

Is my sensor sleeping, hibernating, or broken? A data-driven monitoring system for indoor energy harvesting sensors

Alan Wang
ahw9f@virginia.edu
University of Virginia
Charlottesville, Virginia

Jianyu Su
js9wv@virginia.edu
University of Virginia
Charlottesville, Virginia

Arsalan Heydarian
ah6rx@virginia.edu
University of Virginia
Charlottesville, Virginia

Bradford Campbell
bradjc@virginia.edu
University of Virginia
Charlottesville, Virginia

Peter Beling
pb3a@virginia.edu
University of Virginia
Charlottesville, Virginia

ABSTRACT

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.

CCS CONCEPTS

• Computer systems organization → Sensor networks; • Computing methodologies → Classification and regression trees.

KEYWORDS

cyber-physical systems, energy harvesting sensors, machine-learning

ACM Reference Format:

Alan Wang, Jianyu Su, Arsalan Heydarian, Bradford Campbell, and Peter Beling. 2020. Is my sensor sleeping, hibernating, or broken? A data-driven monitoring system for indoor energy harvesting sensors. In *The 7th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation (BuildSys '20)*, November 18–20, 2020, Virtual Event, Japan. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3408308.3427625>

1 INTRODUCTION

As buildings strive to be not just green but also healthy [1], 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 the needs of users [2–6]. Since the needs of occupants are complex and multifaceted, there is 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 insuring 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 [7]. Not only can harvesting out-perform batteries when devices are smaller than a sugar cube [8], 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.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

BuildSys '20, November 18–20, 2020, Virtual Event, Japan

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-8061-4/20/11...\$15.00

<https://doi.org/10.1145/3408308.3427625>

As energy-harvesting devices become more accessible [9], and as such more used in buildings [10, 11], the set of challenges related to reliability and robustness become more pressing. While a small, “stick-on”, and photovoltaic-powered sensor [12] 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 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 different heartbeat intervals and event-trigger conditions.
- The average prediction accuracy for healthy signal zones is greater than 80% across all investigated device types.

This paper is thus organized as follows: Section 2 describes different modes of fault detection and the specifics of the radio signal attenuation literature. The methodology section introduces the lab space where we conducted our study, the modifications we made to previously applied methods to monitor data transmissions, and how and why we applied feature representation learning on our data. The feature preparation section describes specifically how we extracted the gateway status and device-to-gateway path characteristics from our database to use in our machine learning algorithm. The experiment and data preparation section details our six month study, describing the columns of our published dataset. The results section details the validation of our method, demonstrating an above 80% accuracy in prediction future signal status across all device types. The conclusion sections summarize our findings, expanding on our assumptions and lessons learned during our sensor deployment. Finally, in limitations and future works we provide a description of parameters we did not account for, as well as our future directions.

2 RELATED WORK

Because of its lower cost during upfront installation and better scalability of maintenance compared to battery and mains-powered devices [13], an energy-harvesting (EH) sensor based architecture was proposed as an ideal infrastructure for building monitoring and event detection [14]. *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 [15]. Kavulya et al. extends on this nomenclature and describes different diagnosis techniques, limitations, and examples [16]. 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 [17–20]. 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 [21]. 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 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 3.1 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 3.2 defines the different divisions of time we use, and details when and what gateways were installed on the timeline. In section 3.3, we build upon the defined time definitions to derive a signal health score calculation that resists bias caused by event-triggered sensing. Section 3.4 provides background information on how signal attenuation is modeled for an indoor environment, which is demonstrated to result in measurable data packet loss [21], which we can now detect using the derived signal health score. Lastly, section 3.5 describe the details of how we use feature-representation learning to predict future healthy signal locations for the different aforementioned spatial and temporal arrangements.

3.1 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 1 documents the gateway location, device location, and relevant device specifications. The gray circles drawn around the gateways mark a 25 meter radius.

3.2 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 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 1. For example, the time frame used for light level sensor health score calculations is a constant equal to 30 minutes.

Table 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.

Table 1: Time ranges and gateways configuration

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]	$G_4 = \{GW1, GW2, GW3, GW4\}$

3.3 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




Figure 1: Projected Device Plan and Descriptions

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:

$$H = \frac{1}{N} \sum_t \min(1, r_t) \quad (1)$$

The silent rate of the time period for the device as described in [21] 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.

3.4 Radio Signal Attenuation

The Keegnan-Motley model of logarithmic signal loss L , as described by [19], is:

$$L(d) = L_{FS}(d) + n_w L_w + n_f L_f \quad (2)$$

With $L_{FS}(d)$ the theoretical loss in free space for an isotropically 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 thickness of




Figure 2: Fault Identification Flowchart

the wall [23]. While signal attenuation is then generally calculated using a constant attenuation per unit path length α , in our feature preparation section (section 4.2) we detail our process in tracing the discrete partitions and free space for each device-to-gateway path.

