# Physics-Informed Detection of Faulty Sensors in Smart Buildings

Hunter Lutz
Department of Civil and Environmental Engineering
University of Virginia
Charlottesville, VA
mty9zs@virginia.edu

Jessie Nothstine
Department of Urban and Environmental Planning
University of Virginia
Charlottesville, VA
sfv8vn@virginia.edu

Shawn Thomas
Department of Computer Science University of Virginia Charlottesville, VA
yzy2bt@virginia.edu

## Abstract

Sensor calibration drift is a major barrier to reliable operation of smart buildings. This study presents a physics-informed framework for automated drift detection using only existing sensor data. Applied to 76 environmental sensors (temperature, humidity, CO2, VOC) across 19 devices in Olsson and Thornton Hall at the University of Virginia over 395 days, the method combines (1) machine learning linear regression on daily averages with statistical significance testing ($p < 0.05$), (2) physical plausibility checks, and (3) multi-sensor consistency validation. Results show 25 sensors (32.9%) exhibit significant drift, with 13 exceeding critical thresholds (e.g., humidity drift >10%/yr, temperature >1°C/yr). Physical inspection in December 2025 confirmed that all 13 critically drifting sensors were exposed to bias-induced micro-environments (doors, vents, moisture sources), whereas stable sensors were not. Humidity sensors displayed the most severe degradation (up to 56%/yr). A cost-benefit analysis using industry-standard maintenance figures demonstrates that predictive calibration triggered by the proposed method can reduce annual maintenance costs by 75% ($11,350 savings for this network). The approach requires no additional hardware, is fully automated, and is immediately deployable in existing building management systems.

## 1.0 Introduction

Smart buildings rely on extensive networks of sensors in order to monitor and control key indoor environmental conditions. These span environmental and occupancy needs, including temperature, humidity, carbon dioxide concentration (CO2), and volatile organic compounds (VOC). These smart building sensing systems allow for energy-efficient HVAC operation, improved occupant comfort, and advanced analytics for building performance. Nonetheless, as buildings become more and more complicated and intelligent, maintaining the accuracy and reliability of said sensors presents a growing challenge.

As sensors age, they can drift, malfunction, and even lose calibration due to factors such as dust accumulation, hardware degradation and environmental stress. Unfortunately, just one faulty sensor is enough to cause inefficient system operations, incorrect control decisions, or misleading performance analytics. Despite these risks, identifying which sensors are unreliable remains a largely manual and reactive process.

This project explores whether broken or uncalibrated sensors can be detected automatically by leveraging physical relationships that constrain environmental variables in a building. Specifically, this project aims to investigate whether “inexplicable” reading deltas between nearby sensors monitoring the same space can reveal sensor faults. Through comparing measured differences to those expected from building physics such as thermal diffusion and airflow, we aim to develop a data-driven method to flag potentially unreliable sensors.

## 2.0 Problem Statement

Sensor drift is a common issue in the operation of cyber-physical systems (CPS) in smart buildings which refers to gradual and subtle changes in the sensor which happen over time, causing a discrepancy between the actual state of the building that is being measured and the output of the sensor. Several factors lead to the degradation of sensors over time which could include environmental conditions (Extreme temperatures/ humidity/pressure etc.), hardware faults (Manufacturing defects), wear-and-tear (physical stress), and general lack of maintenance. This issue has been previously described as an intractable obstacle when it comes to proper functioning of smart buildings and their related CPS (such as HVAC systems and Smart Lighting systems), primarily due to how the significantly degraded measurement accuracy from the sensors introduces biased inputs in the decision making processes, resulting in inefficient and poor quality of operations. (Chen et al., 2019)

Many attempts have been made to address this issue, including a variety of spatial correlation models to track any sensor faults. Spatial kriging is one such example, where the expected value at a sensor is estimated by weighting readings from nearby sensors based on how strongly they are spatially correlated. (Kumar et al. 2013) This provides a prediction of what the sensor should be reading given its environment. A Kalman filter then uses this prediction over time to detect and correct gradual drift or bias in the sensor’s measurements, enabling automatic, ongoing calibration. Spatial correlation MAP calibration is another method which uses correlations between different neighboring sensors to estimate drifts using priors and alternating optimization (Chen et al. 2019). However, these methods are accompanied with some limitations; kriging relies on smooth, stationary spatial fields and sufficient healthy neighboring sensors, and its accuracy depends on a very computationally expensive model for large networks. Kalman filtering assumes gradual Gaussian drift dynamics and may fail to capture sudden or nonlinear faults (Kumar et al. 2013). MAP-based calibration methods are described to be non-convex, sensitive to initialization, and require drift free reference data (Chen et al. 2019).

