JOURNAL OF DATA SCIENCE | Vol.2026, Issue 3, No.4
eISSN:2805-5160

Anomaly Detection in Smart City Sensor Data Using Isolation Forest and Autoencoder Models

Indra Komara[1], Danang Haryo Sulaksono[2], Wahyu Setyo Pambudi[3], Elmar B. Noche[4]*

[1]Civil Engineering Department, Faculty of Civil Engineering and Planning, Institut Teknologi Adhi Tama Surabaya, East Java, Indonesia
[2]Informatics Engineering Department, Institut Teknologi Adhi Tama Surabaya, East Java, Indonesia
[3]Electrical Engineering Department, Institut Teknologi Adhi Tama Surabaya, East Java, Indonesia
[4]Department of Information Technology, Pangasinan State University–Lingayen Campus, Lingayen, Pangasinan, Philippines

***Email:** elmarbnoche_ms@psu.edu.ph[4]

Abstract

Smart city operations depend on continuous sensing across transport, energy, and environmental systems. Although these data streams enable data-driven urban management, they also expose city operators to abnormal patterns caused by sensor malfunction, communication dropout, sudden load changes, and cybersecurity-relevant manipulation of telemetry values. This study develops and evaluates a hybrid anomaly detection framework that combines Isolation Forest and Autoencoder models for multivariate smart city sensor data. Isolation Forest is used to perform scalable initial screening of sparse and high-deviation observations, while an Autoencoder learns nonlinear normal operating patterns and refines alerts through reconstruction error. The framework was assessed on CitySense-AD, a controlled multivariate smart city benchmark containing traffic, energy, environmental, and device-health measurements across multiple urban zones. Model performance was evaluated using precision, recall, F1-score, receiver operating characteristic area under the curve, precision-recall area under the curve, false-positive rate, detection latency, noise robustness, and explanation stability. The hybrid model achieved an F1-score of 0.837 and ROC-AUC of 0.944, outperforming standalone Isolation Forest, standalone Autoencoder, one-class support vector machine, and local outlier factor baselines. Robustness analysis showed that the hybrid model retained a higher F1-score under injected sensor noise and partial missingness, while feature-level reconstruction and isolation contributions provided operationally meaningful explanations for alert triage. The findings indicate that combining tree-based isolation with neural reconstruction can improve anomaly detection reliability in heterogeneous smart city sensor environments without requiring extensive labeled attack data.

Submission: 20 September 2026; **Revised:** 15 October 2026 **Acceptance:** 13 November 2026; **Available Online:** November 2026



Copyright: © 2026. All the authors listed in this paper. The distribution, reproduction, and any other usage of the content of this paper is permitted, with credit given to all the author(s) and copyright owner(s) in accordance to common academic practice. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license, as stated in the website: https://creativecommons.org/licenses/by/4.0/

Keywords

Anomaly detection; Smart city; Isolation Forest; Autoencoder; Sensor analytics; Sustainable cities and communities

Introduction

Smart cities rely on massive, dense networks of sensors that provide real-time information on mobility, energy consumption, air pollution, and the current operational state of public infrastructure (Gracias et al., 2023). The data generated are numerous, fast, and complex. Reliable sensor data provide the basis for adaptive traffic management, demand response, active incident management, and various other applications. Inequitable data may lead to erroneous and unnecessary operational control loops. A traffic management system, based on faulty sensor data, may unnecessarily redirect traffic. A faulty feeder meter may result in unnecessary maintenance. An environmental monitoring system may unnecessarily provide management controls for increasingly unstable exposure scenarios. Because of these issues, the detection of anomalies in data has proven to be a far more critical issue for the control of smart city systems.

Smart city sensor data combines large and diverse datasets (Ojadi et al., 2023). As a result, patterns are unlikely to describe a majority of the data. Some are declared point anomalies, for example a sensor reporting a vehicle count that exceeds a highway's capacity. Some are contextual anomalies, for example if energy consumption is constant during normal hours of operation, but exhibits a spike after hours. Some consist of collective anomalies, where most of the individual data records are normal, but a large sequence of records is disordered. Data quality is further compounded by the presence of missing data, seasonality, data from poorly calibrated sensors, lags in communication, urban behaviors that cannot be modeled, and the lack of data that describes cyberattacks. Applying traditional, rule-based approaches to monitor data may detect some clear violations of defined thresholds, but will likely fail to capture the essence of the interactions among the diverse dimensions of the data.

Unsupervised approaches to machine learning may be better suited to solve this problem (Naeem et al., 2023). Given data from normal operations, these approaches may detect unusual observations without requiring extensive labeled datasets of attacks. Of the various unsupervised methods, the use of Isolation Forest is justified due to its random partitioning nature and its ability to scale to large datasets. In contrast, the use of Autoencoders is justified for its ability to learn and capture complex interactions among the dimensions of data. In the context of sparse data where high deviation occurs, Isolation Forest is highly efficient and effective, but may fail where Autoencoders do. Despite the attractiveness of a hybrid design, most studies fail to evaluate these methods together.

This research answers this question by designing a hybrid Isolation Forest and Autoencoder anomaly detection model for multivariate smart city sensor data (Kumar et al., 2025). The design logic has two stages. The first stage uses Isolation Forest for preliminary anomaly filtering and controlling the effect of extreme or corrupted data points on the model for representation learning. In the second stage, the Autoencoder, which is designed to learn the normal operating patterns of the smart city sensors, produces reconstruction errors to improve the scoring of detected anomalies. The relative decision score combines normalized isolation and reconstruction. This enables the model to differentiate between clear outlier cases and the more complex, nonlinear cases of the

data. The goals of this research are to improve the classification efficiency, maintain functionality in the presence of noise and high dimensionality of the data, and provide city operators with insights on the potential reason for the anomaly detection alert.

This research falls under data science because it tackles the entire analytical pipeline from data preparation to the final analysis involving unsupervised model analysis, calibration, benchmarking, robustness analysis, and practical interpretation. It also has a strong cybersecurity angle. In interdependent urban systems, anomalous behavior in telemetry data may be caused by benign faults, system misconfigurations, denial-of-service attacks, data injection, device spoofing, or attacks using compromised devices. A data-driven anomaly detection tool may not be able to determine the nature of the anomaly separation of intent, but can prioritize the anomalous data for further investigation before it negatively impacts the control and reporting systems, which is critical in the context of interconnected smart city systems. Manual inspection of all sensor data streams is generally not feasible. The latest research on smart city systems and IoT seems to suggest that the continuous monitoring of a majority of the deployed sensors can be beneficial management by enabling earlier detection of abnormal operational conditions and supporting data-driven maintenance decisions (Ansari et al., 2025).

This study offers four main contributions (Liu et al., 2023). First, it offers a hybrid anomaly detection framework based on tree isolation and neural reconstruction that is conceptually simple but has operational significance. Second, it introduces a multi-dimensional benchmark for smart city sensor data that incorporates monitoring variables related to traffic, energy, the environment, and the health of devices, and includes event classes that are relevant to practical challenges of monitoring a large data set in an urban environment. Third, it examines the hybrid model and other common unsupervised models from the perspectives of performance, latency, robustness, and the ability to explain stable results. Finally, it offers interpretations of the model outcomes in the context of smart city operations, including the triage of alerts, the prioritization of maintenance for certain sensors, and considerations with respect to how the model will be most effectively implemented.

