

Mobile Traffic Data as a Proxy for Urban Mobility: a Preliminary Study in Paris

Debora Giovannelli^{1,†}, Fabio Pinelli^{2,3,*} and Chiara Pugliese^{4,†}

¹University of Florence, Italy

²IMT School for Advanced Studies Lucca, Italy

³ISTI-CNR, Pisa, Italy

⁴IIT-CNR, Pisa, Italy

Abstract

As 4G and 5G networks become the backbone of urban infrastructure, the granular mobile traffic data they generate offers a high-resolution lens into collective human mobility dynamics. This paper investigates the hypothesis that service-specific mobile traffic can serve as a reliable proxy for human mobility and urban behavior. We model the problem as a Time Series Anomaly Detection (TSAD) task by proposing an unsupervised framework that employs a reconstruction-based representation model to establish a baseline urban activity. Our pipeline identifies significant deviations in social media app usage (e.g., Twitter, Instagram, YouTube) as spatio-temporal anomalies. These anomalies are subsequently clustered to filter noise and reveal coherent patterns corresponding to real-world events. We validate our framework using the NetMob 2023 dataset across the Paris metropolitan area. Our preliminary results demonstrate that this approach effectively captures a diverse spectrum of mobility events, including political protests, sporting matches, and disasters. These findings suggest that service-specific traffic data is a powerful, non-intrusive indicator of urban mobility flows and social disruptions.

Keywords

Urban Event Detection, Time Series Anomaly Detection, Human Mobility Modeling

1. Introduction

The evolution of 4G and 5G networks has fundamentally transformed the digital footprint of urban populations. As mobile applications become deeply embedded in daily routines, they generate massive volumes of service-specific mobile traffic data, i.e., granular records of upload and download activities across gaming, social networking, and messaging platforms. While mobile operators primarily collect this data through passive measurement probes for network optimization, these traces offer a high-resolution, implicit lens through which to observe human dynamics. This paper presents a preliminary study examining the hypothesis that service-specific mobile traffic can serve as a reliable proxy for human mobility and collective behavior within urban spaces. Traditionally, urban event detection has been treated as a search for unusual patterns in city life, often relying on Call Detail Records (CDRs) [1], taxi trajectories, or social media metadata [2]. While effective, these sources often suffer from coarse spatial resolution or sparsity. We argue that these urban disruptions can be effectively captured by modeling human mobility as a **Time Series Anomaly Detection** (TSAD) task. Recent shifts toward network signaling data [3] suggest that app-level consumption patterns may provide a more immediate and nuanced reflection of urban events. Because specific human activities, such as attending a protest or a stadium-sized sporting event, are now linked to the use of certain applications, deviations from "normal" traffic patterns can be identified as spatio-temporal anomalies that signal not only that an event is occurring, but also its nature, location, and social intensity.

In this work, we propose an unsupervised framework designed to demonstrate the feasibility of using big data, such as app-specific traffic, as a proxy for urban mobility through this TSAD lens. Our pipeline

Published in the Proceedings of the Workshops of the EDBT/ICDT 2026 Joint Conference (March 24-27, 2026), Tampere, Finland

*Corresponding author.

†These authors contributed equally.

✉ debora.giovannelli@edu.unifi.it (D. Giovannelli); fabio.pinelli@imtlucca.it (F. Pinelli); chiara.pugliese@iit.cnr.it (C. Pugliese)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

identifies spatio-temporal anomalies in social media app usage, which we represent as geolocated multivariate time series. The methodology employs a reconstruction-based representation step to learn baseline urban behaviors. By calculating reconstruction error, the model identifies significant deviations from expected activity as potential anomalies. These detected points are then clustered spatially to filter out localized noise and reveal coherent patterns that correspond to meaningful events, i.e., human movements. To evaluate our approach, we apply the framework to real-world data from the NetMob 2023 dataset [4], focusing on social network application traffic (e.g., Twitter, Instagram, YouTube, and Facebook) across the Paris metropolitan area. Our preliminary results demonstrate that this pipeline successfully captures a diverse spectrum of urban events, ranging from political protests and disasters to sports matches, by exploiting mobile traffic data and without the need for labeled ground truth. The alignment between our detected anomalies and known real-world occurrences suggests that the intuition that service-specific traffic is a powerful indicator of urban mobility flow and social disruptions is well-founded.

