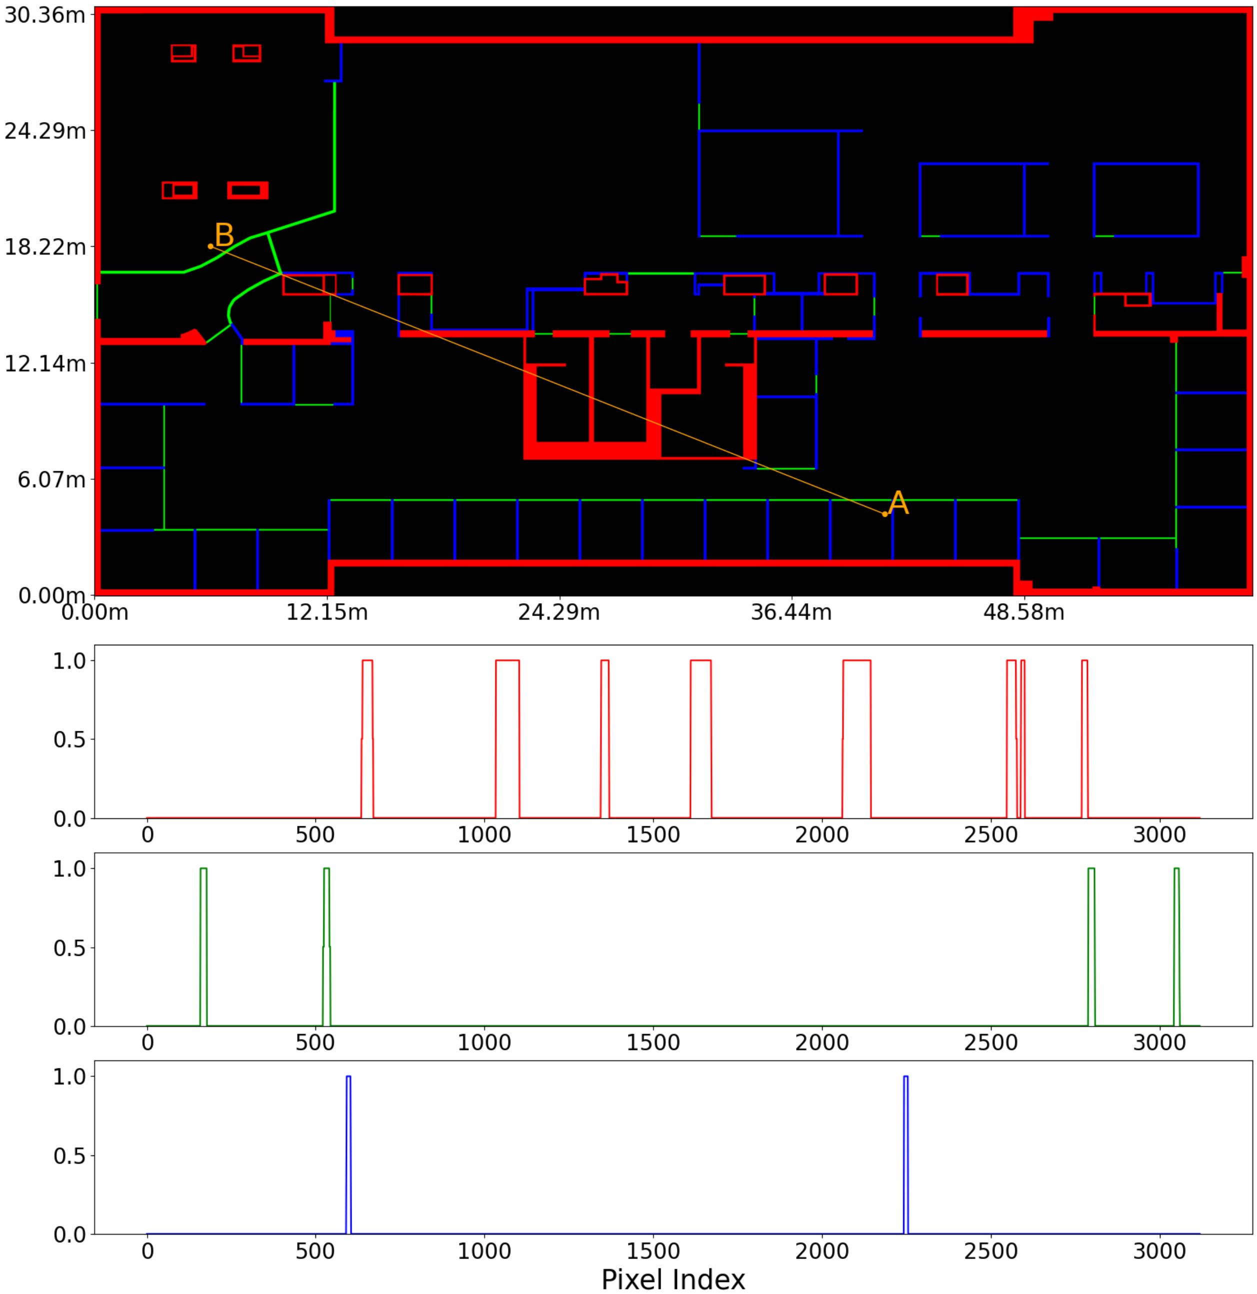


Figure 3.5: Example ray-trace with its corresponding waveform from point A to point B. The number of pulses corresponds with the number of walls traversed, the width of the pulse corresponds with the thickness of the wall, and the color of each signal corresponds with a type of building material.

The *wall_arr* column (Shown in Table 3.2) describe characteristics of the building between the sensor and the gateway, where element 0 is the count of the air pixels, element 1 the count of the red wall pixels, element 2 the count of the green wall pixels, and element 3 the count of the blue wall pixels.

| Name | Example value | Description |
|---|---|---|
| date | 2020-01-31 | YYYY-MM-DD time description for the device |
| device_type | Light Level | Array description of the device category |
| device_name | LL1 | Identifier of the device |
| g[n]_wall_arr | [722, 156, 0, 5] | Description of the building elements between the gateway and the device in [a,r,g,b] for gateway n |
| g[n]_count | 0 | Number of transmission to gateway ¡n¿ |
| g[n]_dist | 57.28 | Distance to gateway[n] in meters |
| health_score | 0.7 | Health score for the day for the device |
| timing | [2020-03-16T01:... | Array detailing the specific timings of each data point for the day |
| g[n]_on | True | Whether or not the gateway[n] was on |

Table 3.2: Data set column descriptions

For our final feature set, we combined the elements of the *wall_arr* for all four gateways (16 elements), and added 8 elements that is the on off statues of each four of our gateways to arrive at a total of 24 features. For ease of comparison, we arbitrarily classify any sensor with a threshold above a 70% signal health score as healthy. The choice of this threshold for future experiments will likely depend on the research question and the lab's capacity for maintenance.

## 3.6 Results

### 3.6.1 Aggregated Signal Health Monitoring

While we found that the proposed signal health score calculations does mitigate the effects of event-triggered data points from biasing the overall health score, we also found it important to note that the score does **not** completely remove the effects of additional event-triggers. For example, for magnetic contact sensors, the health score could be amplified if the installed door is more frequently used. The additional event-triggers make it more likely for the transmission to register, even if the device is located in a more attenuated zone. Therefore, categorizing the occupancy schedule of the space, and studying the system during a period of time with no-occupancy can provide a cleaner reading as to whether or not the health score is due to artificial amplification. Additionally, while this signal amplification can be readily isolated in occupant-triggered devices, as seen in Figure

3.6, the effects are harder to isolate for sensors that are triggered by environmental conditions. To evaluate the signals without event triggers from the temperature humidity sensor, for example, will require a controlled environment of less than 2% humidity and 0.6 degrees Celsius fluctuations. In order for the event-driven amplification of the device signal health score to impact the composite daily health-score, however, the threshold for the device-trigger will need to be exceeded more than once per time frame, across multiple time frames, and also be registered in place of the signal that otherwise would not have been registered.

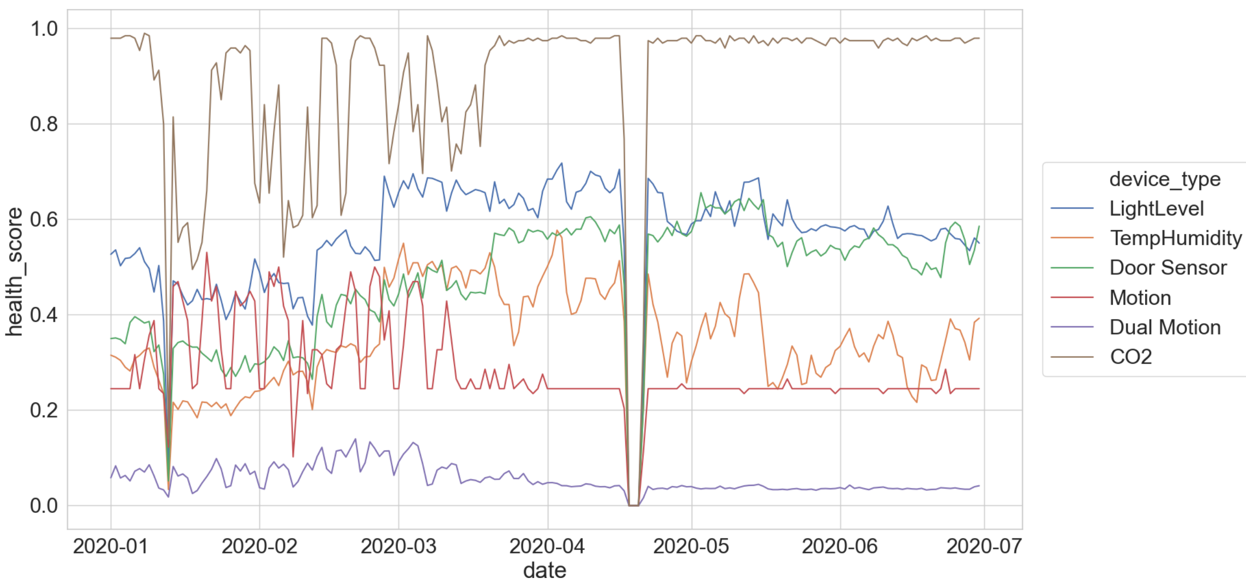


Figure 3.6: Overview of aggregated signal health score traces per type, demonstrating a network-down period in late April

## 3.7 Model Evaluation

To verify our process, for each device type, for each time range, for each of 100 iterations we randomly masked 25% of the devices rounded down as the testing set, using the remaining as the training set. We then trained a representational encoder with the training data to encode the test data features fed into another 100 independently trained classifiers. Finally, we aggregated the classification results using the encoded features to predict whether or not the device location is considered healthy. Our results shown in Figure 3.7 indicate a stabilization of an average of greater than 80% accuracy when predicting health scores for the masked sensor locations over 100 runs. The results of all the runs using a 1 layer linear classifier and a decision tree classifier are summarized in Table 3.3 and Table 3.4, respectively. For all of our following analysis, we elect to use the decision tree classifier because it gives us a better score than the one layer classifier overall.