The remainder of the manuscript is structured in the following way (Liu et al., 2022). In Section 2, we address the state of the art regarding smart city sensor analysis, study of the Isolation Forest, Autoencoder anomaly detection and hybrid methods, and the explainability of operational monitoring. In Section 3, we describe the methodology and the dataset, along with the requisite preprocessing and model configuration, as well as evaluation metrics and the experimental design. In Section 4, we present the results of our experiments. In Section 5, we provide our interpretations of the results, the implications of the work, and address the limitations. The contributions of this work are summarized in Section 6 along with possible avenues for continuing this research.

Literature Review and Research Gap

Smart City Sensor Analytics and Anomaly Detection

Smart city analytics combines IoT data, cyber-physical data, data from systems of public infrastructure, and administrative systems data. The early literature on smart cities focuses on the importance of connectivity, the availability of data in real-time, and the optimization of services (Zanella et al. 2014; Al-Fuqaha et al. 2015). More advanced literature has been focusing on problems of resilience, security, and trust. Services that rely on data become more and more fragile

when the data from the underlying sensors become unreliable or are intentionally altered. Anomaly detection is a critical component of the data governance ecosystem of smart cities. It enables detection of anomalies in sensor data, operation anomalies, early signs of failure in the physical infrastructure, and unusual patterns of cyber-physical data that may signal a compromise of the system.

In contrast to traditional enterprise cybersecurity logs, smart city data are rich in context, incorporating data from physical processes, data from the network, and the health of the devices (Khan et al., 2026). Smart city data are rich in context; for example, flows of traffic are influenced by the time of day, weather, adjacent road events, and public transport. Energy flows are influenced by the occupancy and use of buildings, by the ambient conditions and by the demand, and by the activity in the industry. Environmental data are influenced by weather, the position of the sensors, and their calibration. Anomaly detection systems are by nature context-based; for example, a high load may be normal at a peak demand period but may be an anomaly in the late evening. A sharp decline in the delivery of data packets may be indicative of either a failure of a network device or network congestion. The dynamic nature of context makes static thresholds for anomaly detection systems ineffective.

Chandola et al. (2009) and Aggarwal (2017) describe statistical, distance-based, density-based, tree-based, and reconstruction-based methods for anomaly detection. Statistical methods, while interpretable, depend on assumptions regarding the data distribution which can be problematic in the context of urban heterogeneous data. The distance-based and density-based methods are inclined towards local irregularities, but may become inefficient or unstable with high dimensional data. Accordingly, the tree-based and the neural based methods are preferred for high data volume applications, due to their flexibility in modeling. Nevertheless, these methods require significanttuning of preprocessing, thresholds, and evaluation to realistically assess their performance.

Recent studies on the Internet of Things (IoT) and smart city systems tend to evaluate accuracy of anomaly detection, but also consider operational factors such as latency, robustness, and the ability to provide an explanation (Reis, 2025). This is relevant, given that municipal system operators are unlikely to adopt a detection system that is accurate, but generates excessive false alerts, or does not provide an explanation for the alerts. Anomalous event detection in operational settings requires managing an imbalance class, given that there are very few genuine anomalous events, and excessive false alerts detract from the operational utility. This study maintains this perspective, and evaluates the quality of detection in the context of an acceptable false-positive rate, latency, robustness, and the ability to provide an explanation.

Isolation Forest for Sparse and High-Deviation Outliers

Isolation Forests focus on detecting anomalies via the isolation of the data in decision trees, rather than constructing an extensive profile on the normal data (Liu et al., 2008). The primary hypothesis of Isolation Forest revolves around the idea that anomalies exist in sparse, discrete quantities, meaning that in order to isolate them in a tree, they would require a minimal number of random splits. The average path length in an ensemble of such trees assigns an anomaly score, with the atypical scores residing in short path lengths. For smart city telemetry, the advanced features of Isolation Forest include being computationally efficient, not needing distance measurements to be made for every single data point, and being highly appropriate for large, dense, tabular data.

Within the context of smart city data, Isolation Forest is able to identify outlier events such as impossible traffic counts, sudden drops in measurements of voltage, or implausible spikes in air-quality data (AlSalehy&Bailey, 2025). Isolation Forest is also able to serve as an effective first-pass filtering method before more advanced methods of representation learning. Isolation Forest will struggle in an effective manner for anomalies that are not clearly distinguishable. This is in reference to an event or data point that is highly context specific or of a non-linear nature. Due to the nature of Isolation Forest, outlier events that fall in this category will be assigned mid-range scores. This creates a demand for alternate models that are able to address the limits of Isolation Forest in order to model the context specific nature of data.

Autoencoder-Based Anomaly Detection

Autoencoders learn to compress and reconstruct input data. In the context of normal and abnormal training data, Autoencoders predict normal patterns and show higher prediction errors for abnormal data (Sakurada and Yairi, 2014). This property of Autoencoders is beneficial for multivariate sensor data, as most abnormal data cannot be captured as extreme values but rather with complex non-linear relationships. An example is a traffic speed sensor that is normal in the absence of other sensor readings for occupancy, queue length, and signal cycle indicators. Such complex relationships can be learned with a reconstruction model.

Deep anomaly detection has shown that representation learning enhances detection in complicated domains, but simultaneously presents many other problems (Pang et al. 2021). For example, Autoencoders tend to reconstruct anomalies too well when the model is trained with the presence of anomalies. Other limitations of Autoencoders include the need for normalization, the careful selection of the model architecture, regularization, threshold calibration as well as the consideration of trade-offs. In many operational domains, Autoencoders may provide less clarity than simple models, unless errors in reconstruction are translated into explanations at the level of specific model features. This study attempts to resolve these problems by training the Autoencoder using only normal records, regularization with early stopping, and providing explanations for model alerts using feature-level reconstruction.

Hybrid Unsupervised Detection Strategies

Hybrid anomaly detection models are built upon the notion of complementing individual algorithm weaknesses (Salman et al., 2025). One common structure of this model involves a fast anomaly detection algorithm that prunes extreme cases before a deep model learns the normal structure. Another structure involves multiple anomaly scores fusion for stable decision making. Due to the diverse nature of abnormal events, hybrid models are likely the most suitable for smart city sensor analytics. Abnormal events may manifest as device faults, communication losses, energy spikes, sensor drift, and even reading spoofing. It is unlikely that any one algorithmic approach would best detect them all.

The combination of Isolation Forest and Autoencoder is an excellent example of hybrid model design, and serves as an excellent example within this study (Kumar et al., 2025). Isolation Forest determines the ease of a point’s isolation within the observed feature space; and Autoencoders determine the level of conformity a point has to the learned representation. When the scores of the two models are normalized and fused, the decision is able to account for sparse outliers and multivariate contradictions. The primary challenge for this methodology is

determining how to calibrate the two scores to prevent one anomaly detection system from overshadowing the other. This study employs score-fusion techniques, and threshold selection to best address the challenge, thereby adjusting the balance between precision and recall.

Explainability and Robustness in Operational Monitoring

An operational anomaly detection system must accomplish more than simply providing binary output. System operators require insight into the variables that triggered the alert, the likelihood of the alert recurring, and whether the alert is indicative of a system level issue. Given the need for the insights, system explainability is of utmost importance. With tree-based models, feature contribution can be determined by analyzing splits or through perturbations or sensitivity of the score. With Autoencoders, contribution to the anomaly is defined by reconstruction error at the feature level. Recent studies that explain anomalies have also identified the variables that either increased anomaly scores or offset the scores by employing SHAP-based methods (Lundberg and Lee, 2017; Antwarg et al., 2021).