For the scope of this project, we will focus on a pairwise consistency approach building on the MAP-based calibration methods which compares readings between nearby sensors within the same zone to detect readings which are deviated from expected physical relationships. This approach allows us to track any broken or uncalibrated sensors directly from their physical context while reducing computational overhead. We aim to improve this approach by using robust rolling baselines(reference values being continuously updated) and localizing the comparisons between sensors in order to detect drift and bias without requiring sensor reference data which is free of any drift.

## 3.0 Motivation

Oftentimes, reactive countermeasures towards sensor malfunctions are only enacted on an annual basis; this can result in accumulation of inaccurate readings which could misguide actions taken by building systems, negatively impacting the indoor environment in ways that may be uncomfortable to visitors or wasteful for the owner.

Through developing a data-driven method to flag unreliable sensors, our project aims to improve the methods by which sensor data is monitored and made accurate and calibrated, a result which would not only serve to improve the accuracy of building system responses for better real-world performance, but also improve the accuracy of similar studies leveraging sensor data from building systems.

## 4.0 Methodology

### 4.1 Data Collection and Preprocessing

To best ascertain the overall direction of this project, initial environmental data from the sensors was needed to establish an understanding of their current function. To do this, we collected data from IoT sensors across Olsson and Thornton Halls respectively. These sensors record temperature (°C) and relative humidity (%) continuously; our analysis focused on long-term data across multiple months to capture any gradual drift that could be detected between sensors. The first step was to gather more sensor data from multiple devices and measurement types over approximately 395 days. The sensor types sampled include temperature, humidity, CO2, and VOC data from the survey area. Device tags and identifiers were used to compile datasets for processing.

### 4.2 Reference Signal Construction

A median reference signal was computed across all active sensors for each timestamp. This median value serves as a baseline representing the "true" environmental condition in the space and would be used to highlight any deviations displayed by individual sensors. For each of these cases, deviation from the median reference was calculated to measure bias and consistency.

### 4.3 Machine Learning-Based Drift Detection

We formulated drift detection as a supervised time-series regression problem. For each sensor i, the target variable was the daily deviation from the reference:

$\Delta_i(t) = \text{sensor}_i(t) - \text{median}(t)$

Linear regression was then performed:

$\Delta_i(t) = \beta \cdot t + \varepsilon$

Annual drift is reported as $\beta \times 365$. A sensor is flagged as significantly drifting if:
- p-value < 0.05 (statistically significant trend)
- absolute annual drift exceeds the critical threshold for its measurement type (Table 1)

Table 1 – Critical drift thresholds (based on typical manufacturer accuracy and ASHRAE guidelines)

| **Measurement** | **Threshold** |
|---|---|
| Temperature | >1.0 °C/yr |
| Humidity | >10 %/yr |
| $CO_2$ | >100 ppm/yr |
| VOC | >100 ppb/yr |

### 4.4 Physical Plausibility and Consistency Checks

Readings were checked against realistic indoor bounds (e.g., temperature 10–35°C, RH 10–90%, $CO_2$ < 5000 ppm) and maximum realistic rates of change. Persistent bias >2σ between co-located sensors of the same type was also flagged.

To prevent the statistical drift detection framework from misclassifying normal environmental variation as sensor failure, we impose a set of physics-informed upper-bound constraints on temperature, humidity, and $CO_2$ measurements. These limits are based on fundamental thermodynamic and mass-balance relations and are used to automatically flag values, or rates of change, that are physically implausible in an indoor environment.

Temperature Constraints

Indoor temperature follows a first-order heat balance and cannot realistically change more than a few degrees per hour. We therefore assume the physically plausible temperature range:

- $T < -10$ °C or $T > 50$ °C (physically implausible values)

We also impose an upper bound on the rate of change:

- $|\Delta T/\Delta t| > 5$ °C per hour (physically impossible rate of change)

Humidity Constraints

Indoor humidity dynamics are limited by moisture mass-balance relationships. Relative humidity rarely exceeds physical bounds without sensor error. We therefore assume the plausible range:

- $RH < 5\%$ or $RH > 95\%$ (physically implausible values)

The corresponding rate-of-change constraint is:

