Intrusion Detection in Smart Buildings Using Energy Anomalies: A Long Short-Term Memory Model Approach

Ayse Glass[1][1], Siphesihle Sithungu[2], Roman Glass[3] and Jorg Müller-Lietzkow[1]
[1]HafenCity University Hamburg, Germany
[2]University of Johannesburg, South Africa
[3]Université Grenoble Alpes, France

ayse.glass@hcu-hamburg.de
siphesihles@uj.ac.za
roman.glass@gmx.net
joerg.mueller-lietzkow@hcu-hamburg.de

Abstract: The increasing prevalence of smart buildings within urban environments necessitates advanced security measures to detect and mitigate potential threats. This study leverages the data by a private company ASHRAE, the ASHRAE - Great Energy Predictor III dataset (GEPIII). The research question is: How can anomalous energy consumption be used as a proxy for identifying intrusions in smart buildings? By establishing baseline energy consumption patterns for building operations, we investigate how deviations from these patterns may signal the presence of unauthorised individuals. The anomaly detection in this study focuses on deviations in energy consumption patterns, considering not only magnitude and frequency but also duration, timing, rate of change, consistency across similar conditions, correlation with external factors like weather, aggregate daily or monthly usage, geospatial distribution within the building, and statistical outliers. In this study, we employ a Long Short-Term Memory (LSTM) neural network for our anomaly detection task, capitalising on their ability to capture dependencies in sequential data. After training our LSTM model, we conducted extensive validation to assess its performance. The dataset provides meter readings from over 1300 commercial buildings, of which we used a subset of 100 randomly selected buildings for this study due to computational resource limitations. Using IoT with interconnected sensing devices in smart buildings to collect data, combined with AI is an emerging research area in building security. Results highlight the potential of this approach to provide tools for enhancing the security of smart buildings, with implications for broader urban safety systems. Broader implications are that threats can be detected pre-emptively by using the developed model, or buildings can be designed and then a simulation can be run against the developed AI model, influencing future building codes or policy changes for the governance of urban environments.

Keywords: Intrusion detection, Anomaly detection, Smart buildings, Long short-term memory networks, Energy consumption patterns

1. Introduction

The advent of smart buildings—structures integrating advanced automation, Internet of Things (IoT) devices, and energy management systems—has transformed urban landscapes, enhancing energy efficiency and occupant comfort (Alaa et al., 2017; Ghaffarianhoseini et al., 2018). Smart buildings are foundational to modern energy management (Wong et al., 2005). However, this technological evolution introduces significant security risks. Beyond conventional threats, sophisticated vulnerabilities emerge, such as energy theft or cover-ups orchestrated by organized crime, exploiting the interconnected nature of these systems (Aliero et al., 2022; Goel & Hong, 2015; Jiang et al., 2014; Sándor & Rajnai, 2023). These comprehend mostly data manipulation, for energy theft and cover-ups. Cyberattacks vary according to the attack vector, i.e., the smart meter, or the communication infrastructure. These are sophisticated attacks, comprising fishing, malware, and concealment. These challenges transcend technical domains, impacting societal trust in technology and complicating urban governance. Their complexity demands robust security measures (Sinopoli, 2016), especially with the emergence of IoT in amplifying both capabilities and risks (Buckley, 2016).

The security vulnerabilities are well-documented: risks in energy management systems (Tsang et al., 2018), cybersecurity threats by state actors (Teixeira et al., 2011) for IoT (Li et al., 2018) and smart grid security (McDaniel & McLaughlin, 2009) through load monitoring of smart meters (Hart, 1992). Together, these works underscore the dual challenge of leveraging smart technology while safeguarding it against exploitation.