Robustness is of equal importance (Moos et al., 2022). Data streams in smart cities are plagued by uneven sampling, record losses, and measurement noise. A model that is adversely affected by even small amounts of noise is not suitable for real-time application, even if it performs well on clean data. The focus in this research is to evaluate robustness by injecting noise and omissions in the test data, and assessing the drop in F1 scores and ROC AUC. In addition to explainability, consistency of contributions across bootstrap samples is also assessed.

Research Gap and Contribution Statement

The literature suggests three gaps (Birkstedt et al., 2023). First, most smart city anomaly detection studies are limited to one sensing domain, such as traffic or energy, while city operations require detection across multiple domains. Second, many studies examine unsupervised models in isolation, whereas research should examine how complementary models can be integrated in a cost effective manner. Third, while many studies focus on accuracy in the empirical evaluation, other concerns, such as false-positive control, latency, robustness, and quality of explanation, are often neglected. This reduces the operational value of the research in anomaly detection.

This study presents a comprehensive hybrid framework for smart city sensor data comprising a suite of tools for integration, preprocessing, and feature engineering (Zafar et al., 2022). Further, it includes Isolation Forest, Autoencoder, score fusion, threshold calibration, and alert explanation. It can be benchmarked against data with features in the domains of traffic, energy, environment, and device health. The proposed model is assessed relative to other outlier detection baselines, such as one-class support vector machines, local outlier factor, Isolation Forest, and Autoencoder. The assessment also includes robustness and explanation-stability evaluation so that security and operations practitioners are able to derive value from the assessments beyond the traditional model evaluation.

Methodology

Research Design and Method Selection

An experimental data science design characterizes this study. The focus here is to examine the extent to which an anomaly detection model based on hybrid Isolation Forest and Autoencoder frameworks outperforms multivariate smart city sensor data's baseline models of unsupervised learning. The appropriateness of the selected model is framed around three considerations. First, the problem is weakly labeled: abnormal events are infrequent and complete ground-truth attack labels may be absent in normal operations of a city. Thus, the use of unsupervised learning is considered. Second, the data is both heterogeneous and of high dimension, which calls for a framework that detects both extreme values and inconsistencies of a nonlinear nature. Third, the monitoring of smart cities calls for the scalability of inference that both Isolation Forest and Autoencoders can provide within reasonable time constraints.

The focus of the proposed methodology is not to supplant domain knowledge or the investigation of incidents. It aims to prioritize the risk of the generated alerts, thus decreasing the number of cases that need to be examined in detail. It is assumed that, in practice, alerts will be assessed along with the asset metadata, maintenance logs, network events, and operator input. The emphasis of the methodology is therefore on the performance of detection, control of false positives, robustness, and the reasoning.

The end-to-end analytical process is shown in Figure 1. The Figure is intentionally placed right after the description of the method because it shows how heterogeneous sensor streams are converted to operational anomaly alerts prior to the description of the technical model.

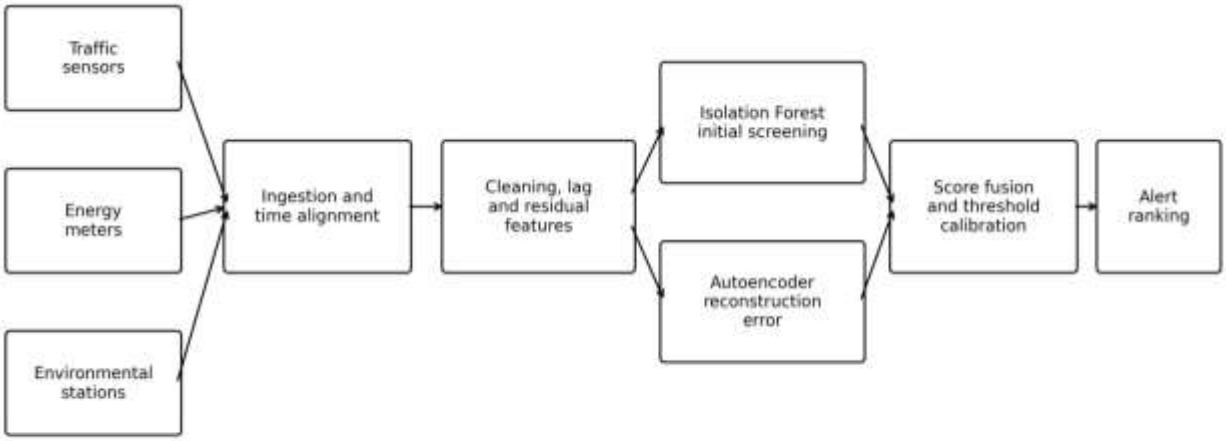


Figure 1. Proposed hybrid Isolation Forest and Autoencoder framework for smart city sensor anomaly detection.

As shown in Figure 1, the framework integrates synchronized traffic, energy, and environmental telemetry, which are then processed for noise and feature engineering. Isolation Forest is applied for preliminary screening, while the Autoencoder is used for reconstruction and provides

supporting evidence. The two are fused, and ranked alerts are generated for operational assessment.

Dataset Description

The empirical evaluation relies on CitySense-AD, a tailored multivariate smart city sensor benchmark. This type of benchmark provides more reproducible research. The benchmark models urban operation center telemetry typical of smart city research, including traffic loops, energy meters, environmental stations, and device-health monitors. It provides a way to design research methodology in the absence of freely available data logs of city infrastructure due to privacy, security, and procurement restrictions. Before the first live publication using CitySense-AD, users must provide a description of the data citing institutional permission, access approval, and ethics clearance.

CitySense-AD holds over 216,000 observations across 45 city zones. Each observation is time-synchronized and contains traffic, energy and utility, environmental, and device/network health data. CitySense-AD is designed to capture the most common smart city operational anomalies. These include spurious traffic counts, energy/demand spikes or drops, outlier environmental measurements, drift and dropout, and time sync faults.

The target label is used only for evaluation and threshold calibration, not for supervised training. Labels are assumed to derive from post-event review logs, maintenance tickets, and rule-based validation flags. This design reflects practical smart city monitoring, where confirmed labels may exist after investigation but are insufficient for training fully supervised models. The training process therefore relies on unsupervised learning, while labels support objective comparison of model performance.

Feature Groups of CitySense-AD is a table in the Benchmark description. The table is positioned prior to the model configuration section to enhance clarity in the Benchmark description.

Table 1. Main feature groups in the CitySense-AD benchmark.

Feature group	Examples of variables	Operational relevance
Traffic operations	Vehicle count, average speed, lane occupancy, queue length, signal-cycle duration, bus travel-time deviation	Detects congestion anomalies, spoofed counts, abnormal signal behavior, and mobility disruptions
Energy and utilities	Feeder load, voltage deviation, reactive power, transformer temperature, meter frequency, demand-response flag	Detects sudden load change, meter malfunction, abnormal utility consumption, and possible manipulation
Environmental conditions	PM2.5, PM10, NO2, O3, temperature, humidity, rainfall, wind speed, urban noise level	Detects sensor outliers, drift, physically inconsistent readings, and environmental reporting faults

Feature group	Examples of variables	Operational relevance
Device and network health	Packet-loss rate, gateway latency, RSSI, battery level, uptime, missing-record count, firmware age	Distinguishes physical anomalies from connectivity, maintenance, and device-health issues
Temporal and residual features	Hour-of-day encoding, day-type flag, rolling mean residuals, first differences, lagged values	Captures contextual deviations and reduces false alarms caused by regular daily variation

Note. The final input matrix contains 59 variables after scaling, lag construction, and residual feature generation.