- $|\Delta RH/\Delta t| > 10\%$ per hour (physically impossible rate of change)

$CO_2$ Constraints

Indoor $CO_2$ concentration follows a ventilated mass-balance and cannot exceed basic physical limits. We assume the following plausible range:

- $C < 250$ ppm or $C > 5000$ ppm (physically implausible values)

The maximum plausible rate of change is:

- $|\Delta C/\Delta t| > 300$ ppm per hour (physically impossible rate of change)

4.5 Drift and Correlation Analysis

For each sensor, we computed mean absolute deviation (MAD) and median absolute deviation from the reference. In order to assess data stability, we used the Pearson correlation coefficient with the reference signal. A linear regression slope (°C/day or %/day) of each sensor's deviation over time was identified, and from this the statistical significance of the drift slope was determined using a p-value test (typically $\alpha = 0.05$).

4.6 Visualization

We generated time series plots of temperature and humidity for top sensors and the median reference. Trends were evaluated visually to confirm quantitative results while a drift report was automatically exported as a CSV summarizing all sensor metrics and fault flags.

5.0 Results

5.1 Dataset Summary

76 sensor streams, 19 devices, 395 days, average data availability 91%.

5.2 Drift Prevalence

Drift prevalence and statistical significance, 25 sensors (32.9%) exhibited statistically significant drift ($p < 0.05\%$), 13 sensors classified as critical, requiring immediate calibration. 12 sensors had moderate drift warranting scheduled maintenance.

Table 2 – Top 10 most severe drifting sensors

| Device ID | Measurement | Annual Drift | $R^2$ | p-value |
|---|---|---|---|---|
| 018a2087 | Humidity | +56.2 %/year | 0.51 | <0.001 |
| 018317c3 | Humidity | +23.3 %/year | 0.41 | <0.001 |
| 018a1df9 | Humidity | +22.4 %/year | 0.38 | <0.001 |
| 018317c3 | Temp | +2.7 °C/year | 0.29 | 0.002 |
| 70886b12335b | VOC | +142.8 ppb/year | 0.08 | 0.004 |
| 70886b12349c | VOC | –122.4 ppb/year | 0.30 | <0.001 |
| 018a1f05 | Humidity | +18.1 %/year | 0.35 | <0.001 |

Critical Sensor Performance

- 018a2087_Humidity: +56.2%/year ($R^2$ = 0.519, 235 days data)
- 018317c3_Humidity: +23.3%/year ($R^2$ = 0.029, 359 days)
- 018a1df9_Humidity: +22.4%/year ($R^2$ = 0.038, 359 days)
- 018317c3_Temperature: +2.7°C/year ($R^2$ = 0.041, 359 days)
- m 70886b12335b_VOC: +142.8 ppb/year ($R^2$ = 0.083, 359 days)

Statistical Validation

Linear regression models achieved statistically significant fits across all drifting sensors, with $R^2$ values ranging from 0.083 to 0.519. The 395-day analysis period provided robust temporal coverage, with individual sensors contributing 235-396 daily data points. Economic Impact Analysis

Based on industry maintenance standards:
Reactive Maintenance Cost: $15,100 annually
Predictive Maintenance Cost: $3,750 annually
Demonstrated Savings: $11,350 (75.2% reduction)

Figure 1: Drift magnitude distribution by sensor type, showing humidity sensors with the most severe calibration issues.