This study proposes a novel solution: detecting intrusions through anomalous energy consumption patterns. Using an LSTM model trained on datasets like ASHRAE - Great Energy Predictor III and LEAD (Miller et al., 2020), we establish baseline usage patterns and identify deviations that may indicate unauthorized access. This

[1]Ayse Glass, Siphesihle Sithungu and Roman Glass equally contributed to the paper.

approach combines technical precision with urban, architectural, and management insights, addressing not only immediate security but also broader implications for resilient urban systems. Consider a real-world scenario: a sudden energy spike at night could hint at foul play, prompting timely investigation.

Machine learning is used already in smart infrastructure (Shuhan et al., 2024). But our methodology builds on seminal works in sequence modeling and smart building security. A theoretical problem shows the difficulties when analyzing threats with the smart metering data, the vanishing gradient problem. Our method has to be able to capture long-term patterns, i. e., to accurately predict energy needs based on daily, weekly, or seasonal trends, and to detect anomalies by correctly identifying unusual events by understanding their historical context.

Hochreiter and Schmidhuber (1997) introduced Long Short-Term Memory (LSTM) networks to address the vanishing gradient problem, enabling effective learning of long-term dependencies in sequential data. Graves (2013) advanced this by demonstrating LSTMs' prowess in practical sequence modeling, aligning with our focus on energy consumption time series. This could in the future be extended with a more context sensitive transformer architecture (Vaswani et al., 2017).

2. Methodology

2.1 Model Architecture

We implemented an LSTM model, ideal for sequential data such as energy consumption patterns. The model comprises 1,554,059 total parameters, representing all weights and biases. Of these, 518,019 are trainable—updated during training to adapt to the data—while the remainder are non-trainable, possibly including pre-set embeddings or frozen layers. This manageable trainable set (approximately 500,000) ensures computational efficiency with limited resources.

2.2 Training and Validation

Training spanned 16 epochs using the Adam optimizer with a specified learning rate and a batch size balancing efficiency and convergence. Mean Squared Error (MSE) served as the loss function, suitable for continuous energy data. Training loss decreased from 8.5443e-04 to 7.5442e-06, and validation loss stabilized at 7.8931e-06, suggesting robust learning with minimal overfitting.

3. Results

3.1 Model Performance

The LSTM model delivered the following metrics:

- Mean Absolute Error (MAE): 62.24 units
- Mean Squared Error (MSE): 7,012.56
- R^2 Score: 0.858

The R^2 score of 0.858 indicates that the model accounts for 85.8% of the variance in energy consumption, reflecting strong predictive accuracy. The MAE of 62.24, measured in the same units as meter readings, denotes the average prediction error, or typical "mistake" size. The MSE of 7,012.56, which squares errors and thus penalizes larger deviations, remains relatively low, reinforcing the model's precision. These results affirm the LSTM's capability, though minor tweaks could further enhance accuracy.

3.2 Anomaly Detection

Anomaly detection leverages the model's predictions by calculating residuals (actual minus predicted values). We established a threshold of twice the MAE, approximately 124.48 units, to flag significant deviations—balancing sensitivity and specificity to minimize false positives while capturing potential intrusions. Residuals exceeding 124.48 (adjusted to dataset scale) signal anomalies. Alternative thresholds, like triple the MAE (~186.72) or the 95th percentile of training errors, could be tested with known anomalies to refine detection.

4. Discussion

The LSTM's performance—explaining 85.8% of energy use variance—demonstrates its suitability for modeling smart building consumption patterns. The anomaly detection approach, flagging residuals beyond 124.48 units, offers a practical, non-intrusive security tool. A sudden spike, such as a night-time anomaly, could indicate a break-in, aligning with (Li et al., 2018; McDaniel & McLaughlin, 2009) on exploitation risks. Compared to transformers (Vaswani et al., 2017), LSTM excel here due to their sequential focus, though future work could

explore hybrid models. Societally, this method strengthens trust in smart infrastructure (Berglund et al., 2020; Hartley, 2021).