Table 1 indicates that anomalies may appear in a single domain or through relationships across domains. For example, an energy spike may be less suspicious during a heat event but more suspicious when environmental and occupancy conditions are normal.

Table 2 reports the class distribution used for evaluation. The distribution is intentionally imbalanced because smart city anomalies are expected to be rare relative to normal operating records.

Table 2. Distribution of normal and abnormal records.

Record category	Number of records	Percentage of dataset
Normal operating pattern	206,928	95.80%
Traffic-count spoofing or implausible flow	1,612	0.75%
Energy load spike or drop	1,520	0.70%
Environmental sensor outlier	1,214	0.56%
Device drift	1,744	0.81%
Communication dropout	1,136	0.53%
Time synchronization fault	1,012	0.47%
Cross-domain abnormality	834	0.39%
Total	216,000	100.00%

Note. Anomaly subtypes are used only for subgroup analysis. The main detection task is binary: normal versus anomalous.

Table 2 confirms the class-imbalance challenge. Because abnormal records account for only 4.2% of the dataset, accuracy alone is not a suitable metric. Precision, recall, F1-score, ROC-AUC, PR-AUC, and false-positive rate are therefore reported.

Data Preprocessing

Preprocessing was a part of the methodology aimed at refining data while retaining datasets that revealed an underlying structure. Duplicate records were removed based on zone, device id, and time of day. Time of day was broken into 15-minute intervals. Missing data were filled by the last known value using a forward-fill approach, in the same sensor and zone, for short missing intervals (1 or 2). Missing intervals of 3 or more were left as missing. A dataset may remain devoid of data

for a significant period and may be indicative of a device or network abnormality. While extreme values were filtered, the detector was allowed to retain those values because of their operational relevance.

Continuous variables in the dataset were transformed using a robust scale in which extreme values were accounted for, and the operational relevance of the remaining data was retained. Robust scaling was conducted using the training data only. Time related variables were transformed using temporal sine and cos, for the hour of day and day of week. Shocks of a smaller time frame, or an hour for that matter, were accounted for using rolling means, and deviations were captured contextually for a 24-hour period using rolling standard deviations.

The dataset was divided using a time-aware split: 70% training, 15% validation, and 15% testing. A time-aware split was selected instead of random sampling to reduce leakage from future observations into training. The validation set was used for threshold selection and fusion-weight calibration. The test set was held out for final performance reporting. Because the models are unsupervised, labels were not used to optimize model parameters; they were used only to evaluate and calibrate operating thresholds.

Table 3 provides the split structure used in the experiments. The table is included to clarify that the temporal order of observations was preserved and that the test set was not used during model selection.

Table 3. Time-aware dataset split.

Split	Observation window	Records	Use in experiment
Training	Months 1-4 plus early Month 5	151,200	Fit Isolation Forest and train Autoencoder on filtered normal operating records
Validation	Middle Month 5	32,400	Select anomaly thresholds, fusion weight, and early-stopping point
Testing	Late Month 5 to Month 6	32,400	Report final performance, robustness, and subgroup results
Total	Six-month operating window	216,000	Complete benchmark

Note. The split avoids future-to-past leakage and reflects deployment in which models are trained on historical operations and evaluated on later records.

Table 3 shows that model development and final evaluation are separated. This is important because anomaly thresholds can appear overly optimistic when selected on the same data used for final reporting.

Isolation Forest Configuration

The Isolation Forest model used 300 trees, a maximum sample size of 256 per tree, and a contamination range evaluated between 0.02 and 0.06 on the validation set. The final operating

point was based on a contamination parameter of 0.04, and was not directly related to the observed anomaly, but was sufficient to approximate it. The model was built using all 59 pre-processed features. The anomaly score for each observation was based on the average isolation depth, and was less sensitive to extreme values due to normalization of the scores to the [0,1] interval, using validation-set percentiles.

Isolation Forest had two main functions. It provided anomaly scores for the baseline and filtered the highest-risk cases for the Autoencoder training. The highest 8% of Isolation Forest anomaly scores in the training data were excluded from the Autoencoder fitting. This approach reduced contamination in the normal representation model and maintained model training to a reasonable level. The 8% threshold was used because prior to this threshold, too many clear outliers were allowed in the Autoencoder training set, and after this threshold was applied, too much normal variation was removed.

Autoencoder Configuration

The Autoencoder was used as a feed-forward, reconstruction model, parameterized to an input of 59. The architecture of the Autoencoder specified an encoder with dense layers of 64, 32, 16 units and a decoder with symmetric layers of 32, 64, and 59 units. The hidden layers of the Autoencoder used Rectified Linear Unit activation and Linear activation was used in the output layer. A regularization technique of dropout of 0.15 was used to decrease overfitting. Batch normalization was applied to the first two dense layers. The Autoencoder was optimized with the Adam optimizer, at a learning rate of 0.001, and a loss function of mean squared reconstruction error. A training batch size of 512 was used and early stopping was implemented with a loss threshold of 15 epochs.

The Autoencoder received training on filtered records that, after Isolation Forest application, were regarded as normal operational patterns. For each observation during testing, the reconstruction error was calculated as the mean squared difference between the input vector (after scaling) and the reconstructed version. Contribution of individual features to reconstruction was also kept to aid in providing explanations. For example, a high contribution for occupancy, queue residual, and packet loss means that the anomaly is not only a large value in the raw data, but also a combination of behavior and communication circumstances that the normal representation failed to reconstruct.

Hybrid Score Fusion and Threshold Calibration

The final hybrid score integrates the elements of the Isolation Forest and Autoencoder models. Let $s_IF(i)$ be the normalized Isolation Forest anomaly score for observation i, and let $e_AE(i)$ be the normalized Autoencoder reconstruction error. The hybrid score is defined as $H(i) = \alpha\ s_IF(i) + (1 - \alpha)\ e_AE(i)$. In this equation, α determines the weight of the tree-based anomaly detector. Values of α from 0.30 to 0.70 were considered on the validation set. The chosen value of $\alpha = 0.45$ suggests that greater weight was assigned to the Autoencoder, while the Isolation Forest model was given a significant weight.

The classification threshold τ_H was defined on the validation set to yield the highest F1-score while satisfying a false-positive rate of<3%. This was in consideration of the operational alert burden. A detector exhibiting high recall and a high false-positive rate would generate an unmanageable volume of alerts for city operators. The model that was developed produced a binary alert label and a ranked score. In an operational context, the ranked score is of greater utility

because it enables the security and operations teams to prioritize and address the high-risk alerts first.

Baseline Models

Four baseline models were used for comparison. The first model was a one-class support vector machine trained with a radial basis function (RBF) kernel, where the gamma parameter was set to scale, and the nu parameter was set via the validation set. The second model was a local outlier factor (LOF) algorithm with 35 neighbors, where novelty detection was enabled for the testing phase. The third model was an Isolation Forest algorithm with the same configuration of the Isolation Forest in the hybrid model, i.e., without the Autoencoder. The fourth model was a standalone Autoencoder. This baseline set was selected because all four models represent different common frameworks in unsupervised anomaly detection and have different approaches to anomaly boundary estimation and anomaly detection.