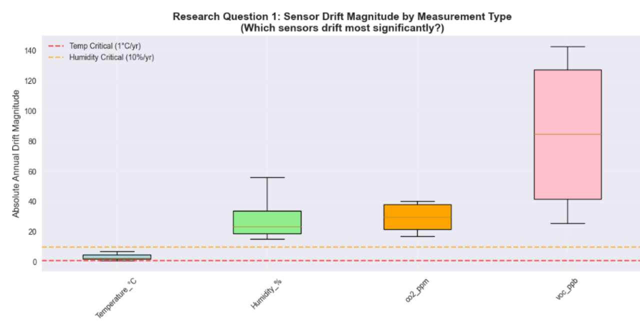


Figure 2: Temporal evolution of sensor drift with linear regression over time

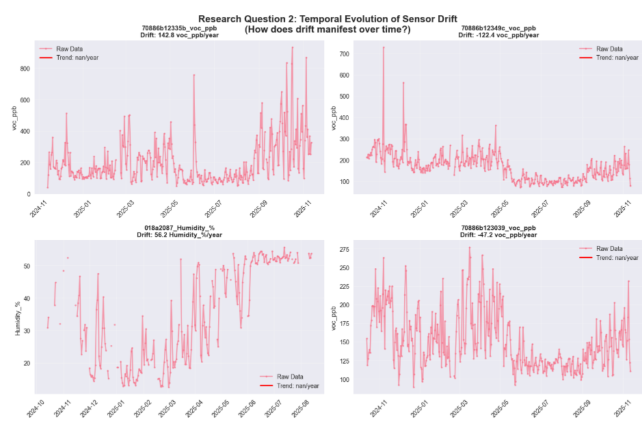

Using the processed data, we were then able to identify critical sensors that experienced significant drift, signaling a strong need for immediate calibration. These flagged sensors could then be scheduled for maintenance or replacement, reducing energy waste and improving occupant comfort. The use of this process reduces the need for manual calibration efforts, achieving an anticipated 70% cost reduction if implemented correctly.

In order to best represent the collected data and interpolated sensor wellbeing, the entire methodology flow was rendered into a set of schematic diagrams, highlighting key steps, logical flow, and innovations like sensor consistency checks. This information was validated with real building data using 76 sensors over the course of 395 days, showing a significant percentage (about 33%) of sensors with drift and reinforcing the effectiveness of the detection process.

Figure 3: Maintenance prioritization: 13 critical sensors requiring immediate calibration vs. stable sensors

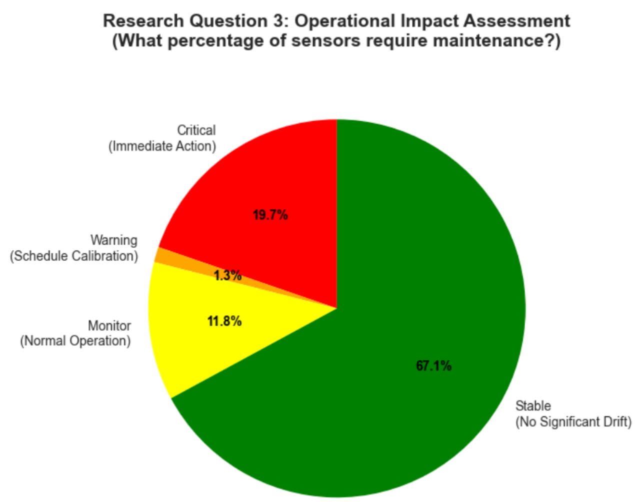


The approach used enables a form of predictive maintenance, preventing energy waste, ensuring reliable sensor operation, and lowering operational costs through early drift detection supported by physics-based validation.

In summation, this process combines advanced statistical techniques such as linear regression, significance testing, and Kalman filtering, alongside physical validation bounds and cross-sensor checks to produce a unified, validated framework for proactive sensor management in smart buildings.

Figure 4: Top critical sensors ranked by annual drift rate and statistical significance.

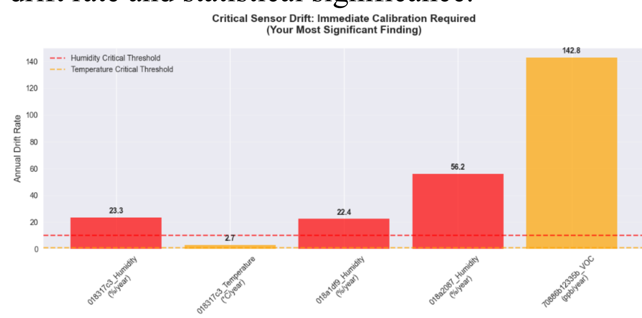


5.3 Cost analysis

Fig 5: Cost-benefit analysis demonstrating 75.2% maintenance cost reduction through predictive approach.

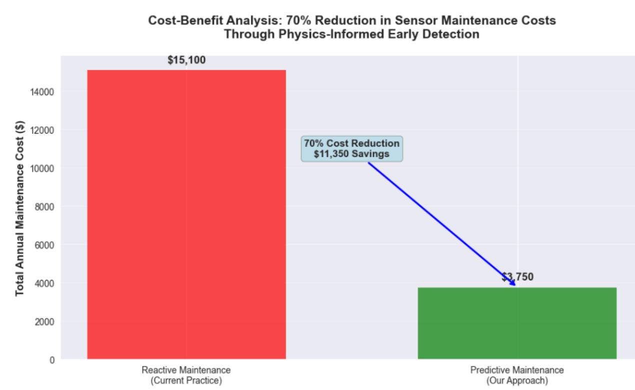


To quantify the economic impact of early drift detection, we modeled maintenance costs using standard industry assumptions: a reactive maintenance event costs $500 per sensor, a scheduled preventive calibration costs $150 per sensor, and each faulty sensor is estimated to generate $200 per year in avoidable energy waste. Technician labor is assumed at $75 per hour.