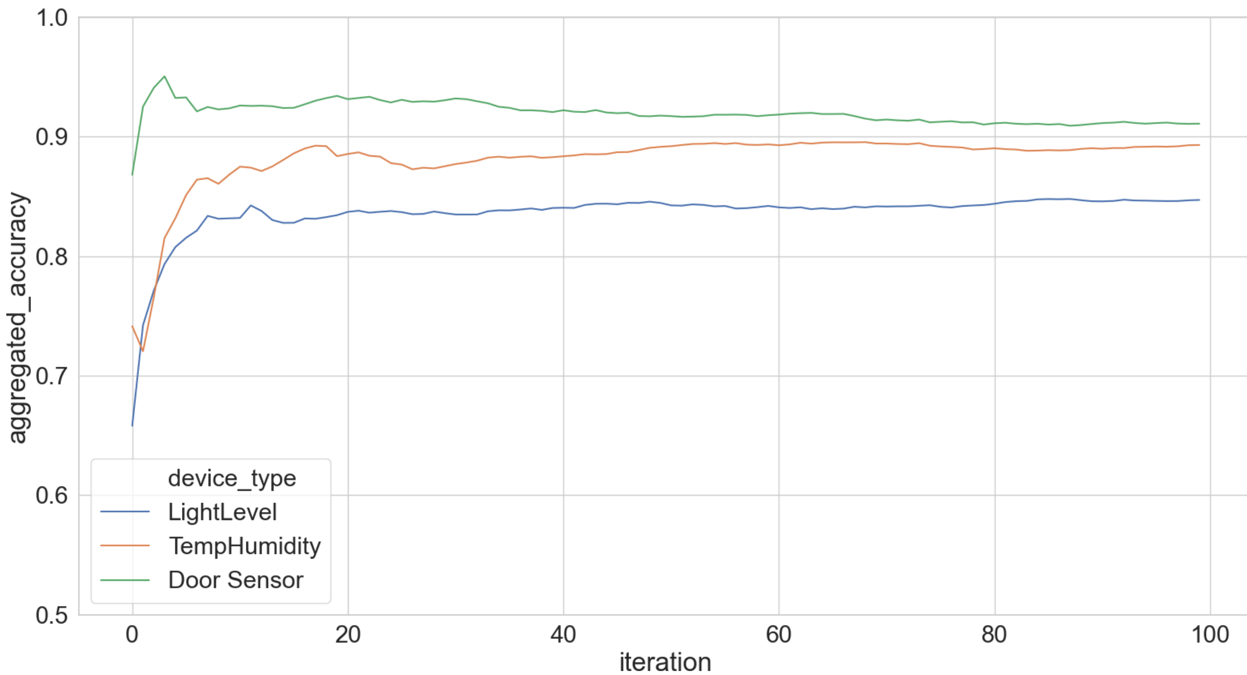


Figure 3.7: Average accuracy over iterations per type for $T_1$ over 100 runs

| Time Range | Device Prefix | Accuracy | Precision | Recall |
|---|---|---|---|---|
| $T_0$ | DS | 0.87 | 0.85 | 0.85 |
| | LL | 0.83 | 0.83 | 0.82 |
| | TH | 0.89 | 0.83 | 0.84 |
| $T_1$ | DS | 0.91 | 0.91 | 0.91 |
| | LL | 0.86 | 0.85 | 0.83 |
| | TH | 0.89 | 0.90 | 0.89 |
| $T_2$ | DS | 0.88 | 0.87 | 0.86 |
| | LL | 0.82 | 0.81 | 0.80 |
| | TH | 0.79 | 0.76 | 0.72 |
| $T_3$ | DS | 0.88 | 0.88 | 0.87 |
| | LL | 0.94 | 0.94 | 0.94 |
| | TH | 0.88 | 0.82 | 0.81 |

Table 3.3: 100 Run Average Decision Tree Classifier Results

| Time Range | Device Prefix | Accuracy | Precision | Recall |
|---|---|---|---|---|
| $T_0$ | DS | 0.81 | 0.77 | 0.77 |
| | LL | 0.81 | 0.82 | 0.79 |
| | TH | 0.89 | 0.81 | 0.79 |
| $T_1$ | DS | 0.91 | 0.92 | 0.91 |
| | LL | 0.84 | 0.83 | 0.81 |
| | TH | 0.89 | 0.90 | 0.88 |
| $T_2$ | DS | 0.88 | 0.87 | 0.86 |
| | LL | 0.80 | 0.80 | 0.79 |
| | TH | 0.76 | 0.73 | 0.71 |
| $T_3$ | DS | 0.81 | 0.81 | 0.80 |
| | LL | 0.91 | 0.91 | 0.91 |
| | TH | 0.76 | 0.67 | 0.68 |

Table 3.4: 100 Run Average 1-Layer Linear Classifier Results

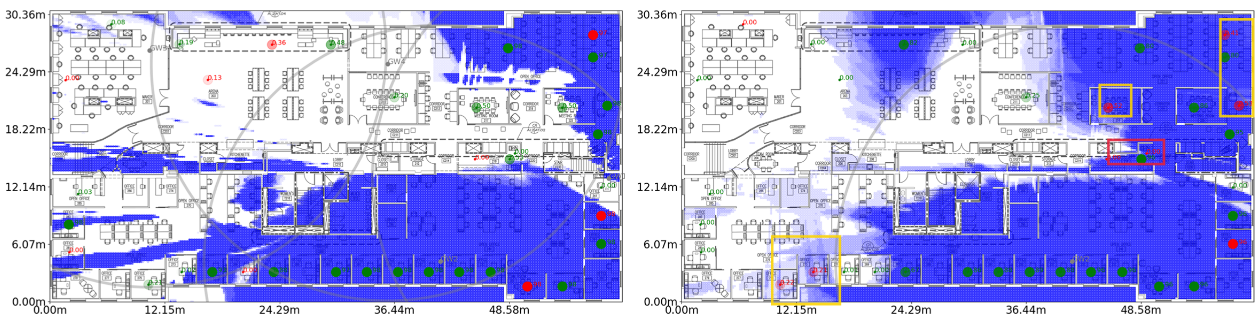


Figure 3.8: $>$ 90% accuracy sampling during $T_3$ (left) versus $<$ 60% accuracy sampling during $T_0$ (right), where the red circles represent the masked test sensors and the green circles the training sensors. The yellow rectangle areas indicate additional attention required, and the red rectangle areas indicate maintenance required.

### 3.7.1 High Accuracy Versus Low Accuracy Assessments

To assess the validity of our model as well as help us determine where are the topographically similar areas with better signal health, we generate a value using the features at each pixel space for its probability to be a healthy location. We demonstrate in Figure 3.8 a comparison between one of the highest-performing sampling and one of the lowest-performing samples for the light level sensors.

The masked sensors are marked in red, and the training sensors are marked in green. The alpha of the red and green represents the health score for the sensor, which is also labeled next to the sensor. The blue background color represents the aggregated prediction percentage for the pixel location. When the model is predicting accurately, as shown in the left image of Figure 3.8, then that means that there are no misalignment between expected signals lost and the actual signals loss. Large misalignment, as shown on the image on the right, indicate that the poorly performing model requires additional diagnosis to detect: 1) whether or not the low signal health score sensor is occurring at the edge of the healthy zones, and 2) whether the signal health is higher or lower than anticipated. When a sensor signal is poor in an area where other sensor signals are healthy, as in the red rectangle, then there is a larger likelihood of abnormal transmission patterns and the sensor should be marked for maintenance. The detected abnormal sensors match the test sensors we initially installed (LL4, LL8, as described in Section 3.5) when they are not selected as part of the training set. The sensors marked in the yellow rectangle areas, while performing sub-optimally, can still be permissible since they are operating out of range of the device specifications or at the edge of the attenuated zones.

## 3.8 Discussion

The value of our current models relies on the assumption that most of the EH devices are operating in stable harvesting conditions. For example, the predicted healthy signal areas using $T_3$ likely included data loss due to the reduced lighting schedule. Additionally, the accuracy of the prediction also relies on the existence of similar topographically placed sensor.

As seen in Figure 3.9, during normal operations of the lab, even when multiple gateways are within the transmission range of the sensor, the topography of the space influences the overall received signal and data can be loss. This influence is sometimes the difference between losing some of the data, and losing all of the data.

Furthermore, since the only way to check if the EH sensor is operating in sufficient lighting once deployed in a dynamically lit area is to check if there are still data transmissions after its operation time (how long it can operate in darkness), disambiguating data loss due to signal attenuation can help the administrator diagnose between expected data loss and data loss that requires maintenance. *EH-HouseKeeper* is demonstrated to capture this discrepancy for a variety of gateway configurations and make indications for which sensors actually require maintenance.

