

Context-Aware Deep Learning Frameworks For Trajectory And Video-Based Anomaly Detection In Smart Urban Systems

Dr. Elena Váradi

Department of Computational Systems Engineering, Eötvös Loránd University, Hungary

Article Received: 05/11/2025, Article Revised: 25/11/2025, Article Accepted: 10/12/2025, Article Published: 01/01/2026

© 2026 Authors retain the copyright of their manuscripts, and all Open Access articles are disseminated under the terms of the [Creative Commons Attribution License 4.0 \(CC-BY\)](https://creativecommons.org/licenses/by/4.0/), which licenses unrestricted use, distribution, and reproduction in any medium, provided that the original work is appropriately cited.

ABSTRACT

Anomaly detection has emerged as a foundational capability for intelligent urban systems, surveillance infrastructures, and complex cyber-physical environments. As cities evolve into highly instrumented, data-intensive ecosystems, the ability to identify abnormal patterns in human movement, vehicular trajectories, and video streams has become essential for ensuring safety, efficiency, and resilience. This research article presents an extensive, theory-driven investigation into contemporary deep learning-based anomaly detection methods, with a particular emphasis on trajectory analysis and video surveillance in smart city contexts. Grounded strictly in the existing scholarly literature, the study synthesizes autoencoder-based models, variational and distributional approaches, attention-driven sequence modeling, and contextual anomaly detection frameworks. Rather than offering a superficial survey, the article develops a unified conceptual narrative that explains why these methods work, how they differ philosophically and technically, and what their implications are for real-world deployment. The methodology section elaborates on representation learning, temporal dependency modeling, and context discovery without relying on mathematical formalism, ensuring conceptual clarity. The results section interprets reported findings across studies in a descriptive and comparative manner, highlighting strengths, weaknesses, and performance trends. The discussion critically examines limitations related to data bias, interpretability, scalability, and ethical deployment, while also identifying promising research directions such as multimodal fusion and adaptive context modeling. The article concludes by arguing that anomaly detection should be understood not merely as a technical task but as a socio-technical capability central to the future of smart cities.

KEYWORDS

Anomaly detection, trajectory analysis, video surveillance, deep learning, smart cities, contextual modeling.

INTRODUCTION

The rapid proliferation of sensing technologies, networked devices, and data-driven infrastructures has transformed contemporary cities into complex, adaptive systems often described under the umbrella term “smart cities.” These environments continuously generate vast streams of heterogeneous data, ranging from video feeds and pedestrian trajectories to vehicular movement logs and maritime navigation paths. Within such data-rich contexts, anomaly detection has emerged as a critical analytical function, enabling systems to identify unusual, suspicious, or potentially harmful behaviors that deviate from established norms (Chandola et al., 2009). The importance of anomaly detection extends beyond security applications; it encompasses traffic

management, public safety, infrastructure monitoring, and urban planning (Batty et al., 2012; Giffinger et al., 2007).

Historically, anomaly detection was approached through statistical methods that relied on predefined thresholds, parametric assumptions, or simple distance-based metrics. While these methods provided early insights, they struggled to cope with the high dimensionality, non-linearity, and contextual variability inherent in modern urban data streams (Iglewicz & Hoaglin, 1993). As a result, false positives and false negatives were common, limiting practical applicability. The emergence of machine learning, and particularly deep learning, has significantly reshaped the anomaly detection landscape

by enabling systems to learn complex representations of normal behavior directly from data (Chandola et al., 2009).

Trajectory-based anomaly detection has become a focal area within this broader field. Trajectories encode rich spatio-temporal information about the movement of entities such as pedestrians, vehicles, ships, or aircraft. Deviations in trajectory patterns may signal dangerous driving, suspicious pedestrian behavior, maritime security threats, or system malfunctions (Shi et al., 2021; Xie et al., 2024). Unlike static data points, trajectories are inherently sequential and context-dependent, meaning that anomalies cannot be identified solely by examining isolated positions. Instead, understanding temporal evolution, interaction with the environment, and adherence to implicit social norms is essential (Kim et al., 2002).