The paper is organized as follows. Section 2 reviews the related work. Section 3 introduces our methodology. Section 4 describes the experimental setup and dataset and presents the results. Finally, Section 5 concludes the paper and discusses limitations and future directions.

2. Related work

As introduced in the previous section, in order to demonstrate that mobile traffic data can serve as a proxy for human mobility, we identify urban events by formulating the problem as a TSAD task. TSAD plays a central role in domains such as finance, healthcare, and cybersecurity, where identifying anomalous temporal patterns is critical to anticipating failures or responding to disruptive events. In the context of mobility data science, this modeling allows us to interpret significant deviations in time series (such as GPS trajectories or CDR records) as a signature of human behavior. In recent years, deep learning models have received considerable attention in TSAD, as they can capture complex temporal dependencies and subtle deviations that traditional methods often fail to detect. Approaches span from predictive modeling, where the model forecasts the expected sequence and anomalies are identified as deviations, to reconstruction-based strategies, where models attempt to reconstruct normal patterns and anomalies emerge from reconstruction errors [5, 6]. Among these, autoencoder-based methods represent a widely adopted paradigm. Convolutional and recurrent autoencoders, as well as probabilistic extensions such as Variational Autoencoders (VAEs), have demonstrated strong performance by exploiting reconstruction error as an anomaly score [7, 8]. Furthermore, Graph Neural Networks (GNNs) and their variants (e.g., GCNs and GATs) have been utilized to model dependencies among multivariate time series as graph structures, capturing relational inductive biases over nodes and their interactions [9, 10, 11]. More recently, transformer-based approaches have emerged as powerful alternatives thanks to their self-attention mechanism, which enables the modeling of long-range temporal dependencies. Models such as Anomaly-Transformer [12] and TranAD [13] leverage association discrepancies and attention-based reconstructions to achieve high scalability and robustness in multivariate settings [14, 15, 16, 17, 18]. Despite these advancements, only a limited number of studies have considered service-specific mobile traffic as a primary source for observing urban mobility and detecting events. This scarcity is largely due to the challenges associated with such data, which are often difficult to access, inherently noisy, and characterized by massive volumes. Early frameworks, such as that by Pinelli et al. [19], identified unusual crowd events from call detail records (CDRs), proving the potential of mobile traces despite their coarse resolution. More recent work has shifted toward network signaling data [3], lightweight real-time anomaly detection in phone traffic [20], or specific event analysis like the Notre-Dame fire [21]. Collectively, these works illustrate the growing interest in using digital footprints to capture complex urban dynamics.

Our work builds upon these foundations but distinguishes itself by leveraging service-specific mobile traffic data to provide a richer, finer-grained view of urban dynamics than traditional CDRs or signaling-only datasets. By adopting a transformer-based autoencoder, we go beyond previous statistical models to

capture both temporal dependencies and inter-feature correlations. Crucially, as a preliminary study to validate our intuition, our framework introduces a structured analysis phase that aggregates anomalies spatially, enabling the robust identification of meaningful urban events and demonstrating that mobile traffic effectively mirrors the movement and activity of the city’s inhabitants.

3. Methodology

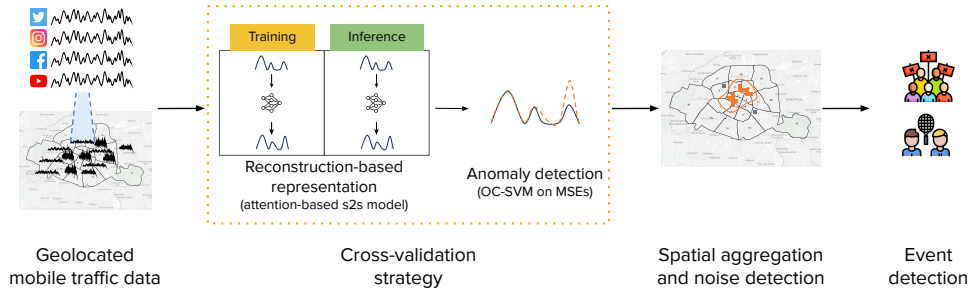


Figure 1: Proposed methodology.

In this section, we describe the proposed methodology, which serves as a framework to evaluate the feasibility of using mobile traffic as a proxy for urban mobility. We first provide an overview of the pipeline’s stages and then detail how each phase is instantiated to identify urban events through the lens of anomaly detection.