Accordingly, a more proactive strategy could have been to install and sample the gateway and

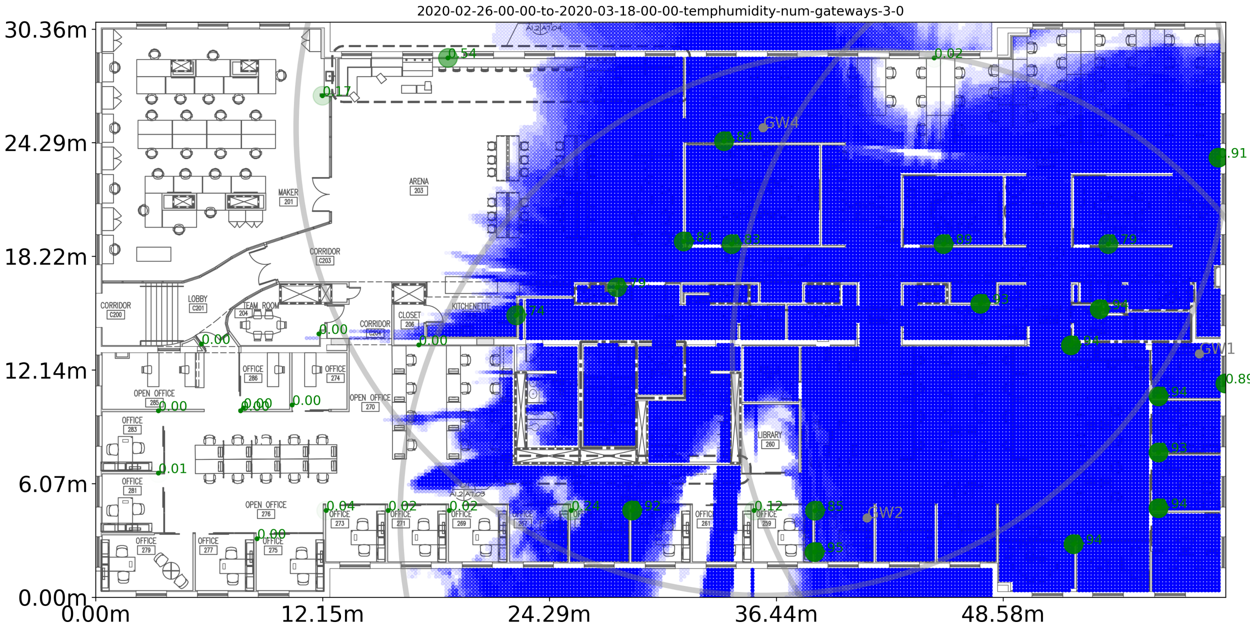


Figure 3.9: Using $T_1$ to predict healthy signal zones for $G_3$ for Temperature Humidity Sensors, showing the predicted signal healthy zone less than the prescribed radii.

EH devices within the operation period of the all the installed devices and sample in that time period to eliminate unhealthy signals due to power issues. With more sensors to be installed, however, the solution would be infeasible, especially if further gateway location optimizations are being performed or sensors being installed at different times. A workable solution, as in our case, then, is to install the charged EH sensors and sample them during normal operations of the building. Conceivably, another solution to further disambiguate connection-related data loss from power-related data loss would be to control the energy source (i.e. lighting) of the space and see if increasing the source output alters the device health score in the location.

While placing EH sensors in range of multiple energy sources (i.e. in view of a window (s) and under artificial light(s)) would assist in the longevity of data transmissions, the power supply of the EH sensors is still variable. For example, for light EH sensors, the consistency of the artificial lights are dependant on the chronotype of the occupants for occupancy sensor triggers, the time periods (i.e. holidays, weekends, workdays), and the weather. While one might argue that a steady source of lighting is guaranteed because it is only relevant to collect data while the occupant is present and therefore the lights are on, it could be worthwhile to consider that the lighting could be sub-optimal in a way that the degradation of data transmission might only be noticeable after a few month's time. Additionally, the variability of the power source while the occupant is absent could also affect the device's transmissions when they return. This is especially true if the time frame of interest lie between the occupant's arrival and the sensor's charge up time, or if abnormal behavior of the occupants increases the energy required to detect the events. Further studies into EH sensors in sub-optimal harvesting conditions is needed to better understand the severity and relevance of this data loss.

Also relevant to future deployments, we mirror the findings of Wagner et al. in chapter 6 regarding the importance of adhesives for sensor installation [6]. Some adhesives we installed degraded over months, and it took additional efforts from the residents in the space to recover. We propose applying more adhesives than considered necessary to reduce future maintenance efforts.

Finally, some outlets do not function as a consistent power source and have their own power schedule (also noted by Hnat et al. [27]). This information is harder to detect in a multiple gateway scenario because the drop in total received data corresponds with lowered occupant activities. If possible, implementing a heartbeat logging mechanism to track the gateway itself on the database can help diagnose the cause of a sensor signal health score drop for future deployments.

## 3.9 Limitations and Future Work

### 3.9.1 Noise Introduced in a Naturalistic Setting

The distance and trace used in our calculations are projections onto a 2D plane, so it does not encompass the complexities of the 3D environment. For instance, additional work needs to be done to extend the system to encompass multiple floors. In addition, *EH-HouseKeeper* does not account for any of the discrepancies between the plan drawing and the real-world environment, nor does it account for any signal attenuation due to the presence of furniture or occupants. The timing of the data could also be further filtered. For example, distinguishing between daytime and night sensor behaviors could further improve our model.

More work can also be done to scrutinize the data value itself (i.e. to identify non-fail-stop failures such as calibration drifting). For example, do those event-triggered data points match the data sheet described value thresholds? Is there a large unaccounted for discrepancy between two data point values in the same proximity? Even in the same zones, the orientation of the sensor device could have a dramatic affect on how much light it receives from the surrounding environment. While the location and orientation can be further optimized by calculating metrics such as Useful Daylight Illuminance for the vertical or horizontal surface that the sensors resided, for our installation we mainly faced the energy harvesting area towards sources of light (i.e. the window, artificial light source).

### 3.9.2 Trying Out New Locations

One future goal for *EH-HouseKeeper* is to start learning patterns for wall typology that we can transfer the attenuation patterns for other gateway locations in the same building, or different buildings. Doing so potentially allows us to reduce the total number of gateways used while increasing the signal health scores across the different devices. In addition, if we can validate model for different spaces using the same techniques, we can begin to optimize for gateway and device location virtually before deploying the system into a new environment.

### 3.9.3 Relating the Sensors to the Occupants

Since the number of sensors and gateways to deploy are limited, considerations must be made about which space is more important to study, and therefore where is the optimal location for the devices and what is an appropriate signal health score threshold. Simply improving the overall coverage of the EH sensors by changing device locations might not sufficiently collect data from true areas of interest that serve the occupant community (e.g. which what space an occupant feels the most creative, the most productive, and why). Moving forward, we plan to conduct interviews with the residents directly within the lab, to investigate what are the most desired attributes within a space as judged by the residents, and to investigate if there are any quantifiable patterns for these spaces.

While there is still more work be done on scrutinizing both the quantity and quality of the data, ultimately, deploying a system like *EH-HouseKeeper* that can continuously check for network, gateway, and sensor compliance and notify the administrators of unexpected faults seems to be a prerequisite to scaling up the number of EH sensors installed, or even just to carry out longitudinal studies with existing EH sensors.

## 3.10 Conclusion

Using energy-harvesting sensors in indoor environments is a promising technique for enabling data-driven and real-time optimization in the millions of existing buildings already constructed. However, these sensors add uncertainty to the data collection process due to intermittent energy availability and unreliable wireless connectivity. To help building managers successfully adopt these emerging sensors, we present *EH-HouseKeeper* to identify when a sensor has actually failed and to help guide deployment upgrades over time. The health score provided by *EH-HouseKeeper* enables building managers to rapidly correct faulty devices without the overhead of periodic inspections or unnecessary maintenance. We demonstrate over the course of half a year in a sensor-rich environment that *EH-HouseKeeper* is effective, and show how it can help guide future deployments. *EH-HouseKeeper* is an important step in making energy-harvesting sensors truly viable at the large scale needed to reduce the energy consumption and increase the occupant utility of the world's buildings.

# Chapter II

# Conclusion

This dissertation integrates contextual data to improve instrumentation, utility, and maintenance for living labs. The *instrumentation* chapter (Chapter 1) combines modern building simulation techniques and joins them with key algorithms to optimize the deployment of IoT sensors. The work demonstrates the approximate optimal location and quantity of sensors to deploy for its application. It also showcases the methodology as a potential venue to limit privacy overreach by providing a score equal to possible inferences. As more and more sensors are being integrated into buildings, the distinction between data ownership and for whose benefit the data is being used will only grow more critical [142]. This chapter demonstrates the first step towards scalably elucidating these trade-offs. After sensor deployment, in the *utility* chapter (Chapter 2), commercial off-the-shelf environmental sensors and medical data from a real hospital were combined to demonstrate statistically significant variables and how those variables can improve patient ($n = 38$) sleep in a hospital environment. The work illustrates how interdisciplinary collaborations and signal processing schemes coupled with statistical models can be combined to inform patient care and recovery. Finally, the *maintenance* chapter (Chapter 3) explores maintaining energy harvesting sensors for long-term environmental sensing. Building information alongside knowledge of signal attenuation are utilized to predict data package loss and diagnose sensing issues. Together, by integrating contextual data *before, during, and after* sensor deployments, this dissertation contributes to practical improvements of living labs and reduces the gap between architectural data, sensors, and computation. For future directions, this dissertation anticipates using the concept of "Observability" [52], combining physical and virtual sensors to make inferences. Additional utility is predicted if such a workflow was brought together in a similar open-source nature as Honeybee and Ladybug [49]. Further, finding ways to sense environmental attributes at larger scales and ways that the indoor environment might relate to more significant societal challenges, similar to [143], could be a useful metric that can help inform local policy decision-making and lead to broader societal impacts. As sensors and simulation platforms become more accessible, the gap between digital and physical, bits and atoms [144] will only continue to shorten, hopefully, towards a better and more healthy future.