Parallel to trajectory analysis, video-based anomaly detection has gained prominence due to the widespread deployment of surveillance cameras in urban spaces. Video data provides dense visual information but also poses substantial challenges related to computational complexity, variability in lighting and viewpoints, and the scarcity of labeled anomalous events (Berroukham et al., 2023). Consequently, unsupervised and self-supervised learning paradigms have become dominant, with models trained primarily on normal behavior and tasked with identifying deviations during inference (Huang et al., 2023).

Despite significant progress, the literature reveals persistent gaps. Many studies focus narrowly on specific data modalities or environments, limiting generalizability. Others achieve high detection accuracy but offer limited interpretability, raising concerns about trust and accountability in safety-critical applications. Furthermore, contextual factors such as time of day, weather, cultural norms, or infrastructure layout are often underexplored, even though they fundamentally shape what constitutes “normal” behavior (Xu et al., 2020; Thorne, 2025).

This article addresses these gaps by providing a comprehensive, theory-rich synthesis of deep learning-based anomaly detection methods for trajectories and video data. Rather than proposing a new algorithm, the study aims to consolidate knowledge, articulate underlying principles, and critically examine the implications of existing approaches. By situating technical methods within the broader smart city and context-aware computing discourse, the article contributes a holistic perspective that is valuable for researchers, system designers, and policymakers alike.

METHODOLOGY

The methodological foundation of this article is

analytical and integrative rather than experimental. The study adopts a qualitative synthesis approach, drawing exclusively on peer-reviewed literature to construct a coherent framework for understanding anomaly detection in trajectory and video data. This approach is particularly suitable given the diversity of models, datasets, and evaluation protocols reported across studies.

At the core of many contemporary methods lies representation learning, a paradigm in which models automatically discover meaningful features from raw data. Autoencoder-based architectures exemplify this approach. An autoencoder consists of an encoder that compresses input data into a latent representation and a decoder that reconstructs the original input. When trained exclusively on normal behavior, the model learns a compact representation that captures typical patterns. Anomalies are detected when reconstruction errors exceed expected levels, indicating that the input deviates from learned norms (Shi et al., 2021). This principle has been applied successfully to both trajectory data and video frames, demonstrating the versatility of autoencoders as unsupervised anomaly detectors.

Variational autoencoders extend this idea by imposing probabilistic structure on the latent space. Instead of learning deterministic encodings, variational models learn distributions that characterize normal behavior variability. This probabilistic framing enables more nuanced anomaly scoring, as deviations can be assessed relative to learned distributions rather than fixed thresholds (Zhang et al., 2023). In trajectory analysis, this allows models to distinguish between rare but acceptable behaviors and truly anomalous movements, a distinction that is critical in dynamic urban environments.

Sequence modeling represents another methodological pillar. Trajectories and videos are inherently temporal, meaning that anomalies may only become apparent when considering long-term dependencies. Attention-based sequence models address this challenge by dynamically weighting different temporal segments based on their relevance to prediction or reconstruction tasks (Wang et al., 2023). By focusing computational resources on salient moments, attention mechanisms enhance the model’s ability to capture subtle deviations that might be overlooked by uniform processing.

A complementary methodological perspective is offered by distributional approaches to trajectory similarity. Instead of relying on pointwise comparisons or handcrafted distance metrics, distributional methods model trajectories as probabilistic entities whose similarities are assessed through principled statistical comparisons (Wang et al., 2024). This approach aligns with the broader trend toward probabilistic reasoning in anomaly detection, emphasizing robustness and interpretability.

Video-based anomaly detection methodologies often integrate motion analysis through optical flow, which captures pixel-level movement between frames. By reconstructing optical flow patterns or predicting future frames, models learn the dynamics of normal scenes. Anomalies manifest as prediction failures or reconstruction inconsistencies, reflecting unexpected motion or appearance changes (Huang et al., 2023). This strategy is particularly effective in crowded scenes where individual object tracking is unreliable.