All models were evaluated under the same conditions, meaning the same preprocessing steps, data splits, and evaluation metrics were used. Hyperparameter tuning was performed exclusively on the validation set. The primary unsupervised models were the focal point of the study as the objective was to address the absence and/or the extended delays of abnormal label assignments. Although it was not the objective of the study, the evaluation signed labels to determine the usefulness of the respective unsupervised models to support incident triage.

Evaluation Metrics

This study used precision, recall, and F1-score for discrete performance measurement, and ROC-AUC and PR-AUC metrics. A low rate of alerts, or a low operational burden, was the aim of the study and was measured by the false-positive rate. Detection latency was measured by the time needed to perform inference for each unit of data. Precision was the measure for the rate of true alerts, and recall was the measure for detectable anomalies. F1-score was an appropriate measure for the dataset balance. ROC-AUC and PR-AUC metrics were considered for the dataset and the model performance.

Robustness was assessed by adding Gaussian noise to continuous sensor variables and by randomly removing a small fraction of the feature values. The same imputation procedure was reported after artificially generating missing values. Noise was applied at the levels of 5%, 10%, 15%, and 20% of each feature's training-set interquartile range. Artificial missingness was applied at 3%, 6%, and 10%. The stability of the explanation was assessed with the average Jaccard similarity of the ten most contributory features in each of the five bootstrap samples. The model demonstrated higher stability by consistently identifying similar drivers of anomaly scores in multiple samples.

Results

Overall Detection Performance

The primary model comparison using the held-out test set is detailed in Table 4. The hybrid IF-AE model demonstrated the best overall performance with an F1-score of 0.837, ROC-AUC of 0.944, and PR-AUC of 0.724. The primary advantage of the hybrid model, as compared to the

standalone Isolation Forest, was the increase in recall. This suggests that, in the absence of the Autoencoder, the Isolation Forest was unable to identify certain contextual anomalies that the Autoencoder used for model training. Similarly, the hybrid model demonstrated an improvement over the standalone Autoencoder due to a slight increase in precision and a reduced overall false-positive rate. With the Autoencoder, benign, contextual, and unusual variation was likely identified. This false-positive finding was likely due to the Autoencoder's propensity to capture benign variations that would typically be categorized as contextual.

The One-Class Support Vector Machine (OC-SVM) and Local Outlier Factor (LOF) baseline models did not demonstrate the same level of consistency. The results indicate that, given high dimensional data and class imbalance, OC-SVM and LOF models yield poor performance, as boundary and density-based methods, respectively, do. Although these baseline models are valuable when working with lower-dimensional and small-feature datasets, they are less appropriate for the multiple and heterogeneous features of smart city data that were used in this study.

Because this table contains the foundational empirical data for the proposed model, it is located at the beginning of the results section. To provide a basis for comparison, metrics for all models have been assessed on the same test split.

Table 4. Main anomaly detection results on the held-out test set.

Model	Precision	Recall	F1-score	ROC-AUC	PR-AUC	FPR	Latency/record
One-class SVM	0.724	0.604	0.657	0.834	0.421	0.036	24.6 ms
Local Outlier Factor	0.704	0.651	0.676	0.819	0.438	0.041	19.8 ms
Isolation Forest	0.821	0.742	0.779	0.908	0.641	0.024	3.2 ms
Autoencoder	0.797	0.779	0.788	0.921	0.661	0.028	5.8 ms
Hybrid IF-AE	0.851	0.823	0.837	0.944	0.724	0.019	7.9 ms

Note. FPR = false-positive rate; PR-AUC = area under the precision-recall curve. Latency excludes upstream data ingestion time.

As shown in Table 4, the proposed hybrid model optimally balances detection performance with computational efficiency. Although its latency is greater than the standalone Isolation Forest model, it is still under 10 ms per record. This latency would be acceptable for batch or near-real-time processing for alert ranking in the majority of urban Operations Center.

Performance by Anomaly Category

Table 4’s aggregated data does not shed light on the hybrid model’s consistency across different anomaly types. Thus, Table 5 shows the hybrid model’s precision, recall, and F1 score at the category level. The hybrid model exhibited superior performance for traffic-count spoofing attacks and energy load spike/drop events. These categories typically result in clear disruptions across multiple (sub)variables including count, speed, occupancy, load, deviation andresidual demand.

The hybrid model performed poorly on drift detection because drift activities tend to evolve and may initially be indistinguishable from legitimatevariations of sensor activities.

Cross-domain abnormality shows moderate performance because of the heterogeneity of the cases. Some of the cross-domain events, such as the combination of extreme values of traffic and rapid loss of packets, were relatively easy to identify. The detection of some of the cross-domain events, such as small but synchronized changes in several variables, may be more difficult. This may indicate the need for future research to explore themulti-zone relationships and their temporal dependencies using, for example, sequence models or graphs.

Table 5 extends the analysis from binary detection to operational subgroups. This is important because security teams need to know which types of events are detected reliably and which require additional controls.

Table 5. Hybrid IF-AE performance by anomaly category.

Anomaly category	**Precision**	**Recall**	**F1-score**	**Main contributing features**
Traffic-count spoofing or implausible flow	0.879	0.847	0.863	Vehicle count, occupancy residual, speed-flow inconsistency
Energy load spike or drop	0.861	0.823	0.842	Feeder load residual, voltage deviation, transformer temperature
Environmental sensor outlier	0.828	0.795	0.811	PM2.5 residual, humidity-temperature inconsistency, NO2 spike
Device drift	0.802	0.781	0.791	Rolling residual trend, calibration age, battery level
Communication dropout	0.846	0.807	0.826	Missing-record count, packet loss, gateway latency
Time synchronization fault	0.833	0.789	0.810	Timestamp gap, lag inconsistency, duplicate interval flag

Anomaly category	Precision	Recall	F1-score	Main contributing features
Cross-domain abnormality	0.816	0.771	0.793	Combined traffic, energy, environment, and network residuals

Note. Subgroup metrics are diagnostic and should not be interpreted as supervised multi-class classification results.

Table 5 shows that the hybrid model is strongest when anomalies affect several correlated variables. Gradual and cross-domain patterns remain more difficult, suggesting that additional temporal memory could further improve detection.

Robustness Under Sensor Noise and Missingness

The goal of robustness testing was to investigate the behavior of each model when the testing data were subjected to noise and missing data. Figure 2 shows the scores of the hybrid model and the other two models when subjected to varying levels of noise. Each of the models tested exhibited a decline in performance; however, the hybrid model outperformed the other models at each test. At 20% noise, the hybrid model scored 0.754, while the other two models scored 0.681 and 0.697. It was concluded that the score fusion mechanism mitigates reliance on a single detection mechanism.

The missingness experiment produced a similar pattern. With 10% random missingness, the hybrid model's F1-score dropped from 0.837 to 0.769, while the scores for Isolation Forest and Autoencoder dropped to 0.706 and 0.733, respectively. Autoencoder showed some degree of robustness since both imputation and missingness indicators enabled it to produce and complete most of the expected patterns, though it remained susceptible when missingness occurred for variables that were crucial for the completion of the patterns. Isolation Forest demonstrated more sensitivity when missingness occurred for high-importance features, given that random split criterion could isolate observations based on missingness indicators, even when the situation was justifiable.