# Bibliography


[1] Alan Wang, Jianyu Su, Arsalan Heydarian, Bradford Campbell, and Peter Beling. Is my sensor sleeping, hibernating, or broken? a data-driven monitoring system for indoor energy harvesting sensors. In *Proceedings of the 7th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, pages 210–219, 2020.

[2] Cyber-Physical Systems (CPS) — new.nsf.gov. https://new.nsf.gov/funding/opportunities/cyber-physical-systems-cps. [Accessed 18-May-2023].

[3] Neil E. Klepeis, William C. Nelson, Wayne R. Ott, John P. Robinson, Andy M. Tsang, Paul Switzer, Joseph V. Behar, Stephen C. Hern, and William H. Engelmann. The national human activity pattern survey (nhaps): a resource for assessing exposure to environmental pollutants. *Journal of Exposure Science & Environmental Epidemiology 2001 11:3*, 11:231–252, 7 2001.

[4] Mohamad Awada, Burcin Becerik-Gerber, Simi Hoque, Zheng O'Neill, Giulia Pedrielli, Jin Wen, and Teresa Wu. Ten questions concerning occupant health in buildings during normal operations and extreme events including the covid-19 pandemic. *Building and Environment*, 188:107480, 1 2021.

[5] US EIA. How much energy is consumed in us residential and commercial buildings? *United States Energy Information Administration, Washington, DC, accessed Dec*, 11:2018, 2017.

[6] Andreas Wagner, William O'Brien, and Bing Dong. Exploring occupant behavior in buildings. *Wagner, A., O'Brien, W., Dong, B., Eds*, 2018.

[7] Mengmeng Wang, Lili Li, Caixia Hou, Xiaotong Guo, and Hanliang Fu. Building and health: Mapping the knowledge development of sick building syndrome. *Buildings*, 12:287, 3 2022.

[8] Amir Baniassadi, Brad Manor, Wanting Yu, Lewis Lipsitz, and Alvaro Pascual-Leone. A platform to study the effects of home environment on health and wellbeing of older adults. *Innovation in Aging*, 5:95, 12 2021.

[9] J. G. Cedenõ-Laurent, A. Williams, P. MacNaughton, X. Cao, E. Eitland, J. Spengler, and J. Allen. Building evidence for health: Green buildings, current science, and future challenges. *Annual Review of Public Health*, 39:291–308, 4 2018.

[10] Ján Drgoňa, Javier Arroyo, Iago Cupeiro Figueroa, David Blum, Krzysztof Arendt, Donghun Kim, Enric Perarnau Ollé, Juraj Oravec, Michael Wetter, Draguna L Vrabie, et al. All you need to know about model predictive control for buildings. *Annual Reviews in Control*, 50:190–232, 2020.

[11] Stephen Verderber and David Reuman. Windows, views, and health status in hospital therapeutic environments. *Journal of Architectural and Planning Research*, pages 120–133, 1987.

[12] Joseph G Allen and John D Macomber. *Healthy buildings: How indoor spaces drive performance and productivity*. Harvard University Press, 2020.

[13] Pieter Ballon, Jo Pierson, and Simon Delaere. Test and experimentation platforms for broadband innovation: Examining european practice. *Available at SSRN 1331557*, 2005.

[14] Veli-Pekka Niitamo, Seija Kulkki, Mats Eriksson, and Karl A Hribernik. State-of-the-art and good practice in the field of living labs. In *2006 IEEE international technology management conference (ICE)*, pages 1–8. IEEE, 2006.

[15] BHMSA Bergvall-Kareborn, M Hoist, and Anna Stahlbrost. Concept design with a living lab approach. In *2009 42nd Hawaii international conference on system sciences*, pages 1–10. IEEE, 2009.

[16] R Jacoby Cureau, Ilaria Pigliautile, Anna Laura Pisello, Mateus Bavaresco, Christiane Berger, Giorgia Chinazzo, Zs Deme Belafi, A Ghahramani, Arsalan Heydarian, D Kastner, et al. Bridging the gap from test rooms to field-tests for human indoor comfort studies: A critical review of the sustainability potential of living laboratories. *Energy Research & Social Science*, 92:102778, 2022.

[17] A Habibipour. Living lab research: A state-of-the-art review and steps towards a research agenda. In *OLLD18—OpenLivingLabs Days Research and Innovation Conference, Geneva, Switzerland*, 2018.

[18] Julien Nembrini and Denis Lalanne. Human-building interaction: When the machine becomes a building. In *Human-Computer Interaction-INTERACT 2017: 16th IFIP TC 13 International Conference, Mumbai, India, September 25-29, 2017, Proceedings, Part II 16*, pages 348–369. Springer, 2017.

[19] Emiel Por, M v Kooten, and Vanja Sarkovic. Nyquist–shannon sampling theorem. *Leiden University*, 1(1), 2019.

[20] Steven W Smith et al. The scientist and engineer's guide to digital signal processing, 1997.

[21] Ming Jin, Shichao Liu, Yulun Tian, Mingjian Lu, Stefano Schiavon, and Costas Spanos. Indoor environmental quality monitoring by autonomous mobile sensing. In *Proceedings of the 4th ACM International Conference on Systems for Energy-Efficient Built Environments*, pages 1–4, 2017.

[22] Eric Schulz, Maarten Speekenbrink, and Andreas Krause. A tutorial on gaussian process regression: Modelling, exploring, and exploiting functions. *Journal of Mathematical Psychology*, 85:1–16, 2018.

[23] Dinesh Kumar Gautam, Prakash Kotecha, and Senthilmurugan Subbiah. Efficient k-means clustering and greedy selection-based reduction of nodal search space for optimization of sensor placement in the water distribution networks. *Water Research*, page 118666, 2022.

[24] Liyang Yu, Neng Wang, and Xiaoqiao Meng. Real-time forest fire detection with wireless sensor networks. In *Proceedings. 2005 International Conference on Wireless Communications, Networking and Mobile Computing, 2005.*, volume 2, pages 1214–1217. Ieee, 2005.

[25] Jiakang Lu, Dagnachew Birru, and Kamin Whitehouse. Using simple light sensors to achieve smart daylight harvesting. In *Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-efficiency in Building*, pages 73–78, 2010.

[26] Emelieke RCM Huisman, Ernesto Morales, Joost van Hoof, and Helianthe SM Kort. Healing environment: A review of the impact of physical environmental factors on users. *Building and environment*, 58:70–80, 2012.

[27] Timothy W Hnat, Vijay Srinivasan, Jiakang Lu, Tamim I Sookoor, Raymond Dawson, John Stankovic, and Kamin Whitehouse. The hitchhiker's guide to successful residential sensing deployments. In *Proceedings of the 9th ACM Conference on Embedded Networked Sensor Systems*, pages 232–245, 2011.

[28] Guillermo Barrenetxea, François Ingelrest, Gunnar Schaefer, and Martin Vetterli. The hitchhiker's guide to successful wireless sensor network deployments. In *Proceedings of the 6th ACM conference on Embedded network sensor systems*, pages 43–56, 2008.

[29] Da Yan, Xiaohang Feng, Yuan Jin, and Chuang Wang. The evaluation of stochastic occupant behavior models from an application-oriented perspective: Using the lighting behavior model as a case study. *Energy and Buildings*, 176:151–162, 2018.

[30] Wei Zhou, Yan Jia, Anni Peng, Yuqing Zhang, and Peng Liu. The effect of iot new features on security and privacy: New threats, existing solutions, and challenges yet to be solved. *IEEE Internet of things Journal*, 6(2):1606–1616, 2018.

[31] Ali Ghahramani, Jovan Pantelic, Matthew Vannucci, Lorenza Pistore, Shichao Liu, Brian Gilligan, Soheila Alyasin, Edward Arens, Kevin Kampshire, and Esther Sternberg. Personal co2 bubble: Context-dependent variations and wearable sensors usability. *Journal of Building Engineering*, 22:295–304, 2019.

[32] Jovan Pantelic, Shichao Liu, Lorenza Pistore, Dusan Licina, Matthew Vannucci, Sasan Sadrizadeh, Ali Ghahramani, Brian Gilligan, Esther Sternberg, Kevin Kampschroer, et al. Personal co2 cloud: laboratory measurements of metabolic co2 inhalation zone concentration and dispersion in a typical office desk setting. *Journal of exposure science & environmental epidemiology*, 30(2):328–337, 2020.

[33] Ming Jin, Nikolaos Bekiaris-Liberis, Kevin Weekly, Costas Spanos, and Alexandre Bayen.

Sensing by proxy: Occupancy detection based on indoor co2 concentration. *UBICOMM 2015*, 14, 2015.

[34] Andrzej Szczurek, Monika Maciejewska, and Tomasz Pietrucha. Occupancy determination based on time series of co2 concentration, temperature and relative humidity. *Energy and Buildings*, 147:142–154, 2017.

[35] Young Ran Yoon, Ye Rin Lee, Sun Ho Kim, Jeong Won Kim, and Hyeun Jun Moon. A non-intrusive data-driven model for detailed occupants' activities classification in residential buildings using environmental and energy usage data. *Energy and Buildings*, 256:111699, 2022.

[36] Jana Huchtkoetter and Andreas Reinhardt. On the impact of temporal data resolution on the accuracy of non-intrusive load monitoring. In *Proceedings of the 7th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, pages 270–273, 2020.

[37] Neal Wadhwa, Hao-Yu Wu, Abe Davis, Michael Rubinstein, Eugene Shih, Gautham J Mysore, Justin G Chen, Oral Buyukozturk, John V Guttag, William T Freeman, et al. Eulerian video magnification and analysis. *Communications of the ACM*, 60(1):87–95, 2016.