Contextual anomaly detection constitutes a critical methodological extension. Context-aware models explicitly incorporate auxiliary information such as spatial zones, temporal cycles, or environmental conditions. By conditioning anomaly detection on context, these models reduce false alarms and improve adaptability (Xu et al., 2020; Thorne, 2025). For example, a pedestrian running may be anomalous in a residential street at night but normal near a sports facility during the day.

Importantly, the methodologies reviewed emphasize unsupervised or weakly supervised learning. This reflects practical constraints, as anomalous events are rare, diverse, and difficult to label exhaustively. By learning from normal data, these systems align with real-world deployment scenarios where labeled anomalies are scarce or unavailable (Berroukham et al., 2023).

RESULTS

The findings reported across the reviewed literature collectively demonstrate that deep learning-based anomaly detection methods significantly outperform traditional statistical and rule-based approaches in complex, high-dimensional settings. Autoencoder-based models consistently show strong baseline performance, particularly in environments with relatively stable patterns of normal behavior (Shi et al., 2021). Their ability to compress and reconstruct data enables effective detection of gross deviations, making them suitable for initial deployment in surveillance systems.

Variational autoencoders introduce additional robustness by modeling uncertainty. Studies applying variational approaches to trajectory data report improved discrimination between normal variability and true anomalies, especially in scenarios with heterogeneous movement patterns (Zhang et al., 2023). This suggests that probabilistic latent representations are better aligned with the stochastic nature of human and vehicular behavior.

Attention-based sequence models demonstrate superior performance in online and real-time settings. By selectively focusing on informative temporal segments, these models achieve higher sensitivity to subtle anomalies without incurring prohibitive computational

costs (Wang et al., 2023). This is particularly relevant for applications such as traffic monitoring, where timely detection is critical.

Distributional trajectory similarity methods offer a different kind of result: enhanced interpretability and theoretical grounding. By framing anomaly detection as a problem of statistical divergence, these methods provide clearer explanations for why a given trajectory is considered anomalous (Wang et al., 2024). While computationally more demanding, they offer valuable insights for domains where accountability and transparency are paramount.

In video anomaly detection, frameworks based on optical flow reconstruction and frame prediction consistently report high detection accuracy in benchmark datasets. These methods are especially effective in capturing motion-related anomalies, such as sudden running or unusual crowd dispersion (Huang et al., 2023). However, their performance may degrade in scenes with significant visual clutter or camera motion.

Context-aware approaches show a marked reduction in false positives across multiple studies. By integrating contextual cues, these models better align system outputs with human expectations, enhancing practical usability (Xu et al., 2020; Thorne, 2025). This result underscores the importance of moving beyond purely data-driven definitions of normality toward richer, semantically informed models.

DISCUSSION

The collective evidence reviewed in this article highlights both the promise and the complexity of anomaly detection in smart urban systems. Deep learning has undeniably expanded the representational capacity of detection models, enabling them to capture intricate spatio-temporal patterns that were previously inaccessible. However, this increased power also introduces new challenges related to interpretability, data dependency, and ethical deployment.

One critical issue is the reliance on historical data to define normal behavior. Urban environments are dynamic, shaped by evolving social norms, infrastructure changes, and external events. Models trained on past data may struggle to adapt to new conditions, leading to concept drift and degraded performance. Context-aware and adaptive learning mechanisms offer partial solutions, but their design and validation remain open research questions (Xu et al., 2020).

Another concern relates to bias and fairness. If training data reflects existing inequalities or surveillance priorities, anomaly detection systems may disproportionately target certain populations or behaviors. This risk is particularly salient in video

surveillance, where cultural and contextual nuances are difficult to encode explicitly (Berroukham et al., 2023). Addressing these issues requires interdisciplinary collaboration and transparent evaluation practices.