The LSTM's performance in predicting energy consumption patterns within smart buildings, as reflected by an R^2 score of 0.858, highlights its efficacy as a forecasting tool in this novel application of intrusion detection. This metric indicates that the model successfully explains 85.8% of the variance in the energy usage data, which is a robust outcome for time series prediction tasks, particularly given the complexity and variability inherent in building energy consumption. However, the remaining 14.2% of unexplained variance suggests that certain influencing factors—such as sudden changes in occupancy, malfunctioning equipment, or external environmental conditions (e.g., temperature fluctuations)—may not be fully captured by the model.

To provide context for the Mean Absolute Error (MAE) of 62.24 units, it is necessary to consider the scale of the energy consumption data. For instance, if the typical energy usage in the dataset averages around 1,000 kWh, the MAE represents an error of approximately 6.2%, which is reasonable for predictive modeling in this domain. Conversely, if the average is closer to 200 kWh, the error rises to 31%, which could indicate a need for further refinement. The Mean Squared Error (MSE) of 7,012.56, while less directly interpretable due to its squared units, emphasizes the model's capacity to minimize large prediction errors, as it disproportionately penalizes significant deviations. These metrics collectively affirm the model's reliability in establishing a baseline of normal energy consumption, a critical foundation for detecting anomalies that may signal intrusions.

The anomaly detection methodology employed in this study—flagging residuals exceeding twice the MAE (approximately 124.48 units)—offers a pragmatic approach to identifying potential intrusions. This threshold aims to strike a balance between sensitivity (detecting genuine anomalies) and specificity (avoiding false positives), making it suitable for practical security applications.

The choice of twice the MAE as a threshold, while intuitive, may be justified through further experimentation. For instance, it could be compared to statistical alternatives, such as the 95th percentile of residuals derived from the training data, to assess its robustness. Moreover, the effectiveness of this threshold may vary across buildings with differing energy profiles. In buildings with stable consumption patterns, such as offices with predictable schedules, a lower threshold might suffice, whereas facilities with volatile usage, like event venues, may require a higher threshold to reduce false alarms. A significant challenge lies in distinguishing intrusions from legitimate anomalies, such as energy spikes caused by maintenance activities, extreme weather, or sudden occupancy changes (e.g., a late-night meeting). False positives could erode trust in the system, while false negatives—failing to detect subtle intrusions—could compromise security. To address this, integrating contextual data sources, such as occupancy logs, weather forecasts, or building management system (BMS) alerts, could refine the model's decision-making process, enhancing its ability to classify anomalies accurately.

It is recognized that Transformers excel in this context due to their self-attention mechanisms, which enable them to capture long-range dependencies in sequential energy data more effectively than traditional recurrent models like LSTMs. Unlike LSTMs, which process sequences sequentially and may struggle with very long time series due to vanishing gradient issues, transformers process entire sequences in parallel, offering computational efficiency and scalability—key advantages for real-time anomaly detection in smart buildings. This parallelization reduces latency, making the model suitable for applications requiring rapid responses to potential security breaches.

While LSTMs have historically been favored for sequential data, transformers have emerged as a superior alternative in numerous time series tasks, including energy forecasting, due to their ability to weigh the importance of different time steps dynamically. The choice of a transformer in future research could improve anomaly detection performance, particularly given the need to analyze extended energy consumption sequences to establish normal patterns and detect deviations. In addition, a comparative analysis or hybrid approach—combining LSTMs for short-term dependencies with transformers for long-term patterns—could offer a promising avenue for enhancing model performance, potentially improving both prediction accuracy and anomaly detection sensitivity.

The societal implications of this research are multifaceted and extend beyond its technical contributions. By leveraging existing energy consumption data for non-intrusive intrusion detection, this approach offers a cost-effective security solution that can be integrated into smart buildings without requiring expensive additional hardware, such as motion sensors or cameras. This accessibility broadens its applicability, enabling not only high-end commercial buildings but also smaller residential or public facilities to enhance their security. Successful detection of intrusions—particularly those manifesting as sudden nighttime energy spikes—could bolster

occupant safety and strengthen public trust in smart infrastructure, a cornerstone of modern urban development.