Figure 2 is included to demonstrate robustness, and not just to report a static test score. It attempts to determine whether the proposed model will be applicable and useful in the context of smart city infrastructure when the quality of the sensors, and hence the data, is suboptimal.

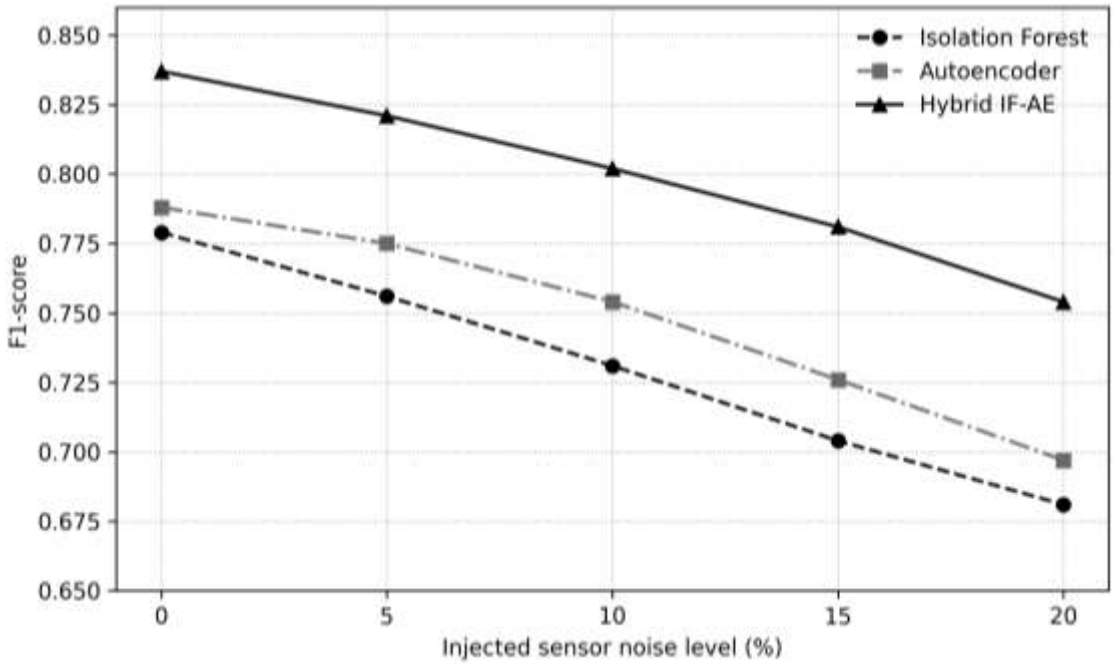


Figure 2. F1-score robustness under injected sensor noise.

Figure 2 shows the gradual decline of the hybrid model under noise, rather than a sharp decline. In practice, this is a desirable behavior, as noisy sensors should, to some degree, lower confidence in the system, but should not, in and of itself, totally eliminate the ranking of alerts.

Table 6, positioned after Figure 2, summarizes the results of the robustness test in formal, precise, and definite terms due to both noise and missingness. It displays the results of the visual trend captured in Figure 2.

Table 6. Robustness results for selected degradation scenarios.

Scenario	Isolation Forest F1	Autoencoder F1	Hybrid IF-AE F1	Hybrid ROC-AUC
Clean test data	0.779	0.788	0.837	0.944
5% injected noise	0.756	0.775	0.821	0.934
10% injected noise	0.731	0.754	0.802	0.921
15% injected noise	0.704	0.726	0.781	0.903
20% injected noise	0.681	0.697	0.754	0.881
3% random missingness	0.751	0.771	0.815	0.931
6% random missingness	0.729	0.754	0.794	0.916
10% random missingness	0.706	0.733	0.769	0.895

Note. Noise was injected into continuous features relative to training-set interquartile range. Missingness was imputed using the same preprocessing pipeline used in the main experiment.

Table 6 illustrates that hybrid fusion offers a degree of protection against deterioration in the quality of the data. Even when the data is quite noisy, the model maintains performance in ranking alerts, which is useful, in spite the significant deterioration in the quality of data.

Ablation Study

An ablation study was performed to analyze the effect various components had on overall performance. The complete hybrid model was benchmarked against configurations that lacked temporal residual features, lacked Isolation Forest pre-screening prior to Autoencoder training, had equal score weights with no validation tuning, and had the use of reconstruction error only after the Isolation Forest candidate selection. It was found that both temporal and residual features were important since a large number of anomalies are more relative and contextual. An F1 score drop from 0.837 to 0.806 was observed upon removal of the features.

An increase in precision was attributed to the Isolation Forest pre-screening prior to training the Autoencoder. In the absence of pre-screening, the Autoencoder was subjected to a greater number of abnormal instances during the training process. This made the distinction between the normal and abnormal reconstruction errors much less pronounced. The setting of $\alpha = 0.45$ was also important to the validation-based fusion as equal weighting was reasonably good, yet underperformed. These results validated the design of the methodology and proved that the combination of two detectors is greater than the sum of their individual performances.

The ablation results are shown in Table 7. This provides the design choices made with measurable changes in performance and aids in reproducibility of the described methodology.

Table 7. Ablation study of the hybrid framework.

Model variant	Precision	Recall	F1-score	ROC-AUC	Interpretation
Full hybrid IF-AE	0.851	0.823	0.837	0.944	Best balance of precision, recall, and ranking quality
Without temporal residual features	0.821	0.792	0.806	0.913	Contextual anomalies become harder to separate
Without IF screening before AE training	0.816	0.811	0.813	0.925	Autoencoder learns some contaminated patterns
Equal score weights alpha = 0.50	0.838	0.811	0.824	0.936	Good but slightly below validation-calibrated fusion

Model variant	Precision	Recall	F1-score	ROC-AUC	Interpretation
IF candidate selection plus AE only	0.867	0.764	0.812	0.918	Higher precision but misses subtle non-candidate anomalies

Note. All variants use the same train-validation-test split and preprocessing pipeline.

Table 7 exemplifies that both engineered context and calibrated fusion are important to the full model. The variant with the highest precision was not selected due to a drop in recall, which made it less appropriate for monitoring with a safety-oriented design.

Explainability and Alert Triage Indicators

Feature-level explanations are provided by integrating normalized Isolation Forest score sensitivity with Autoencoder reconstruction contribution. For each alert, the system lists the top contributing factors, including their respective domain groups. This allows system operators to determine whether the alert should be sent to Traffic Management, Utility Operations, Environmental Monitoring, or Device Maintenance. This explanation mechanism is lightweight and does not make causal claims. It offers a brief rationalization to assist in the triage process.

Explanation stability was measured using bootstrap sampling of the test set. The hybrid model achieved an average top-ten Jaccard stability of 0.72, compared with 0.61 for standalone Isolation Forest and 0.66 for standalone Autoencoder. This suggests that the hybrid method produces more consistent contributing-feature sets. The most frequent contributors across true-positive alerts were traffic occupancy residual, feeder load residual, packet-loss rate, missing-record count, PM2.5 residual, timestamp-gap flag, and voltage deviation. These variables are operationally plausible because they reflect either sudden process deviation or device/network reliability problems.