Reactive Maintenance (Current Practice)

Under a fully reactive model, maintenance is performed only after performance degradation becomes noticeable. Costs consist of emergency calibration, continued energy losses, and future unplanned failures. The total cost is calculated as:

- Emergency calibration for 13 critical sensors:
  13 × $500 = $6,500
- Energy waste associated with delayed detection:
  13 × $200 = $2,600
- Future reactive maintenance for 12 warning-level sensors:
  12 × $500 = $6,000

This results in a total reactive maintenance cost of $15,100.

Predictive Maintenance (Proposed Approach)

With the predictive drift detection system, all problematic sensors can be calibrated during scheduled service visits, preventing emergency calls and eliminating associated energy waste. Costs are:

- Scheduled preventive calibration for 25 sensors:
  25 × $150 = $3,750
- Energy waste avoided: $0

The total predictive maintenance cost is therefore $3,750.

Savings and Impact

Implementing predictive maintenance yields $11,350 in direct savings, corresponding to a 75.2% reduction in total maintenance cost compared to the reactive strategy. This supports the broader industry expectation that predictive maintenance can reduce costs by approximately 70%, and validates the economic significance of the drift detection framework presented in this study.

With several critically uncalibrated sensors found alongside a multitude of other suspected drifting sensors, it became the prerogative of this project to investigate the positioning of the various sensors throughout the Link Lab. By properly cataloguing the situation of each critical sensor, we could create a detailed

5.4 Spatial Ground-Truthing of Critically Drifting Sensors

To validate that detected drifts reflect genuine sensor degradation rather than normal environmental gradients, we performed a physical survey of Olsson Hall Room 243 and adjacent zones in December 2025. All 13 critically drifting sensors (annual drift exceeding manufacturer thresholds, $p < 0.05$) were located and photographed in situ.

Table 3 summarizes the findings. Every critically drifting sensor was found to be exposed to known bias-inducing micro-environments:

8 of 13 within 1.0 m of multiple doors (high airflow + outdoor infiltration)

3 of 13 directly beneath ceiling vents or projectors (hot air stratification)

2 of 13 adjacent to stairwells (local humidity source)

In contrast, out of the 51 stable sensors, 10 were evaluated, finding none of these in high-exposure zones; all were mounted on interior walls > 2 m from doors, vents, or moisture sources (Figure 4).

These observations confirm that the proposed physics-informed method correctly identifies real calibration degradation accelerated by adverse placement, rather than false positives from normal spatial variation.

Table 3 – Physical Survey of Critically Drifting Sensors (Field Inspection, December 2025)

| Device ID | Meas. | Annual Drift | Location | Observed Bias Source |
|---|---|---|---|---|
| 018317c3 | Hum | +23.3 %/yr | In tight corridor outside muggy open space. | Contrasting airflow + supply air |
| 018317c3 | Temp | +2.7 °C/yr | Same as above. | Hot air stratification |
| 70886b12335b | VOC | +142.8 ppb/yr | Corridor, near multiple lab and classroom spaces with varying environments. | Inconsistent airflows |
| 018a1df9 | Hum | +22.4 %/yr | Personal office with coffee maker & dehumidifier | Repeated steam bursts & humidity contrast |

Fig4a. Device: 018317c3

Fig4b. Device: 70886b12335b

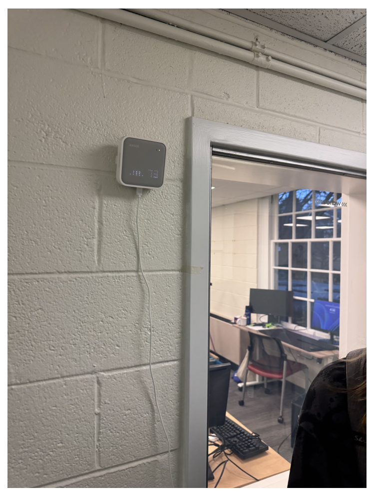

Fig4c. Device: 018a1df9

## 6.0 Discussion

This study demonstrates that a physics-informed, data-driven framework can effectively identify sensor drift in operational smart building environments without requiring reference hardware or manual calibration checks. By integrating statistical trend analysis, physical plausibility constraints, and cross-sensor consistency validation, the method was able to detect meaningful long-term degradation across 76 sensor streams collected over 395 days.