3.1. Overview

We propose the processing pipeline illustrated in Figure 1. This framework is designed to demonstrate that service-specific mobile traffic can mirror urban dynamics by formulating the detection of urban events as a TSAD problem applied to a regular grid-partitioned city.

As represented in Figure 1, the methodology begins by transforming raw mobile traffic into normalized, service-specific time series, segmented into fixed-length windows (e.g., daily 24-hour sequences) for each tile, i.e., a cell of the regular grid. These tiles are linked to administrative districts, providing spatial context. In a cross-validation strategy (second step), we split the administrative districts into several folds of training and validation sets with a held-out strategy. On the training set, a reconstruction-based model is then trained to learn regular patterns of multivariate activity. At test time, sequences from held-out districts are reconstructed, and deviations — measured by reconstruction error — signal potential anomalies. At the end of the cross-validation process, tiles that are accounted as anomalies in a sufficient number of folds are subsequently spatially aggregated to highlight coherent and contiguous urban events (third step). In the following, we motivate the different steps and provide a possible implementation of them.

Mobile traffic time series data. We rely on service-specific mobile traffic data, which provides the temporal evolution of application-level network usage within the city. The study area is first tessellated into non-overlapping regular tiles, i.e., a grid. For each tile, we analyze the relative multivariate time series representing the traffic volume associated with different mobile services or applications. These time series are spatially anchored to their respective tiles, thus linking traffic dynamics to specific urban areas. Prior to subsequent steps, the data undergo basic preprocessing, including temporal segmentation into fixed-length windows (e.g., 24-hour slots) and normalization.

Cross-validation strategy. We design this step of our methodology around a cross-validation framework based on a held-out split of administrative districts. The districts of the city are divided into several folds, each including distinct districts used alternately for training and validation. More details about this methodology are provided in Section 4.

Reconstruction-based representation. Within each training fold, a reconstruction-based model is trained to reproduce the input time series, thereby learning the regular temporal dynamics of multivariate activity in an unsupervised manner. By minimizing reconstruction loss, the model captures typical patterns and interdependencies among variables, enabling it to generate coherent multivariate sequences that reflect normal behavior. Consequently, when the reconstruction error is high, this indicates that the observed sequence deviates from the learned normal patterns, implicitly signaling the presence of an anomaly. Following this idea, at test time, sequences from the held-out districts are reconstructed, and deviations – quantified through reconstruction error – signal potential anomalies. Moreover, the learned latent representation provides an informative embedding of the time series that can be leveraged for other downstream tasks, such as clustering or similarity analysis.

Anomaly detection. The anomaly detection stage uses the reconstruction model’s output to flag anomalies in new, unseen data. After training the reconstruction model, we apply it to test time series – data from tiles belonging to districts held out of training – to obtain reconstructed estimates. For each sequence, the reconstruction error, computed as the discrepancy between the observed and reconstructed multivariate time series, serves as an anomaly score. An anomaly detector (such as a one-class classifier or a thresholding method) is then employed to decide if the reconstruction error of the current multivariate time series is significantly higher than what is expected under normal conditions. This should ensure that anomalies are identified without any information from the test set leaking into the model training. Details on the design of our cross-validation scheme will be provided below. Therefore, anomalies, corresponding to potential events, are identified as instances where the reconstruction significantly diverges from the input signal.

Spatial aggregation and noise reduction. A single significant urban event may trigger anomalies at multiple adjacent locations. We therefore aggregate individual anomalous observations into clusters that correspond to real-world events. Thus, if multiple neighboring locations are flagged as anomalous, we group them into a single cluster.

The end result of the pipeline is a set of detected events, each characterized by spatial contiguity, indicating where the mobile traffic deviated significantly from normal. This general methodology ensures a clear separation between learning normal behavior (training phase) and detecting anomalous events (testing phase), with robust spatial cross-validation to prevent overfitting or data leaks.

In the next section, we provide a detailed description of the implementation of each step of this proposed framework.

3.2. Implementation

We instantiate the above general pipeline with specific modeling choices tailored to the characteristics of our mobile traffic data. In our implementation, an attention-based sequence-to-sequence model is used for the reconstruction-based representation learning stage, a One-Class Support Vector Machine (OC-SVM) serves as the anomaly detector, and a custom polygon-based clustering method groups anomalies spatially to identify events.