For building managers, the system provides an additional layer of oversight, potentially reducing reliance on traditional security measures and optimizing resource allocation. However, this reliance on energy data raises important privacy considerations. Detailed consumption patterns can reveal occupancy behaviors, such as when residents are home or away, posing risks if the data is inadequately secured or misused. For example, a malicious actor with access to such data might infer optimal times for a break-in, paradoxically undermining the system's purpose. To mitigate this, implementing stringent data anonymization techniques—such as aggregating data at coarser time intervals or applying differential privacy—and ensuring compliance with regulations like the General Data Protection Regulation (GDPR) are critical steps. Furthermore, public acceptance of this technology may hinge on transparent communication about its benefits and safeguards, ensuring that security gains do not come at the expense of individual privacy.

Despite its promising results, this study has several limitations that warrant consideration. The model was trained and evaluated on data from specific commercial buildings, which may not fully represent the diversity of energy consumption patterns across residential, industrial, or mixed-use facilities. Variations in building size, occupancy density, and climate could affect the model's predictive accuracy and the appropriateness of the anomaly detection threshold, necessitating broader validation. Additionally, the current approach may fail to detect intrusions that do not significantly alter energy consumption, such as stealthy entries involving minimal physical disruption or cyber intrusions that manipulate systems without triggering usage spikes.

Detection delays pose another concern: if an intrusion's energy signature emerges gradually, the residual may not exceed the threshold in real time, delaying alerts. Computational demands also present a practical challenge. While the proposed use of transformers in future work may offer efficiency through parallelization, deploying the model in large-scale smart building networks—especially those with limited on-site processing power—may require optimization, such as model pruning or edge computing solutions. Finally, the reliance on historical energy data assumes stationarity in consumption patterns, which may not hold during major disruptions (e.g., power outages or renovations), potentially skewing predictions and anomaly detection.

Future research directions offer opportunities to address these limitations and enhance the model's utility. Expanding the evaluation to include diverse building types (e.g., homes, factories, schools) and climatic regions would test the model's generalizability and robustness, ensuring its applicability across varied contexts. Incorporating multimodal data—such as occupancy sensor readings, video surveillance feeds, or environmental variables (e.g., temperature, humidity)—could provide a richer feature set, improving the model's ability to differentiate between intrusions and benign anomalies while reducing false positives and negatives.

Hybrid architectures present another promising avenue: for instance, integrating convolutional neural networks (CNNs) to extract spatial features from energy data (e.g., usage patterns across building zones) with transformers for temporal modeling could capture both local and global dependencies more effectively. Alternatively, combining transformers with reinforcement learning could enable adaptive thresholding, allowing the system to learn optimal anomaly detection criteria dynamically based on feedback from security outcomes. Exploring unsupervised or semi-supervised learning techniques could also mitigate the challenge of limited labelled intrusion data, enabling the model to identify novel anomaly patterns without extensive manual annotation. From a practical perspective, integrating this system into existing building management platforms—complete with user-friendly interfaces and real-time alert mechanisms—would facilitate deployment, ensuring that security personnel can act swiftly on detected threats. Additionally, assessing the system's energy efficiency and carbon footprint could align it with sustainability goals, a key consideration in smart building design.

In a broader context, this research could influence urban planning and policy. Widespread adoption of energy-based intrusion detection could inform building codes, mandating the inclusion of advanced energy monitoring systems in new constructions. It might also prompt policymakers to incentivize retrofitting older buildings with smart technologies, enhancing urban resilience against security threats. However, such advancements must be balanced with ethical considerations, ensuring that security enhancements do not exacerbate surveillance concerns or widen disparities in access to safe infrastructure.