[38] Khairul Rijal Wagiman, Mohd Noor Abdullah, Mohammad Yusri Hassan, and Nur Hanis Mohammad Radzi. A new optimal light sensor placement method of an indoor lighting control system for improving energy performance and visual comfort. *Journal of Building Engineering*, 30:101295, 2020.

[39] I LM. Approved method: Ies spatial daylight autonomy (sda) and annual sunlight exposure (ase). *Illuminating Engineering Society. https://www. ies. org/product/ies-spatial-daylight-autonomy-sda-and-annual-sunlight-exposure-ase*, 2013.

[40] Minjae Shin and Jeff S Haberl. Thermal zoning for building hvac design and energy simulation: A literature review. *Energy and Buildings*, 203:109429, 2019.

[41] Chao Wang and Jian Kang. Development of acoustic computer simulation for performance

spaces: A systematic review and meta-analysis. In *Building Simulation*, pages 1–17. Springer, 2022.

[42] Michael Batty. Digital twins, 2018.

[43] Lamya Abdeljalil Belhaj, Sacha Gosselin, Hélia Pouyllau, and Yann Semet. Smart-sensor placement optimization under energy objectives. In *2016 Global Information Infrastructure and Networking Symposium (GIIS)*, pages 1–7. IEEE, 2016.

[44] Nan Li, Burcin Becerik-Gerber, Bhaskar Krishnamachari, and Lucio Soibelman. A bim centered indoor localization algorithm to support building fire emergency response operations. *Automation in Construction*, 42:78–89, 2014.

[45] María Teresa Aguilar-Carrasco, Julia Díaz-Borrego, Ignacio Acosta, Miguel Ángel Campano, and Samuel Domínguez-Amarillo. Validation of lighting parametric workflow tools of ladybug and solemma using cie test cases. *Journal of Building Engineering*, page 105608, 2022.

[46] Christoph F Reinhart. Lightswitch-2002: a model for manual and automated control of electric lighting and blinds. *Solar energy*, 77(1):15–28, 2004.

[47] Neal E Young. Greedy set-cover algorithms (1974-1979, chvátal, johnson, lovász, stein). *Encyclopedia of algorithms*, pages 379–381, 2008.

[48] Ron KC Cheng. *Inside Rhinoceros 5*. Cengage Learning, 2013.

[49] Mostapha Sadeghipour Roudsari, Michelle Pak, Adrian Smith, et al. Ladybug: a parametric environmental plugin for grasshopper to help designers create an environmentally-conscious design. In *Proceedings of the 13th international IBPSA conference held in Lyon, France Aug*, pages 3128–3135, 2013.

[50] Arsalan Heydarian, Evangelos Pantazis, Alan Wang, David Gerber, and Burcin Becerik-Gerber. Towards user centered building design: Identifying end-user lighting preferences via immersive virtual environments. *Automation in Construction*, 81:56–66, 2017.

[51] Gregory J Ward. The radiance lighting simulation and rendering system. In *Proceedings of*

*the 21st annual conference on Computer graphics and interactive techniques*, pages 459–472, 1994.

[52] Anshul Agarwal, Vitobha Munigala, and Krithi Ramamritham. Observability: A principled approach to provisioning sensors in buildings. In *Proceedings of the 3rd ACM International Conference on Systems for Energy-Efficient Built Environments*, pages 197–206, 2016.

[53] Mark Sheinin, Yoav Y Schechner, and Kiriakos N Kutulakos. Computational imaging on the electric grid. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6437–6446, 2017.

[54] Biren B Kamdar, Dale M Needham, and Nancy A Collop. Sleep deprivation in critical illness: its role in physical and psychological recovery. *Journal of intensive care medicine*, 27(2):97–111, 2012.

[55] E Ely, S Gautam, R Margolin, J Francis, L May, T Speroff, B Truman, R Dittus, G Bernard, and S Inouye. The impact of delirium in the intensive care unit on hospital length of stay. *Intensive care medicine*, 27:1892–1900, 2001.

[56] Selina Dobing, Natalia Frolova, Finlay McAlister, and Jennifer Ringrose. Sleep quality and factors influencing self-reported sleep duration and quality in the general internal medicine inpatient population. *PloS one*, 11(6):e0156735, 2016.

[57] Randall S Friese. Sleep and recovery from critical illness and injury: a review of theory, current practice, and future directions. *Critical care medicine*, 36(3):697–705, 2008.

[58] Santi Kulpatcharapong, Pol Chewcharat, Kiat Ruxrungtham, Sutep Gonlachanvit, Tanisa Patcharatrakul, Busarakum Chaitusaney, Dittapol Muntham, Sirimon Reutrakul, and Naricha Chirakalwasan. Sleep Quality of Hospitalized Patients, Contributing Factors, and Prevalence of Associated Disorders. *Sleep Disorders*, 2020:1–7, 2020.

[59] Xiao Tan, Lieve van Egmond, Markku Partinen, Tanja Lange, and Christian Benedict. A narrative review of interventions for improving sleep and reducing circadian disruption in medical inpatients. *Sleep medicine*, 59:42–50, 2019.

[60] Carlos Rício and ApaFrancesca Panin. Acute Wards : a Literature Review. (February), 2020.

[61] Jun-Ki Min, Afsaneh Doryab, Jason Wiese, Shahriyar Amini, John Zimmerman, and Jason I Hong. Toss'n'turn: smartphone as sleep and sleep quality detector. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 477–486, 2014.

[62] Hagen Fritz, Kerry A Kinney, Congyu Wu, David M Schnyer, and Zoltan Nagy. Data fusion of mobile and environmental sensing devices to understand the effect of the indoor environment on measured and self-reported sleep quality. *Building and Environment*, 214:108835, 2022.

[63] Brandon M Booth, Karel Mundnich, Tiantian Feng, Amrutha Nadarajan, Tiago H Falk, Jennifer L Villatte, Emilio Ferrara, and Shrikanth Narayanan. Multimodal human and environmental sensing for longitudinal behavioral studies in naturalistic settings: Framework for sensor selection, deployment, and management. *Journal of medical Internet research*, 21(8):e12832, 2019.

[64] Andrew D Krystal and Jack D Edinger. Measuring sleep quality. *Sleep medicine*, 9:S10–S17, 2008.

[65] Teofilo L Lee-Chiong. *Sleep: a comprehensive handbook*. John Wiley & Sons, 2005.

[66] Sallie Ann Keller, Stephanie S Shipp, Aaron D Schroeder, and Gizem Korkmaz. Doing data science: A framework and case study. *Harvard Data Science Review*, 2(1), 2020.

[67] João Ramos, Joana Belo, Dário Silva, Carlos Diogo, Susana Marta Almeida, and Nuno Canha. Influence of indoor air quality on sleep quality of university students in lisbon. *Atmospheric Pollution Research*, 13(2):101301, 2022.

[68] Wan-qing He, Ai-jun Shi, Xia Shao, Lei Nie, Tian-yi Wang, and Guo-hao Li. Insights into the comprehensive characteristics of volatile organic compounds from multiple cooking emissions and aftertreatment control technologies application. *Atmospheric Environment*, 240:117646, 2020.

[69] Farman Hassan, Muhammad Hamza Mehmood, Babar Younis, Nasir Mehmood, Talha Imran, and Usama Zafar. Comparative analysis of machine learning algorithms for classification of environmental sounds and fall detection. *Science and Technology*, 4(1):163–174, 2022.

[70] Ankush Manocha. Iot-assisted irregular environmental event determination for health analysis of pregnant females. *Transactions on Emerging Telecommunications Technologies*, 33(1):e4392, 2022.

[71] TA Bedrosian and RJ Nelson. Timing of light exposure affects mood and brain circuits. *Translational psychiatry*, 7(1):e1017–e1017, 2017.

[72] C Jarboe, J Snyder, and MG Figueiro. The effectiveness of light-emitting diode lighting for providing circadian stimulus in office spaces while minimizing energy use. *Lighting Research & Technology*, 52(2):167–188, 2020.

[73] Helene Emsellem, K Knutson, D Hillygus, O Buxton, H Montgomery-Downs, M LeBourgeois, and J Spilsbury. sleep in america poll: Sleep in the modern family. *Arlington, VA: National Sleep Foundation*, 2014.

[74] Christoph Frank Reinhart. *Daylighting handbook: fundamentals, designing with the sun.* Christoph Reinhart, 04 2014.

[75] David L. DiLaura, Kevin W. Houser, Richard G. Mistrick, and Gary R. Steffy. *The lighting handbook*. 2011.

[76] Kathy Missildine, Nancy Bergstrom, Janet Meininger, Kathy Richards, and Marquis D Foreman. Sleep in hospitalized elders: a pilot study. *Geriatric nursing*, 31(4):263–271, 2010.

[77] Amy Stafford, Amy Haverland, and Elizabeth Bridges. Noise in the icu. *AJN The American Journal of Nursing*, 114(5):57–63, 2014.

[78] Kenneth I Hume, Mark Brink, Mathias Basner, et al. Effects of environmental noise on sleep. *Noise and health*, 14(61):297, 2012.

[79] Monica S Hammer, Tracy K Swinburn, and Richard L Neitzel. Environmental noise pollution in the united states: developing an effective public health response. *Environmental health perspectives*, 122(2):115–119, 2014.

[80] G Brooke Anderson, Michelle L Bell, and Roger D Peng. Methods to calculate the heat index

as an exposure metric in environmental health research. *Environmental health perspectives*, 121(10):1111–1119, 2013.

[81] Peder Wolkoff. Indoor air humidity, air quality, and health–an overview. *International journal of hygiene and environmental health*, 221(3):376–390, 2018.

[82] Cong Song, Tingting Zhao, Zhiyuan Song, and Yanfeng Liu. Effects of phased sleeping thermal environment regulation on human thermal comfort and sleep quality. *Building and Environment*, 181:107108, 2020.