We begin by preprocessing the mobile network traffic dataset to extract normalized time series of daily traffic volumes, aggregated into hourly intervals for each tile across the urban area. Each tile is associated with an administrative district, following the official municipal partitioning of the city. We adopt a district-based cross-validation strategy: in each fold, held-out districts are reserved for testing, so that each appears at least τ times in the test set. This spatial partitioning ensures clear separation between training and test data, enabling robust generalization to unseen urban areas.

For the **reconstruction-based representation** phase, we adopt an attention-based sequence-to-sequence model that takes a time series segment as input and attempts to reconstruct it at the output. As illustrated in Figure 2, the model consists of an encoder and a decoder, both built with Transformer layers leveraging multi-head self-attention. The encoder receives the sequence $X = [\mathbf{x}_1, \dots, \mathbf{x}_T]$ for a batch of districts along with positional encodings that indicate the temporal order of the sequence. Through stacked self-attention layers, the encoder produces a set of latent representations $H = [\mathbf{h}_1, \dots, \mathbf{h}_d]$ capturing the salient patterns in the input sequence, where d represents the dimension of the embeddings.

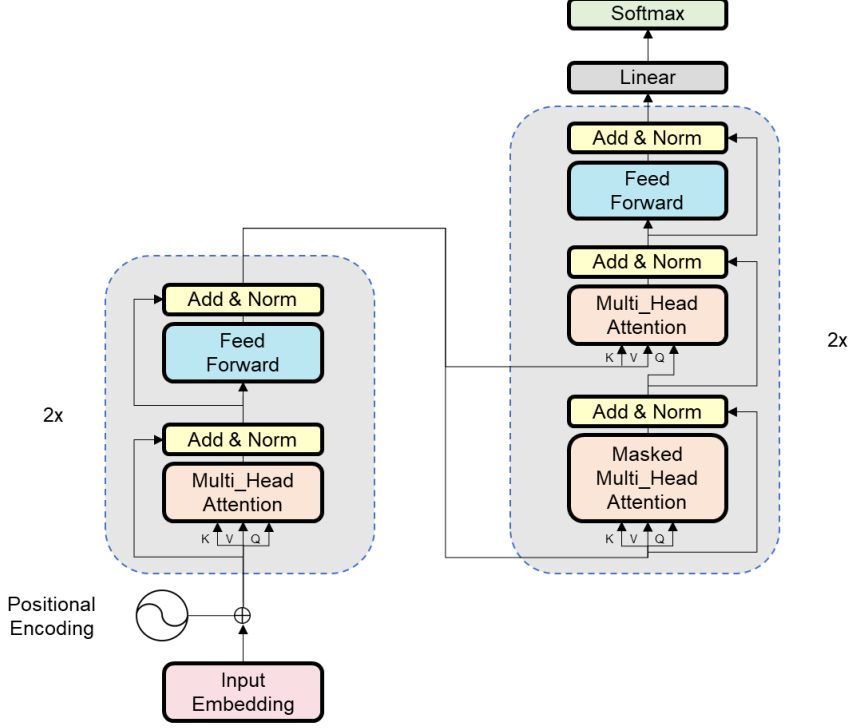


Figure 2: The attention-based sequence-to-sequence model architecture used as the reconstruction-based model.

The decoder then takes H and learns to produce an output sequence $\hat{X} = [\hat{x}_1, \dots, \hat{x}_T]$ of the same length, ideally reproducing the original input X .

We train the model by minimizing a reconstruction loss between X and \hat{X} . In particular, we adopted the Mean Squared Error (MSE) averaged over batch size B , sequence length T , and number of services N :

$$L = \frac{1}{B \times T \times N} \sum_{b=1}^B \sum_{t=1}^T \sum_{n=1}^N (X_{b,t,n} - \hat{X}_{b,t,n})^2$$

where $X_{b,t,n}$ represented the original value at time step t for service n in batch sample b , and $\hat{X}_{b,t,n}$ was the corresponding reconstruction. This formulation encouraged accurate reproduction of the multivariate time series by penalizing large deviations in the reconstruction.