Scalability is also a practical limitation. While deep models perform well in controlled settings, deploying them across city-scale infrastructures demands careful consideration of computational resources, latency constraints, and system integration. Lightweight models and edge computing strategies may mitigate some of these challenges, but trade-offs between accuracy and efficiency are inevitable.

Looking forward, multimodal anomaly detection represents a promising direction. By integrating trajectories, video, sensor readings, and contextual metadata, future systems could achieve more holistic situational awareness (Zhang et al., 2021). Such integration, however, raises new questions about data fusion, synchronization, and privacy protection.

CONCLUSION

Anomaly detection stands at the intersection of data science, urban studies, and social governance. The literature reviewed in this article demonstrates that deep learning-based methods, particularly those grounded in representation learning, sequence modeling, and contextual awareness, have significantly advanced the state of the art. Yet, these technical achievements must be understood within a broader socio-technical framework that acknowledges dynamic contexts, ethical considerations, and real-world constraints.

By synthesizing diverse approaches to trajectory and video-based anomaly detection, this article contributes a comprehensive perspective that emphasizes understanding over optimization. The future of smart cities depends not only on detecting anomalies accurately but on interpreting them wisely and responding responsibly.

REFERENCES

1. Batty, M., Axhausen, K. W., Giannotti, F., et al. (2012). Smart cities of the future. *The European Physical Journal Special Topics*, 214(1), 481–518.
2. Berroukham, A., Housni, K., Lahraichi, M., & Boulfrifi, I. (2023). Deep learning-based methods for anomaly detection in video surveillance: A review. *Bulletin of Electrical Engineering and Informatics*, 12, 314–327.
3. Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys*, 41(3), 1–58.
4. Giffinger, R., Fertner, C., Kramar, H., et al. (2007). *Smart cities—Ranking of European medium-sized cities*. Vienna University of Technology.
5. Huang, H., Zhao, B., Gao, F., Chen, P., Wang, J., & Hussain, A. (2023). A novel unsupervised video anomaly detection framework based on optical flow reconstruction and erased frame prediction. *Sensors*, 23, 4828.
6. Iglewicz, B., & Hoaglin, D. C. (1993). *How to detect and handle outliers*. Sage Publications.
7. Kim, K. Y., Kim, H. C., & Oh, S. H. (2002). Characteristic analysis of pedestrian behavior on local street in residential area. *Journal of the Korean Society of Civil Engineers D*, 22, 197–205.
8. Shi, H., Xu, X., Fan, Y., Zhang, C., & Peng, Y. (2021). An auto encoder network based method for abnormal behavior detection. In *Proceedings of the International Conference on Software Engineering and Information Management*, 243–251.
9. Thorne, E. V. (2025). Enhancing anomaly detection in complex systems through context discovery. *The Pinnacle Research Journal of Scientific and Management Sciences*, 2(10), 1–8.
10. Wang, C., Li, K., & Chen, L. (2023). Deep unified attention-based sequence modeling for online anomalous trajectory detection. *Future Generation Computer Systems*, 144, 1–11.
11. Wang, Y., Wang, Z., Ting, K. M., & Shang, Y. (2024). A principled distributional approach to trajectory similarity measurement and its application to anomaly detection. *Journal of Artificial Intelligence Research*, 79, 865–893.
12. Xie, Z., Bai, X., Xu, X., & Xiao, Y. (2024). An anomaly detection method based on ship behavior trajectory. *Ocean Engineering*, 293, 116640.
13. Xu, M., Zhang, J., & Liu, X. (2020). Contextual anomaly detection in real-time data streams for smart cities. *IEEE Transactions on Industrial Informatics*, 16(6), 4087–4096.
14. Zhang, L., Lu, W., Xue, F., & Chang, Y. (2023). A trajectory outlier detection method based on variational auto-encoder. *Mathematical Biosciences and Engineering*, 20, 15075–15093.
15. Zhang, Y., Li, Q., & Zhang, S. (2021). Multi-modal machine learning for anomaly detection in smart cities: A survey. *Future Generation Computer Systems*, 115, 265–279.