[83] Md Dilshad Manzar, Mani Sethi, and M Ejaz Hussain. Humidity and sleep: a review on thermal aspect. *Biological Rhythm Research*, 43(4):439–457, 2012.

[84] Li Lan, K Tsuzuki, YF Liu, and ZW Lian. Thermal environment and sleep quality: A review. *Energy and Buildings*, 149:101–113, 2017.

[85] Zachary A Caddick, Kevin Gregory, Lucia Arsintescu, and Erin E Flynn-Evans. A review of the environmental parameters necessary for an optimal sleep environment. *Building and environment*, 132:11–20, 2018.

[86] Xiaojing Zhang, Guanzhang Luo, Jingchao Xie, and Jiaping Liu. Associations of bedroom air temperature and co2 concentration with subjective perceptions and sleep quality during transition seasons. *Indoor air*, 31(4):1004–1017, 2021.

[87] Bowen Du, Marlie C Tandoc, Michael L Mack, and Jeffrey A Siegel. Indoor co2 concentrations and cognitive function: A critical review. *Indoor Air*, 30(6):1067–1082, 2020.

[88] Peter Strøm-Tejsen, D Zukowska, Pawel Wargocki, and David Peter Wyon. The effects of bedroom air quality on sleep and next-day performance. *Indoor air*, 26(5):679–686, 2016.

[89] Lance A Wallace, Edo D Pellizzari, Tyler D Hartwell, Roy Whitmore, Charles Sparacino, and Harvey Zelon. Total exposure assessment methodology (team) study: personal exposures, indoor-outdoor relationships, and breath levels of volatile organic compounds in new jersey. *Environment International*, 12(1-4):369–387, 1986.

[90] Brandon E Boor, Michal P Spilak, Jelle Laverge, Atila Novoselac, and Ying Xu. Human exposure to indoor air pollutants in sleep microenvironments: A literature review. *Building and Environment*, 125:528–555, 2017.

[91] Hagen Fritz. Data fusion of mobile and environmental monitoring devices to understand the effects of the indoor environment on sleep quality. *Whole Communities-Whole Health-Published Research*, 2021.

[92] Xiaochen Tang, Pawel K Misztal, William W Nazaroff, and Allen H Goldstein. Volatile organic compound emissions from humans indoors. *Environmental science & technology*, 50(23):12686–12694, 2016.

[93] Chien-Cheng Jung, Pei-Chih Wu, Chao-Heng Tseng, and Huey-Jen Su. Indoor air quality varies with ventilation types and working areas in hospitals. *Building and Environment*, 85:190–195, 2015.

[94] Energy Corps. California air resources board. 2019.

[95] Lei Li, Weituo Zhang, Li Xie, Sinong Jia, Tienan Feng, Herbert Yu, Jie Huang, and Biyun Qian. Effects of atmospheric particulate matter pollution on sleep disorders and sleep duration: a cross-sectional study in the uk biobank. *Sleep medicine*, 74:152–164, 2020.

[96] Environmental Protection Agency. REVISED AIR QUALITY STANDARDS FOR PARTICLE POLLUTION AND UPDATES TO THE AIR QUALITY INDEX (AQI) . https://www.epa.gov/sites/default/files/2016-04/documents/2012_aqi_factsheet.pdf, 2012. [Online; accessed 12-June-2023].

[97] Liqiong Chang, Jiaqi Lu, Ju Wang, Xiaojiang Chen, Dingyi Fang, Zhanyong Tang, Petteri Nurmi, and Zheng Wang. Sleepguard: Capturing rich sleep information using smartwatch sensing data. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 2(3), sep 2018.

[98] Florian Wahl and Oliver Amft. Data and expert models for sleep timing and chronotype estimation from smartphone context data and simulations. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(3):1–28, 2018.

[99] Ryan M. May, Kevin H. Goebbert, Jonathan E. Thielen, John R. Leeman, M. Drew Camron, Zachary Bruick, Eric C. Bruning, Russell P. Manser, Sean C. Arms, and Patrick T. Marsh. Metpy: A meteorological python library for data analysis and visualization. *Bulletin of the American Meteorological Society*, 103(10):E2273 – E2284, 2022.

[100] Rakesh Malladi, Giridhar P Kalamangalam, and Behnaam Aazhang. Online bayesian change point detection algorithms for segmentation of epileptic activity. In *2013 Asilomar Conference on Signals, Systems and Computers*, pages 1833–1837. IEEE, 2013.

[101] Arash Tavakoli, Shashwat Kumar, Xiang Guo, Vahid Balali, Mehdi Boukhechba, and Arsalan Heydarian. Harmony: A human-centered multimodal driving study in the wild. *IEEE Access*, 9:23956–23978, 2021.

[102] Shashwat Kumar, Debajyoti Datta, Guimin Dong, Lihua Cai, Laura Barnes, and Mehdi Boukhechba. Leveraging mobile sensing and bayesian change point analysis to monitor community-scale behavioral interventions: a case study on covid-19, 12 2021.

[103] Charles Truong, Laurent Oudre, and Nicolas Vayatis. Selective review of offline change point detection methods. *Signal Processing*, 167:107299, 2020.

[104] Ariel B Neikrug and Sonia Ancoli-Israel. Sleep disturbances in nursing homes. *The journal of nutrition, health & aging*, 14(3):207–211, 2010.

[105] David Clark-Carter. z scores. *Encyclopedia of statistics in behavioral science*, 2005.

[106] B Jason. Smote for imbalanced classification with python, 2021.

[107] Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16:321–357, 2002.

[108] Guillaume Lemaître, Fernando Nogueira, and Christos K Aridas. Imbalanced-learn: A python toolbox to tackle the curse of imbalanced datasets in machine learning. *The Journal of Machine Learning Research*, 18(1):559–563, 2017.

[109] J. Bruin. newtest: command to compute new test @ONLINE, February 2011.

[110] Takeshi Imura, Yuji Iwamoto, Tetsuji Inagawa, Naoki Imada, Ryo Tanaka, Haruki Toda, Yu Inoue, Hayato Araki, and Osamu Araki. Decision tree algorithm identifies stroke patients likely discharge home after rehabilitation using functional and environmental predictors. *Journal of Stroke and Cerebrovascular Diseases*, 30(4):105636, 2021.

[111] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

[112] Jehad Ali, Rehanullah Khan, Nasir Ahmad, and Imran Maqsood. Random forests and decision trees. *International Journal of Computer Science Issues (IJCSI)*, 9(5):272, 2012.

[113] Kyle Konis. A novel circadian daylight metric for building design and evaluation. *Building and Environment*, 113:22–38, 2017.

[114] Shari L Forbes, Katelynn A Perrault, Pierre-Hugues Stefanuto, Katie D Nizio, and Jean-François Focant. Comparison of the decomposition voc profile during winter and summer in a moist, mid-latitude (cfb) climate. *PLoS One*, 9(11):e113681, 2014.

[115] Clete A Kushida, Arthur Chang, Chirag Gadkary, Christian Guilleminault, Oscar Carrillo, and William C Dement. Comparison of actigraphic, polysomnographic, and subjective assessment of sleep parameters in sleep-disordered patients. *Sleep medicine*, 2(5):389–396, 2001.

[116] Girardin Jean-Louis, Hans von Gizycki, Ferdinand Zizi, Jeffrey Fookson, Arthur Spielman, Joao Nunes, Robert Fullilove, and Harvey Taub. Determination of sleep and wakefulness with the actigraph data analysis software (adas). *Sleep*, 19(9):739–743, 1996.

[117] George Havenith, Ingvar Holmér, and Ken Parsons. Personal factors in thermal comfort assessment: clothing properties and metabolic heat production. *Energy and buildings*, 34(6):581–591, 2002.

[118] Leah B. Peirce, Nicola M. Orlov, Amarachi I. Erondu, Samantha L. Anderson, Michael Chamberlain, David Gozal, and Vineet M. Arora. Caregiver and staff perceptions of disruptions to pediatric inpatient sleep. *Journal of Clinical Sleep Medicine*, 2018.

[119] Carrie A Redlich, Judy Sparer, and Mark R Cullen. Sick-building syndrome. *The Lancet*, 349(9057):1013–1016, 1997.

[120] Gail Brager, Gwelen Paliaga, and Richard De Dear. Operable windows, personal control and occupant comfort. 2004.

[121] MG Figueiro, M Kalsher, BC Steverson, J Heerwagen, K Kampschroer, and MS Rea. Circadian-effective light and its impact on alertness in office workers. *Lighting Research & Technology*, 51(2):171–183, 2019.

[122] Piers MacNaughton, Usha Satish, Jose Guillermo Cedeno Laurent, Skye Flanigan, Jose Vallarino, Brent Coull, John D Spengler, and Joseph G Allen. The impact of working in a green certified building on cognitive function and health. *Building and Environment*, 114:178–186, 2017.

[123] Zhun Yu, Benjamin CM Fung, Fariborz Haghighat, Hiroshi Yoshino, and Edward Morofsky. A systematic procedure to study the influence of occupant behavior on building energy consumption. *Energy and buildings*, 43(6):1409–1417, 2011.

[124] Sujesha Sudevalayam and Purushottam Kulkarni. Energy harvesting sensor nodes: Survey and implications. *IEEE Communications Surveys & Tutorials*, 13(3):443–461, 2010.

[125] Lohit Yerva, Brad Campbell, Apoorva Bansal, Thomas Schmid, and Prabal Dutta. Grafting energy-harvesting leaves onto the sensornet tree. In *Proceedings of the 11th international conference on Information Processing in Sensor Networks*, pages 197–208, 2012.

[126] Krishna Veni Selvan and Mohamed Sultan Mohamed Ali. Micro-scale energy harvesting devices: Review of methodological performances in the last decade. *Renewable and Sustainable Energy Reviews*, 54:1035–1047, 2016.

[127] Yi-Chang Li and Seung Ho Hong. Bacnet–enocean smart grid gateway and its application to demand response in buildings. *Energy and buildings*, 78:183–191, 2014.