After training the model on the training set, we proceed to the **anomaly detection** phase, which is carried out exclusively on test data. For each input sequence in the test set, the trained model produces a reconstruction, and the corresponding reconstruction error is computed as the Mean Squared Error (MSE) between the original and reconstructed sequences. These errors serve as indicators of deviations from normal behavior, with unusually high values suggesting the presence of anomalous patterns. To distinguish anomalous from regular sequences, we employ a *One-Class Support Vector Machine* (OC-SVM) [22], trained on the distribution of reconstruction errors observed in the training set. The OC-SVM learns a decision boundary that tightly encloses the majority of training samples and flags outliers in the test set as potential anomalies, thereby enabling unsupervised event detection without requiring labeled examples. This training–testing procedure, coupling reconstruction and anomaly detection, is repeated over all cross-validation folds. As each time series appears in the test folds the same number of times, we count how often it is classified as anomalous across folds. Time series identified as anomalous in a sufficient number of folds are then considered *robust anomalies* and passed to the subsequent analysis stage.

The final step is to interpret tile-related anomalies in the context of real-world events. Isolated anomalies might be false alarms or minor fluctuations, whereas when multiple anomalies occur in proximity, they likely correspond to a significant event. We therefore perform clustering on the detected anomalous sequences, aggregating adjacent tiles into contiguous clusters and discarding when isolated.

In particular, to identify spatially coherent events, we apply a polygon-based clustering method to daily anomalies. For each day, we construct an undirected graph whose nodes represent anomalous tiles and whose edges connect tiles sharing a polygon boundary (excluding vertex-only adjacency). Connected components of this graph correspond to contiguous anomaly clusters. Clusters with fewer than a threshold number of contiguous tiles are labeled as noise.

This combination of temporal consistency and spatial aggregation provides a reliable characterization of anomalies, enabling the distinction between significant large-scale events and smaller fluctuations while supporting their interpretation in terms of spatial diffusion and temporal evolution.

4. Experiments

In this section, we describe the experimental evaluation of the proposed methodology. We first introduce the dataset and the preprocessing steps, and outline the experimental setup used to train and validate the model under a spatial cross-validation scheme. Finally, we present and discuss the results, illustrating how the detected anomalies correspond to real-world urban events observed during the study period.

4.1. Dataset

We base our experiments on the NetMob 2023 dataset [4], which contains anonymized service-specific mobile traffic data collected over 77 consecutive days in 20 metropolitan regions, including major French cities such as Paris, Lyon, Marseille, Toulouse, and Nice. The data are recorded at a fine spatial resolution of 100×100 m and at a temporal resolution of 15 minutes. Each city is represented by a grid-based structure, dividing the urban area into non-overlapping tiles. The dataset captures activity from 68 services, including widely used mobile applications such as YouTube, Netflix, Facebook, Instagram, and Gmail. For our analysis, we aggregated the data to hourly intervals and focused on the Paris metropolitan area, specifically the 20 central districts (Figure 3). Among all available services, we selected four representative ones (Twitter, YouTube, Instagram, Facebook), since social networks are particularly informative of different kinds of events, ranging from natural disasters to concerts.

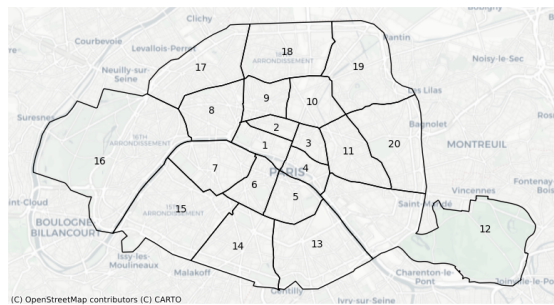


Figure 3: Map of the 20 administrative districts of Paris.

4.2. Experimental setup

Input data were normalized and then segmented into fixed-length multivariate sequences of 24 timestamps, corresponding to one day of traffic for each tile. This enables the model to learn temporal patterns within each service as well as interactions between services over time. Feeding the reconstruction-based model with sequences of 24 timestamps across all four services per tile allows it to jointly model these dynamics and to detect deviations from normal behavior.

The reconstruction-based model was trained using the AdamW optimizer with a learning rate of 1×10^{-4} and weight decay of 1×10^{-5} . Training was conducted for up to 50 epochs, with early stopping based on the validation loss to prevent overfitting. Gradient clipping with a maximum norm of 1.0 was applied to stabilize training and avoid exploding gradients.