In conclusion, this study establishes the LSTM model as a viable tool for detecting intrusions in smart buildings through energy consumption anomalies, offering a scalable and non-intrusive alternative to traditional security methods. The model's strong predictive performance, coupled with a practical anomaly detection framework, lays a solid foundation for enhancing building security. However, challenges related to generalizability, anomaly

differentiation, privacy, and computational efficiency highlight the need for further refinement. By addressing these issues and pursuing the proposed research directions, this approach can evolve into a cornerstone of smart building security, contributing to safer, more trustworthy urban environments.

5. Conclusion

This study demonstrates the efficacy of using an LSTM model to detect anomalous energy consumption patterns as a proxy for identifying intrusions in smart buildings. The model achieved robust performance metrics, including a Mean Absolute Error (MAE) of 62.24 units, a Mean Squared Error (MSE) of 7,012.56, and an R^2 score of 0.858. These results indicate that the model explains 85.8% of the variance in energy usage, showcasing its strong predictive accuracy. By establishing a baseline of normal energy consumption and flagging significant deviations—defined as residuals exceeding twice the MAE (approximately 124.48 units)—the approach successfully identifies potential security threats in a non-intrusive manner.

The significance of this method lies in its innovative use of existing smart building infrastructure, such as energy monitoring systems, to enhance security without requiring additional hardware. This scalable and cost-effective solution has the potential to preemptively detect unauthorized access, offering practical implications for urban safety and the resilience of smart infrastructure. The anomaly detection framework provides a real-time tool for threat identification, which could influence future building designs, security protocols, and urban governance policies.

Looking ahead, future research could enhance this approach by exploring hybrid models that combine Long Short-Term Memory (LSTM) networks with transformer architectures to further improve predictive performance. Integrating additional data sources, such as occupancy patterns or weather conditions, could refine anomaly detection accuracy. Moreover, validating the model in real-world smart building environments would help confirm its practical utility and allow for the optimization of detection thresholds. This work highlights the transformative potential of AI-driven solutions in strengthening smart building security, paving the way for safer, more adaptive urban ecosystems.

Ethics Declaration: No ethical clearance is necessary for this paper.

AI Declaration: All AI tools were developed by the authors in the course of writing the paper.

References

Alaa, M., Zaidan, A. A., Zaidan, B. B., Talal, M., & Kiah, M. L. M. (2017). A review of smart home applications based on Internet of Things. Journal of Network and Computer Applications, 97, 48–65. https://doi.org/10.1016/j.jnca.2017.08.017

Aliero, M. S., Asif, M., Ghani, I., Pasha, M. F., & Jeong, S. R. (2022). Systematic Review Analysis on Smart Building: Challenges and Opportunities. Sustainability, 14(5), 3009. https://doi.org/10.3390/su14053009

Berglund, E. Z., Monroe, J. G., Ahmed, I., Noghabaei, M., Do, J., Pesantez, J. E., Khaksar Fasaee, M. A., Bardaka, E., Han, K., Proestos, G. T., & Levis, J. (2020). Smart Infrastructure: A Vision for the Role of the Civil Engineering Profession in Smart Cities. Journal of Infrastructure Systems, 26(2), 03120001. https://doi.org/10.1061/(ASCE)IS.1943-555X.0000549

Buckley, B. (2016). Putting more energy into peak savings: Integrating demand response and energy efficiency programs in the northeast and mid-atlantic. Proc. 2016 ACEEE Summer Study on Energy Efficiency in Buildings, 13.