[128] Francesco Fraternali, Bharathan Balaji, Yuvraj Agarwal, Luca Benini, and Rajesh Gupta. Pible: battery-free mote for perpetual indoor ble applications. In *Proceedings of the 5th Conference on Systems for Built Environments*, pages 168–171, 2018.

[129] Qingfeng Lin, Hongtao Huang, Yan Jing, Huiying Fu, Paichun Chang, Dongdong Li, Yan Yao, and Zhiyong Fan. Flexible photovoltaic technologies. *Journal of Materials Chemistry C*, 2(7):1233–1247, 2014.

[130] Bradford Campbell and Prabal Dutta. An energy-harvesting sensor architecture and toolkit for building monitoring and event detection. In *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings*, pages 100–109, 2014.

[131] Jean-Claude Laprie. Dependable computing: Concepts, limits, challenges. In *Special issue of the 25th international symposium on fault-tolerant computing*, pages 42–54, 1995.

[132] Soila P Kavulya, Kaustubh Joshi, Felicita Di Giandomenico, and Priya Narasimhan. Failure diagnosis of complex systems. In *Resilience assessment and evaluation of computing systems*, pages 239–261. Springer, 2012.

[133] William F Young, Christopher L Holloway, Galen Koepke, Dennis Camell, Yann Becquet, and Kate A Remley. Radio-wave propagation into large building structures—part 1: Cw signal attenuation and variability. *IEEE Transactions on Antennas and Propagation*, 58(4):1279–1289, 2010.

[134] Kate A Remley, Galen Koepke, Christopher L Holloway, Chriss A Grosvenor, Dennis Camell, John Ladbury, Robert T Johnk, and William F Young. Radio-wave propagation into large building structures—part 2: Characterization of multipath. *IEEE transactions on antennas and propagation*, 58(4):1290–1301, 2010.

[135] Jonas Medbo and J-E Berg. Simple and accurate path loss modeling at 5 ghz in indoor environments with corridors. In *Vehicular Technology Conference Fall 2000. IEEE VTS Fall VTC2000. 52nd Vehicular Technology Conference (Cat. No. 00CH37152)*, volume 1, pages 30–36. IEEE, 2000.

[136] Daniel B Faria et al. Modeling signal attenuation in ieee 802.11 wireless lans-vol. 1. *Computer Science Department, Stanford University*, 1, 2005.

[137] Johannes Tiusanen. Wireless soil scout prototype radio signal reception compared to the attenuation model. *Precision Agriculture*, 10(5):372–381, 2009.

[138] Bradford Campbell, Branden Ghena, Ye-Sheng Kuo, and Prabal Dutta. Swarm gateway: Demo abstract. In *Proceedings of the 3rd ACM International Conference on Systems for Energy-Efficient Built Environments*, pages 217–218, 2016.

[139] André GM Lima and Luiz F Menezes. Motley-keenan model adjusted to the thickness of the wall. In *SBMO/IEEE MTT-S International Conference on Microwave and Optoelectronics, 2005.*, pages 180–182. IEEE, 2005.

[140] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.

[141] Gregory Koch, Richard Zemel, and Ruslan Salakhutdinov. Siamese neural networks for one-shot image recognition. In *ICML deep learning workshop*, volume 2. Lille, 2015.

[142] Tu Le, Alan Wang, Yaxing Yao, Yuanyuan Feng, Arsalan Heydarian, Norman Sadeh, and Yuan Tian. Exploring smart commercial building occupants' perceptions and notification preferences of internet of things data collection in the united states. *arXiv preprint arXiv:2303.04955*, 2023.

[143] Mary Angelique G Demetillo, Colin Harkins, Brian C McDonald, Philip S Chodrow, Kang Sun, and Sally E Pusede. Space-based observational constraints on no2 air pollution inequality from diesel traffic in major us cities. *Geophysical Research Letters*, 48(17):e2021GL094333, 2021.

[144] Hiroshi Ishii and Brygg Ullmer. Tangible bits: towards seamless interfaces between people, bits and atoms. In *Proceedings of the ACM SIGCHI Conference on Human factors in computing systems*, pages 234–241, 1997.

# Chapter A

# Appendix

**@Algorithm 1** Greedy Set Cover Approximation (GSCA)

**Input**:

- Universe of elements $U_A$
- Collection of sub universes $\mathbb{U} = \{U_1, U_2, ...U_n\}$
- Value function $V(S)$

**Output**: Array $P$, containing position(s) $p$ of sensors as referenced by index in $\mathbb{S}$

$I = \{\}$ ▷ Viewable Universe

$P = \{\}$ ▷ The final indices of the selected points stored here

$\mathbb{S}' = \mathbb{S}$ ▷ Retain the original collection

**while** $I \neq U$ **do**

  $i = (\frac{V(S_0)}{|S_0 - I|}, \frac{V(S_1)}{|S_1 - I|}, \cdots \frac{V(S_{|\mathbb{S}|})}{|S_{|\mathbb{S}|} - I|})$ ▷ Find index of highest value in the remaining set of $\mathbb{S}$

  **if** $S_i \nsubseteq I$ **then** ▷ Only add the location if it is not a subset of the existing viewable universe $I$

    $I = I \bigcup S_i$

  **end if**

  $\mathbb{S}' = \mathbb{S}' \setminus S_i$ ▷ Remove selected subset from consideration

  Append index of $S_i$ in relation to $\mathbb{S}$ to $P$

  **if** $\mathbb{S}' = \emptyset$ **then** ▷ If you reach the end without completing the universe

    **return** No set cover exists

```
    end if
end while
return P
```

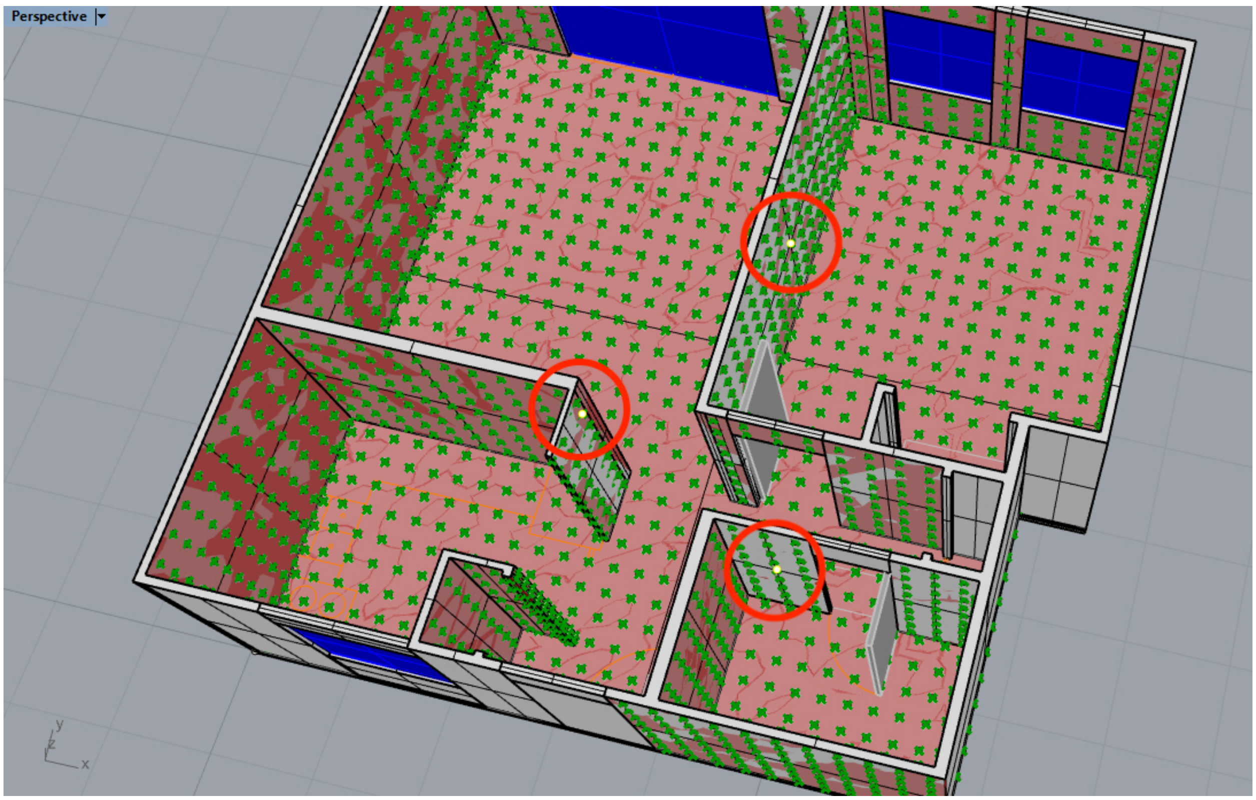


Figure A.1: Example manual selection of 3 sensor positions, with the complex scenario (see Section 1.2.3) in mind, achieving a 95% (540/567) total inferable states.

FallRiskTotal, NUDESC-Disorientation, NUDESC-InappropriateBehavior, NUDESC-InappropriateCommunication, NUDESC-IllusionsandHallucinations, NUDESC-PsychomotorRetardation, NuDESCScore, PainScale, PainRating, BERTCall, BERTReason, METCall, Fall, DidFallCauseanInjury, O2DeliveryDevice, O2FlowRateLmin, SpO2, Nonnarcotic-PainMedName, NonnarcoticPainMedRoute, NonnarcoticAction, OpioidMedName, OpioidMedRoute, OpioidAction, BenzodiazepineReceptorAgonistsName, BenzodiazepineReceptorAgonists, SedativeHypnoticMedName, SedativeHypnoticMedRoute, OtherMedsthatCouldAffectSleepandorDeliriumName, OtherMedsthatCouldAffectSleepandorDeliriumRoute, OtherMedsActionType, OtherInpatientSleepMedName, OtherInpatientSleepMedRoute, InpatientMelatoninUse, MedNameDoseCodes, InpatientMedsforDelirium, InpatientMedsforDeliriumNameincludesdose, InpatientMedsforDeliriumRoute, NighttimeEvent19000700VitalSignsTaken, NighttimeEventEKG, NighttimeEventBloodGlucoseMeasured, NighttimeLabEvent, NighttimeAdministrationofMedsEXCLUDINGSPECIALMEDS, MedName, MedicationAction, NighttimeADTEvent, NighttimeConsult, NighttimeProcedure, NighttimeLinesDrainsAirwayWound, OtherNighttimeEvent, NighttimeEventCarePlanInterventionEvent, IOEventatNight, NeuroWDLB, IfnotWDL, LevelofConsciousness, Orientated, CommentRelatedtoSpecificEvent, CommentRelatedtoSubjectnotrelatedtoaspecificdatetime.