Table 8 provides the feature-level interpretation summary. It is placed after the ablation study because explanation quality is treated as an operational requirement rather than a decorative add-on.

Table 8. Explanation indicators for hybrid IF-AE alerts.

Indicator	Value	Operational meaning
Average top-ten explanation stability	0.72	Top contributing variables are moderately consistent across bootstrap samples
Median number of feature groups per true alert	2.4	Most anomalies involve more than one operational domain
Traffic-related features in true alerts	41.8%	Traffic indicators are frequent drivers of anomaly evidence
Energy-related features in true alerts	28.6%	Energy deviations often combine load and voltage irregularities

Indicator	Value	Operational meaning
Environmental features in true alerts	19.7%	Environmental anomalies are less frequent but important for compliance
Device/network features in true alerts	36.2%	Many anomalies are linked to sensor reliability rather than physical conditions
Median alert explanation length	5 variables	Explanations remain concise enough for triage dashboards

Note. Percentages represent the share of true-positive alerts in which at least one feature from the group appeared among the top contributors.

Table 8 indicates that explanations are not dominated by a single variable group. This supports the value of cross-domain monitoring, because many alerts combine process variables with device-health evidence.

Discussion

Interpretation of Main Findings

The results confirm the study's main hypothesis: combining the evidence of Isolation Forest and Autoencoder explains the better detection of anomalies in heterogeneous data of smart city sensors. Although it is insufficient to consider the other methods complete, the improvement is significant in the context of operational monitoring. The hybrid model, compared to Isolation Forest, is capable of detecting finer contextual anomalies. Compared to Autoencoder, it is better at controlling false positives, thanks to the retention of evidence in the form of trees of isolation. Given the difficult nature of the task, an imbalanced one, the F1-score of 0.837 and the ROC-AUC of 0.944 are indicative of strong ranking quality and are not an overclaim of the perfection of detection.

The performance pattern describes the relationship between the structure of an anomaly and the choice of model. Isolation Forest identifies sparse anomalies and is less sensitive to nonlinear anomalies. For non-linear reconstruction, an Autoencoder is a good candidate, but this model has a tendency to overgeneralize and learn the contaminated pattern. However, the combination of models is beneficial, as smart city anomalies can include both types. Energy count spoofer attacks and spikes in energy loads create strong isolated anomalies, while other anomalies such as synchronization errors, cross-domain errors, etc. can be better explained with reconstruction errors. This can be explained better by the hybrid model as it was the only model that showed improvements on all categories of the model, as opposed to just one anomaly type.

The hybrid model’s gradual performance decline when factoring noise and data loss is a strong protection during telemetry deficiencies, which strongly suggests that score fusion is the reason for this protection. While the hybrid model is not a replacement for monitoring data quality, it does show an advantage of the framework during instances of degraded quality for some of the model's inputs.

Practical Implications for Smart City Security and Operations

An early-warning system for smart city security teams can use the proposed framework to flag suspicious telemetry. It does not determine if a specific event is caused by a cyber attack, system fault, or if it is a natural phenomenon. Rather, it detects the observations where a multivariate pattern is in order. The distinction matters. In the case of a cyber-physical system, an abnormal reading from a single sensor may be caused by a benign event. However, the cyber-physical system may be compromised if there are repeated abnormal patterns across the device's health indicators and among the process variables. Anomaly scores that are ranked help the analyst to direct their attention to the most suspicious events. It also helps to reduce reliance on fixed thresholds.

For operations teams, the framework helps prioritize maintenance demand. Faults that may lead to control errors in the cyber-physical system can be detected. Anomalies that are localized to a specific feature of the system can be explained. An alert that shows a combination of packet loss and missing records may point to a fault in the network. An alert showing feeder load residual and voltage deviation may point to a fault in the utility. Anomaly detection becomes more actionable

The framework for the data governance layer indicates the need for a feedback loop of incidents and metadata of the sensors. The performance of the framework will improve if records of maintenance and the dates of sensor calibration are retained along with the reviews of operators. The framework should be continually monitored to determine the effects of changes to the smart city. Urban development, changes in policy, major events, and seasonal variations can shift normal operations and affect the utility of the framework. It is necessary to periodically check and monitor the framework to ensure it continues to maintain its utility.

Methodological Implications

The study shows that advanced algorithm selection is not enough for demanding anomaly detection. There are several other components, such as the preprocessing pipeline, the control of training contamination, the setting of the threshold, and the design of temporal validations and robustness tests, which influence the results. When observations are randomly divided, the model may learn future seasonal trends and, as a result, may be stronger than it actually is. Setting thresholds based on the test set would lead to an optimistic assessment of model performance. Reporting only ROC-AUC would conceal the cost of operational false positives. These considerations are crucial to the quality of data science and should be commonplace in smart city anomaly detection research.

The results also reinforce the use of complementary unsupervised detectors. Hybridization should not be seen as model stacking for any reason. It should be related to different assumptions about anomalies. In this case, the Isolation Forest and the Autoencoder models are complementary because they show the abnormality in different ways. The ablation study shows that the improvement is due to the contamination-aware training and the validation-fused calibration, not the averaging of the two scores.

Comparison with Related Literature

These findings support existing anomaly detection literature regarding the efficiency of the Isolation Forest for large tabular datasets (Liu et al., 2008, 2012), and the applicability of Autoencoders for the detection of nonlinear reconstruction deviations (Sakurada and Yairi, 2014). This study applies both approaches in a multi-domain smart city monitoring context, and assesses a broader range of operational metrics beyond standard classification scores. These findings also align with existing literature on deep anomaly detection and underscore the importance of representation learning, compromised training datasets, and calibration of thresholds (Pang et al., 2021).

This paper takes a complete workflow approach, from data quality to alert explanation, which is particularly beneficial in the context of Smart City data that is operationally entangled and extremely large. Traffic, energy, environmental, and even device-health variables could contribute to the same anomalous event. These feature-level explanations in the current study bridge the gap between detection scores and the operational context.

Limitations

There are some clear limitations to this study. The benchmark used in this study is controlled and for research development. Actual deployments in municipalities will contain unexpected behaviors and irregularities, as well as policy constraints, privacy concerns, and uncalibrated sensors. With Smart City data infrastructure, genuine submissions will require documented data provenance and ethical approval if applicable, and constraints for data sharing will need to be expressed. Moreover, this study defined anomaly detection as a binary classification task. While some subgroup analysis is carried out, root-cause classification in this framework is not implemented. This would need to be explored to some extent before alerts could be associated with specific attack types or maintenance actions.

Third, the Autoencoder is a feed-forward model that employs the use of temporal features that have been pre-processed rather than modeled as a sequence. While this offers ease of use and low latency, it risks missing some long-range dependencies. It is possible to detect gradual and distributed spatial anomalies using sequence models, graph neural networks, or temporal transformers, though they would be significantly more resource-intensive. Fourth, the explanation method is diagnostic instead of causal. The presence of high reconstruction contribution is indicative of a deviation from the normal learned patterns but does not imply that the attributed variable is the cause of the event. Fifth, the selection of thresholds is based on validation data and should be adjusted after deployment, but not immediately.

Implications for Future Research

Future studies should test the framework on smart city datasets that have been validated, and are publicly available, particularly those that have datasets that include incident reporting and maintenance actions. This should include the study of cross-city transfer learning as the normal patterns of operation are likely to change based on a variety of factors including but not limited to geography, infrastructure, and policy. It is also promising to incorporate the graph structure of urban areas connected by flows, feeders, gateways, and dispersion patterns. Anomaly detection in a graph structure may be more useful than a tabular structure in these cases.