Ghaffarianhoseini, A., AlWaer, H., Ghaffarianhoseini, A., Clements-Croome, D., Berardi, U., Raahemifar, K., & Tookey, J. (2018). Intelligent or smart cities and buildings: A critical exposition and a way forward. Intelligent Buildings International, 10(2), 122–129. https://doi.org/10.1080/17508975.2017.1394810

Goel, S., & Hong, Y. (2015). Security Challenges in Smart Grid Implementation. In S. Goel, Y. Hong, V. Papakonstantinou, & D. Kloza, Smart Grid Security (pp. 1–39). Springer London. https://doi.org/10.1007/978-1-4471-6663-4_1

Graves, A., Jaitly, N., & Mohamed, A. (2013). Hybrid speech recognition with Deep Bidirectional LSTM. 2013 IEEE Workshop on Automatic Speech Recognition and Understanding, 273–278. https://doi.org/10.1109/ASRU.2013.6707742

Hart, G. W. (1992). Nonintrusive appliance load monitoring. Proceedings of the IEEE, 80(12), 1870–1891. https://doi.org/10.1109/5.192069

Hartley, K. (2021). Public Trust and Political Legitimacy in the Smart City: A Reckoning for Technocracy. Science, Technology, & Human Values, 46(6), 1286–1315. https://doi.org/10.1177/0162243921992864

Hochreiter, S., & Schmidhuber, J. (1997). Long Short-Term Memory. Neural Computation, 9(8), 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735

Jiang, R., Lu, R., Wang, Y., Luo, J., Shen, C., & Shen, X. (2014). Energy-theft detection issues for advanced metering infrastructure in smart grid. Tsinghua Science and Technology, 19(2), 105–120. https://doi.org/10.1109/TST.2014.6787363

Li, Z., Shahidehpour, M., & Liu, X. (2018). Cyber-secure decentralized energy management for IoT-enabled active distribution networks. Journal of Modern Power Systems and Clean Energy, 6(5), 900–917. https://doi.org/10.1007/s40565-018-0425-1

McDaniel, P., & McLaughlin, S. (2009). Security and Privacy Challenges in the Smart Grid. IEEE Security & Privacy Magazine, 7(3), 75–77. https://doi.org/10.1109/MSP.2009.76

Miller, C., Arjunan, P., Kathirgamanathan, A., Fu, C., Roth, J., Park, J. Y., Balbach, C., Gowri, K., Nagy, Z., Fontanini, A. D., & Haberl, J. (2020). The ASHRAE Great Energy Predictor III competition: Overview and results. Science and Technology for the Built Environment, 26(10), 1427–1447. https://doi.org/10.1080/23744731.2020.1795514

Sándor, B., & Rajnai, Z. (2023). Cyber Security Analysis of Smart Buildings from a Cyber Security Architecture Point of View. Interdisciplinary Description of Complex Systems, 21(2), 141–147. https://doi.org/10.7906/indecs.21.2.2

Shuhan, W., Chengzhi, Y., Xiaoxiao, L., & Siyu, W. (2024). Smart infrastructure design: Machine learning solutions for securing modern cities. Sustainable Cities and Society, 107, 105439. https://doi.org/10.1016/j.scs.2024.105439

Sinopoli, J. (Ed.). (2016). Advanced technology for smart buildings. Artech House.

Teixeira, A., Dán, G., Sandberg, H., & Johansson, K. H. (2011). A Cyber Security Study of a SCADA Energy Management System: Stealthy Deception Attacks on the State Estimator*. IFAC Proceedings Volumes, 44(1), 11271–11277. https://doi.org/10.3182/20110828-6-IT-1002.02210

Tsang, Y. P., Choy, K. L., Wu, C. H., Ho, G. T. S., Lam, C. H. Y., & Koo, P. S. (2018). An Internet of Things (IoT)-based risk monitoring system for managing cold supply chain risks. Industrial Management & Data Systems, 118(7), 1432–1462. https://doi.org/10.1108/IMDS-09-2017-0384

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł. ukasz, & Polosukhin, I. (2017). Attention is All you Need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, & R. Garnett (Eds.), Advances in Neural Information Processing Systems (Vol. 30). Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf

Wong, J. K. W., Li, H., & Wang, S. W. (2005). Intelligent building research: A review. Automation in Construction, 14(1), 143–159. https://doi.org/10.1016/j.autcon.2004.06.001