Figure A.2: The columns used from EPIC chosen by our nursing team

METCall,Fall,O2DeliveryDevice

Figure A.3: Individual Events Columns

NonnarcoticPainMedName,OpioidMedName,BenzodiazepineReceptorAgonistsName, SedativeHypnoticMedName,OtherMedsthatCouldAffectSleepandorDeliriumName, OtherInpatientSleepMedName,InpatientMelatoninUse, InpatientMedsforDelirium

Figure A.4: Medication Administration Events Columns

NighttimeEvent19000700VitalSignsTaken,NighttimeEventEKG, NighttimeEventBloodGlucoseMeasured, NighttimeLabEvent,NighttimeAdministrationofMedsEXCLUDINGSPECIALMEDS, NighttimeADTEvent,NighttimeConsult,NighttimeLinesDrainsAirwayWound, OtherNighttimeEvent,NighttimeEventCarePlanInterventionEvent,IOEventatNight

Figure A.5: Other Nighttime Events Columns

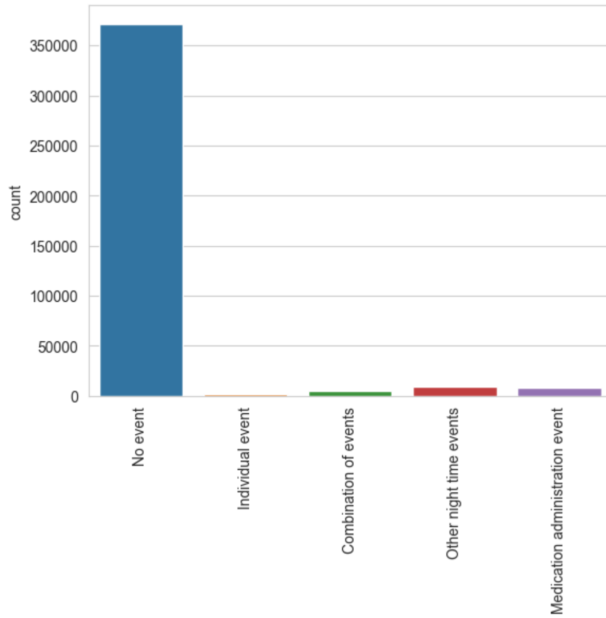


Figure A.6: Imbalanced class labels before SMOTE, where the no events class represents the majority class.

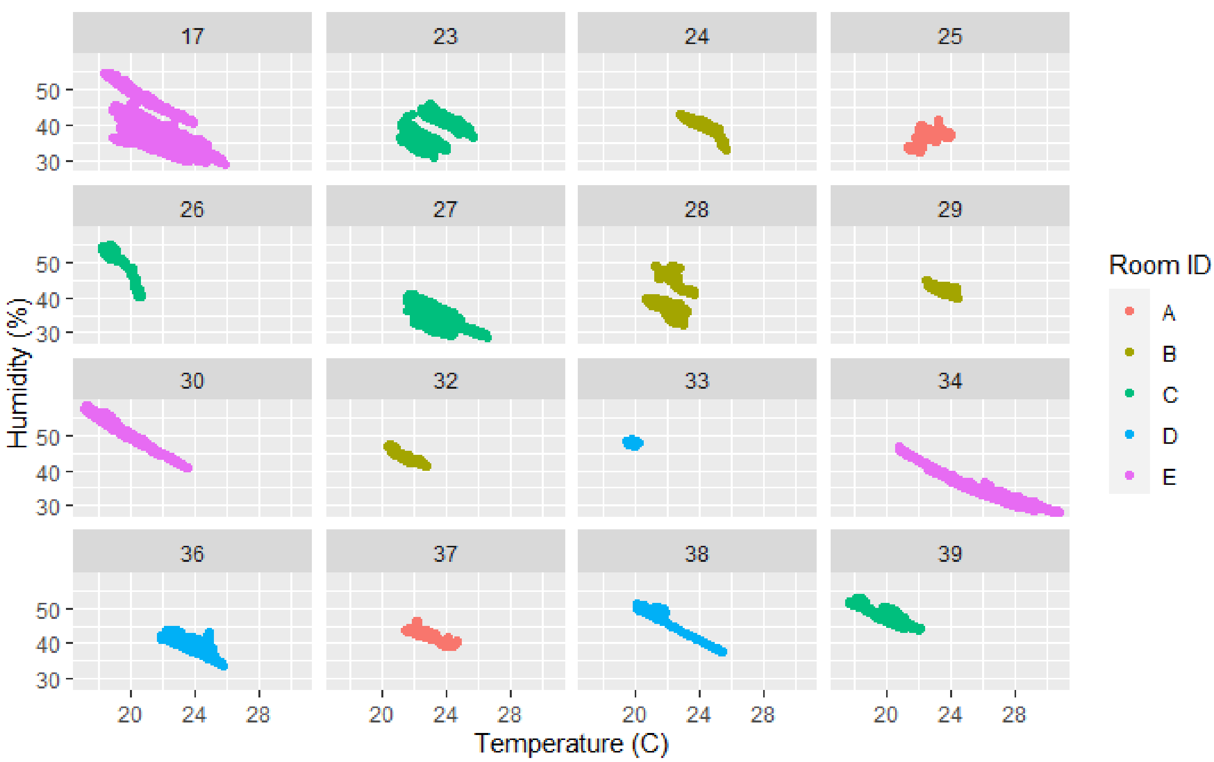


Figure A.7: A strong negative relationship between temperature and humidity across patients and room, leading to humidity's removal in the final model. The number on top of each subgraph refers to the patient id.

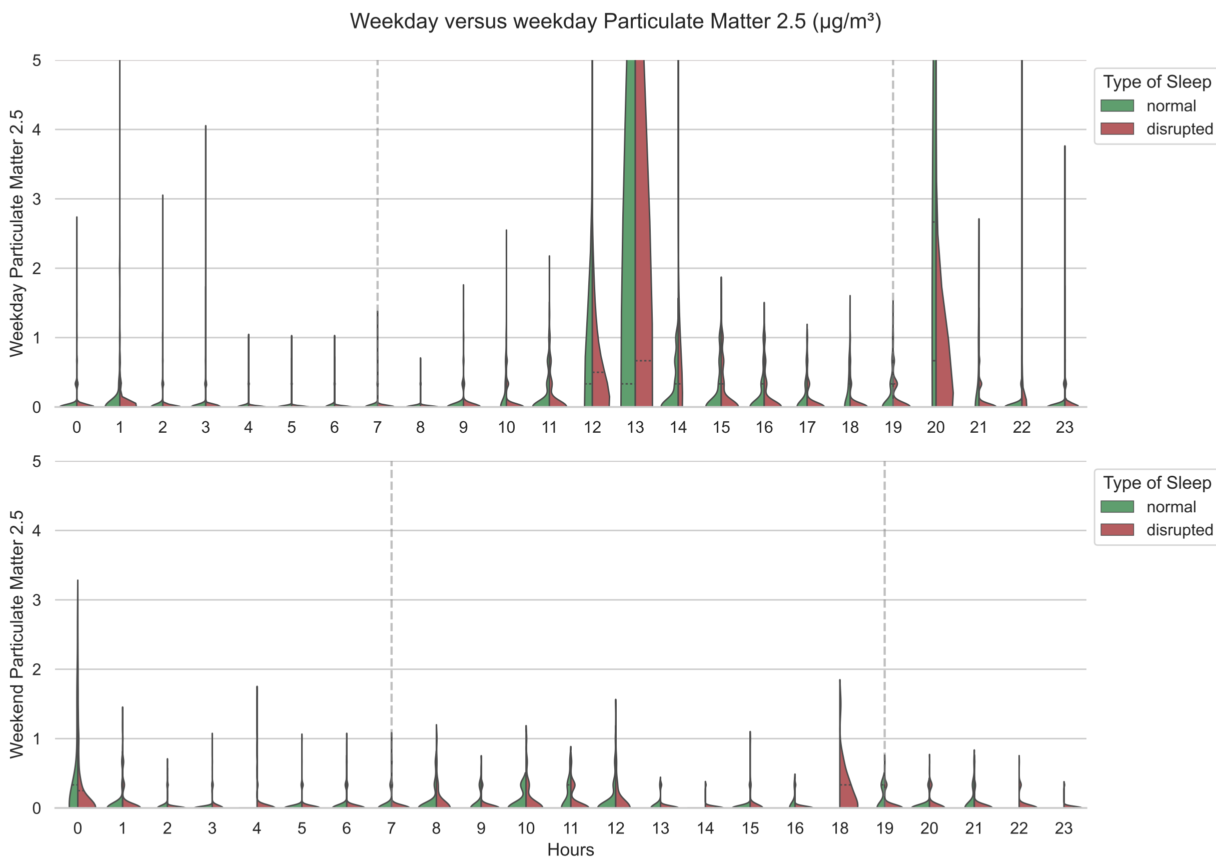


Figure A.8: Weekend Versus Weekday $PM_{2.5}$ and Sleep Disruption, showing peaks around 12 pm, 6 pm, and 12 am, indicating a relationship between perturbation and food consumption times.