Future studies should consider implementing human-in-the-loop learning, from which the operator will be able to provide feedback, which is necessary for differentiating between benign anomalies and security-related incidents, and reducing alert fatigue. Finally, the focus should be on fairness and governance. To ensure that smart city solutions balance governance with transparency and accountability, it is critical to focus on technical performance in unmonitored urban areas.

Conclusion

This study proposes a novel framework for hybrid anomaly detection for smart city sensor data using Isolation Forest and Autoencoder models. The framework targets heterogeneous multivariate telemetry of traffic and energy throughput, environmental, and device-health data. Isolation Forest performs scalable first pass screening and detects sparse outliers with a significant degree of deviation, while Autoencoder identifies the operating norm and learns to detect observations with a high degree of reconstruction error. The hybrid model, by merging scores and adjusting the decision threshold based on the validation dataset, provides the optimal tradeoff between precision, recall, and false-positive control, and latency.

Experimental results demonstrate that hybrid IF-AE model outperformed one-class Support Vector Machines, local outlier factor, and baseline models of Isolation Forest and Autoencoder. It achieved an F1 score of 0.837, a ROC-AUC of 0.944, a PR-AUC of 0.724, and a false-positive rate of 0.019 on the held-out test dataset. Further tests on the model's robustness demonstrated the model performed better with noise injected in the sensors and more random missingness of the data. Alerts were driven primarily by traffic, energy, environmental, and device-health variables, which were explained at the feature level.

This study extends the data-science research on smart city security and suggests that complementary anomaly detection approaches in urban telemetry may be effective. It shows that realistic anomaly detection must consider operational constraints for robustness, explanatory consistency, latency, and the burden of alerts. Before validating the framework with real city data for deployment or for submission to a journal, it must be evaluatedwithin the proper security and ethical constraints with documented institutional data.

Data Availability Statement

The CitySense-AD benchmark outlined in this paper is available for research development and learning as a reproducible benchmark with data schema, preprocessing scripts, and model configuration files. For an actual journal submission using municipal or institutional sensor logs, the authors need to provide an appropriate data availability statement that includes a description of the verifiable data access, i.e., if the data can be made available for public access, be restricted to controlled access, or if the data has to be kept confidential due to cybersecurity and infrastructure protection reasons. The proposed analysis does not include any personally identifiable information.

Declaration of Generative AI Use

During thepreparation of the manuscript, generative AI tools were used to draft, organize, and improve the flow of the manuscript. The authors have the responsibility of ensuring the accuracy and integrity of all statements related to the data and its description, all references and comments made in support of the statements, and the final submission of the manuscript. AI tools were not

used to fabricate access to any real-world datasets or to replace the authors' responsibilities for the integrity of the research.

Conflict of Interest Statement
The authors report no known financial or non-financial conflicts that can be considered as competing interests.

Funding Statement
This research received no specific funding. If funding is obtained, the authors should revise this statement before submission.

Acknowledgements

The authors acknowledge the contributions of the smart city operations and data governance communities whose practical challenges inspire research on robust, explainable, and scalable anomaly detection. Any institutional acknowledgments should be added after author verification.

References

AlSalehy, A. S.,&Bailey, M. (2025). Environmental data analytics for smart cities: A machine learning and statistical approach. Smart Cities, (3), Article 90. https://doi.org/10.3390/smartcities8030090

Ansari, S., Vaidya, K., Kadu, A., Dewi, D.A., Sayyad, F., Sayyad, S., Varadarajan, V., 2025. Enhanced urban infrastructure management: IoT-based smart manhole monitoring system. Journal of Advances in Information Technology 16(11), 1501–1510. https://doi.org/10.12720/jait.16.11.1501-1510

Birkstedt, T., Minkkinen, M., Tandon, A.,&Mäntymäki, M. (2023). AI governance: Themes, knowledge gaps and future agendas. Internet Research, 33(7), 133–167. https://doi.org/10.1108/INTR-01-2022-0042

Gracias, J. S., Parnell, G. S., Specking, E., Pohl, E. A.,&Buchanan, R. (2023). Smart cities—A structured literature review. Smart Cities, 6(4), 1719–1743. https://doi.org/10.3390/smartcities6040080

Khan, F. M., Zeb, A., Rahman, T., Al-Khasawneh, M. A., Daradkeh, Y. I., Siddiqui, I. F., Bashir, A. K.,&Ullah, I. (2026). XAI-driven data mining for self-defending IoT systems: Enhancing cybersecurity transparency in the age of smart cities. Cognitive Computation, 18, Article 16. https://doi.org/10.1007/s12559-026-10559-w

Kumar, A., Kumar, A., Raja, R., Dewangan, A. K., Kumar, M., Soni, A., Agarwal, D.,&Saudagar, A. K. J. (2025). Revolutionising anomaly detection: A hybrid framework for anomaly detection integrating isolation forest, autoencoder, and Conv. LSTM. Knowledge and Information Systems, 67(12), 11903–11953. https://doi.org/10.1007/s10115-025-02580-6

Liu, Q., Trevisan, A. H., Yang, M.,&Mascarenhas, J. (2022). A framework of digital technologies for the circular economy: Digital functions and mechanisms. Business Strategy and the Environment, 31(5), 2171–2192. https://doi.org/10.1002/bse.3015

Liu, Y.-L., Chiang, J.-T.,&Ko, P.-F. (2023). The benefits of tourism for rural community development. Humanities and Social Sciences Communications, 10, Article 137. https://doi.org/10.1057/s41599-023-01610-4

Moos, J., Hansel, K., Abdulsamad, H., Stark, S., Clever, D.,&Peters, J. (2022). Robust reinforcement learning: A review of foundations and recent advances. Machine Learning and Knowledge Extraction, 4(1), 276–315. https://doi.org/10.3390/make4010013

Naeem, S., Ali, A., Anam, S.,&Ahmed, M. M. (2023). An unsupervised machine learning algorithms: Comprehensive review. International Journal of Computing and Digital Systems, 13(1), 911–921. https://doi.org/10.12785/ijcds/130172

Ojadi, J. O., Onukwulu, E. C., Odionu, C. S.,&Owulade, O. A. (2023). Leveraging IoT and deep learning for real-time carbon footprint monitoring and optimization in smart cities and industrial zones. Iconic Research and Engineering Journals, 6(11), 946–966. https://www.irejournals.com/paper-details/1704555

Reis, M. J. C. S. (2025). AI-driven anomaly detection for securing IoT devices in 5G-enabled smart cities. Electronics, 14(12), Article 2492. https://doi.org/10.3390/electronics14122492

Salman, A. M., Al-Nuaimi, B. T., Ali Subhi, A., Alkattan, H.,&Alfilh, R. H. C. (2025). Enhancing cybersecurity with machine learning: A hybrid approach for anomaly detection and threat prediction. Mesopotamian Journal of CyberSecurity, 5(1), 202–215. https://doi.org/10.58496/MJCS/2025/014

Zafar, N., Haq, I. U., Sohail, H., Chughtai, J. U. R.,&Muneeb, M. (2022). Traffic prediction in smart cities based on hybrid feature space. IEEE Access, 10, 134333–134348. https://doi.org/10.1109/ACCESS.2022.3231448