3.5 Feature Representation Learning

Feature representation learning is commonly used for applications such as natural language processing and one-shot image recognition [24]. 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 [25] 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:

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

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

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

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:

$$P = \sigma\left(\sum_{j=1}^{N_L} \alpha_j |h_{1,L-1}^j - h_{2,L-1}^j|\right) \quad (3)$$

Here, α denotes weights in the final layer, and σ 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.

4 FEATURE PREPARATION

4.1 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 2 shows a flowchart for our fault identification process.

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 shows our identification of a gateway powered by an occupancy-controlled outlet through the health score of all the devices.




Figure 3: Occupancy-controlled-outlet plugged gateway 4 related transmission count per device type over time

4.2 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.




Figure 4: Relationship between signal health score and distance to gateway for light level sensors during T_0

Plotting the distance to health score relationships for the light level sensors, shown in Figure 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.

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 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 are able to represent them using a siamese network. In the next section, we detail how we adopted this information




Figure 5: Example ray-trace with its corresponding waveform from A to B

into our machine learning model that accounts for all the different device to gateway raytraces.

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 4.1) 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 2 summarizes the columns of the data frame with an example row.

The *wall_arr* column (Shown in Table 2) describe characteristics of the building between the sensor and the gateway, where element

Table 2: Data set column descriptions

Name	Example value	Description
date	2020-01-31	YYYY-MM-DD time description for the device
device_type	LightLevel	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

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.

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.

6 RESULTS

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

6.2 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 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 and Table 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 7: Average accuracy over iterations per type for T_1 over 100 runs

Table 3: 100 Run Average Decision Tree Classifier Results

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 4: 100 Run Average 1-Layer Linear 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

6.3 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 8 a comparison between one of the highest performing sampling and one the lowest performing sample 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 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 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.

7 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 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 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




Figure 8: > 90% accuracy sampling during T_3 (left) versus < 60% accuracy sampling during T_0 (right), where the red circles represents 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.




Figure 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.

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 [26]. Some adhesives we installed degraded over the period of months, and 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. [13]). 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.

8 LIMITATIONS AND FUTURE WORK

8.1 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).

8.2 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.

8.3 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 to 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.

9 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.

ACKNOWLEDGMENTS

This work was supported in part by the the National Science Foundation under grant #1823325. Special thanks to Jacob Rantas, the undergraduate researcher who assisted with the installation the sensors, and Amy Eichenberger, our contact at the UVa FMS for support with the model and drawings of the testbed. Our gratitude also goes out to all the residents of the UVa Link Lab, for supporting the Living Link Lab research infrastructure.

REFERENCES

- [1] Joseph G Allen and John D Macomber. *Healthy buildings: How indoor spaces drive performance and productivity*. Harvard University Press, 2020.
- [2] Carrie A Redlich, Judy Sparer, and Mark R Cullen. Sick-building syndrome. *The Lancet*, 349(9057):1013–1016, 1997.
- [3] Gail Brager, Gwelen Paliaga, and Richard De Dear. Operable windows, personal control and occupant comfort. 2004.
- [4] MG Figueiro, M Kalscher, 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.
- [5] 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.
- [6] 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.
- [7] Sujesha Sudevalayam and Purushottam Kulkarni. Energy harvesting sensor nodes: Survey and implications. *IEEE Communications Surveys & Tutorials*, 13(3):443–461, 2010.
- [8] Lohit Yerva, Brad Campbell, Apoorva Bansal, Thomas Schmid, and Prabal Dutta. Grafting energy-harvesting leaves onto the sensor tree. In *Proceedings of the 11th international conference on Information Processing in Sensor Networks*, pages 197–208, 2012.
- [9] 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.
- [10] 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.
- [11] 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.
- [12] 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.
- [13] 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.
- [14] 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.
- [15] Jean-Claude Laprie. Dependable computing: Concepts, limits, challenges. In *Special issue of the 25th international symposium on fault-tolerant computing*, pages 42–54, 1995.
- [16] 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.
- [17] William F Young, Christopher L Holloway, Galen Koepke, Dennis Camell, Yann Bequet, 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.
- [18] 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.
- [19] 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.
- [20] Daniel B Faria et al. Modeling signal attenuation in ieee 802.11 wireless lans-vol. 1. *Computer Science Department, Stanford University*, 1, 2005.
- [21] Johannes Tiusanen. Wireless soil scout prototype radio signal reception compared to the attenuation model. *Precision Agriculture*, 10(5):372–381, 2009.
- [22] Bradford Campbell, Brandon 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.
- [23] 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.
- [24] 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.
- [25] Gregory Koch, Richard Zemel, and Ruslan Salakhutdinov. Siamese neural networks for one-shot image recognition. In *ICML deep learning workshop*, volume 2. Lille, 2015.
- [26] Andreas Wagner, William O'Brien, and Bing Dong. Exploring occupant behavior in buildings. *Wagner, A., O'Brien, W., Dong, B., Eds*, 2018.

ONLINE RESOURCES

The data set that includes all relevant information to replicate our analysis has been archived on [livinglinklab.github.io](https://github.com/livinglinklab)