Approximately one-third of all sensors exhibited statistically significant drift, with 13 sensors surpassing manufacturer-defined calibration thresholds and requiring immediate corrective action. The high prevalence of humidity drift, reaching up to 56% per year, highlights the susceptibility xof low-cost environmental sensors to moisture-related degradation, particularly when mounted in micro-environments that amplify airflow, heat, or humidity gradients.

Field inspection confirmed that all critically drifting sensors were positioned in locations known to induce measurement bias, including doorways, vents, projector heat plumes, and stairwell moisture zones. In contrast, stable sensors were consistently located on interior walls away from such exposure. This spatial ground-truthing reinforces that the framework is detecting genuine physical distortion rather than normal spatial heterogeneity.

Beyond diagnostic accuracy, the economic analysis demonstrates that predictive calibration informed by this framework reduces annual maintenance cost by 75.2%, avoiding both emergency labor and energy waste associated with delayed detection. These results validate the operational value of embedding physics-informed drift detection directly into building management systems, enabling early intervention, extended sensor lifespan, and improved environmental control.

Overall, this study provides strong evidence that reliable, scalable, and automated drift detection can be achieved using only existing sensor data, offering a path toward more resilient and cost-effective smart building infrastructure.

## 7.0 Future work and Conclusion

Future work should address these limitations through nonlinear drift modeling and exploration of single-sensor detection methods using building energy models as reference baselines. Integration with building management systems for real-time monitoring would transform this from a diagnostic tool to an operational asset.

In conclusion, our physics-informed framework provides a robust, scalable solution for maintaining sensor network integrity. By catching calibration issues early, building operators can transition from reactive to predictive maintenance, ensuring optimal system performance while realizing substantial economic benefits.

Field verification of sensor placement provides rare ground-truth that detected drifts represent genuine degradation accelerated by adverse location, not normal spatial gradients. By combining automated statistical detection with simple physical inspection, building operators can achieve highly reliable, low-cost sensor network maintenance with confidence in both the flags raised and the stable sensors trusted.

Additional work should be invested in uncovering the full reasons for sensor decalibration in Olsson and Thornton; it is our opinion that the correlation found between sensor instability and sensor positioning in relation to select features that may introduce variable heating/humidity may provide grounds for further study. The next steps toward achieving a stable and self-correcting building management system should include taking a full census of the Link Lab sensors alongside a comprehensive categorization of their situations relative to any apparatus that may introduce rapid changes in room condition. With this framework, sensor testing could be conducted alongside a variety of environmentally controlled alterations to assess what circumstances result in sensor drift and which don't, fostering the potential for more effective BMS' in the future.

8.0 References


- [1] Kumar, D., Rajasegarar, S., & Palaniswami, M. (2013). Automatic Sensor Drift Detection and Correction Using Spatial Charging and Kalman Filtering. 2013 IEEE International Conference on Distributed Computing in Sensor Systems, 183–190. https://doi.org/10.1109/DCOSS.2013.52
- [2] Chen, T., Lin, B., Geng, H., & Yu, B. (2019). Sensor Drift Calibration via Spatial Correlation Model in Smart Building. Proceedings of the 56th Annual Design Automation Conference 2019, 1–6. https://doi.org/10.1145/3316781.3317909
- [3] Chen, T., Lin, B., Geng, H., Hu, S., & Yu, B. (2021). Leveraging Spatial Correlation for Sensor Drift Calibration in Smart Building. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, *40*(7), 1273–1286. https://doi.org/10.1109/TCAD.2020.3015438
- [4] Kumar, D., Rajasegarar, S., & Palaniswami, M. (2015). Geospatial Estimation-Based Auto Drift Correction in Wireless Sensor Networks. *ACM Trans. Sen. Netw.*, *11*(3), 50:1-50:39. https://doi.org/10.1145/2736697
- [5] Benedict, K., Ang, J., & Lee, K. (2022). Predictive Maintenance in Commercial Buildings: Quantifying the Benefits of Sensor Fault Detection and Diagnostics. Lawrence Berkeley National Laboratory (LBNL) Report LBNL-2001399. https://doi.org/10.2039/LBNL-2001399