Since we operate in a fully unsupervised setting, the design of the validation strategy is critical to avoid data leakage between the training and test sets. To ensure spatial consistency and avoid random sampling, we implemented the district-based cross-validation scheme described in Section 3, assigning each tile to its corresponding district. The 20 districts of Paris served as partition units, with four districts assigned to the test set and the remaining sixteen to the training set in each fold. A subset of 15 folds was generated (using a fixed random seed), ensuring that each district appeared in the test set at least three times. Moreover, this cross-validation design mitigates the effect of potential anomalies present in the training data. Since each fold is trained on a different subset of districts, occasional anomalous patterns in the training set are unlikely to systematically bias the model, thus enhancing robustness in a fully unsupervised setting.

After training, anomaly detection was carried out on the reconstruction errors (MSEs) produced by the reconstruction-based model. To classify errors, we employed an OC-SVM, setting the ν parameter to 0.01, controlling the upper bound on the fraction of anomalies and the lower bound on the fraction of support vectors.

The outputs from all cross-validation folds were aggregated, where each tile was assigned a daily anomaly score over the 77-day period. A tile could therefore be flagged as anomalous in 0, 1, 2, or 3 folds on a given day, depending on its level of consensus across models. A conservative threshold of 3 was adopted to define robust anomalies, though less stringent thresholds are also possible.

As described above, to identify spatially coherent events, we applied a polygon-based clustering method to the daily anomalies described above.

Clusters with less than five contiguous tiles were labeled as *noise*, removing spurious detected anomalies with limited spatial diffusion.

4.3. Results

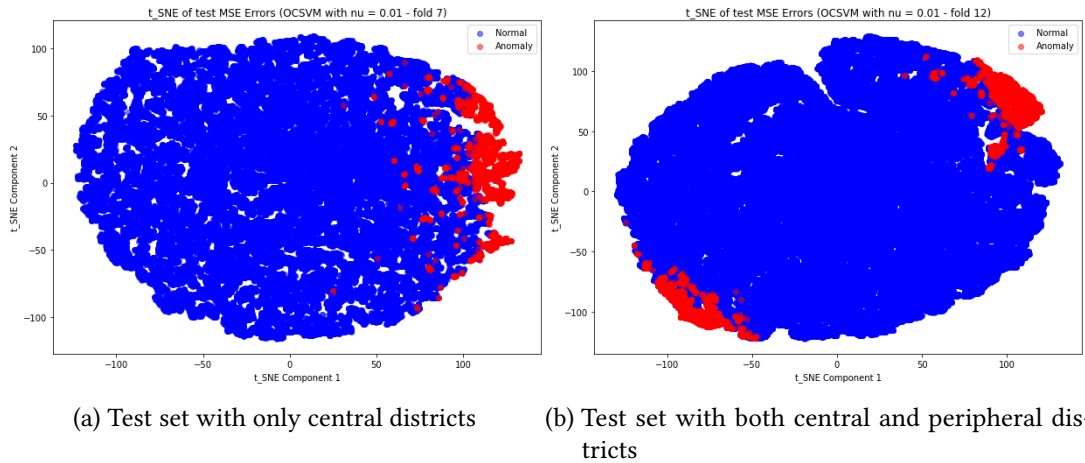


Figure 4: Comparison of t-SNE plots for different test folds.

We begin by analyzing the embeddings across cross-validation folds. To assess whether the composition of testing districts (central, peripheral, or mixed) influences the learned representations, we visualize them using *t-Distributed Stochastic Neighbor Embedding* (t-SNE), a dimensionality-reduction technique. Concretely, we project the high-dimensional reconstruction-error vectors into two dimensions; t-SNE preserves local neighborhoods and reveals cluster structure, yielding a clear separation between normal and anomalous sequences [23]. Comparing t-SNE plots across cross-validation folds shows an influence of spatial context (Figure 4): when only central districts form the test set, anomalous

MSEs form a single compact group clearly separated from normal samples; when peripheral districts are included, anomalous MSEs split into two groups at opposite extremes of the embedding, indicating heterogeneous behaviors by location.

Now, we analyze notable events within the observation period by leveraging the scores obtained from the combination of district-based cross-validation and spatial clustering described above. Daily anomalies are first identified via cross-validation consensus and then aggregated spatially into contiguous clusters, distinguishing spatially contiguous clusters from sporadic *noise*. After clustering anomalous daily time series, we identify noisy ones by reducing the number of initial anomalous clusters by 69.32%.

We begin by examining the anomalies and reference the DBUE dataset introduced by [3], which provides a structured classification of common urban events in Paris in the same period as the Netmob 2023 dataset. A preliminary comparison shows close alignment between our detected anomalies and categories such as disaster-related incidents, protests, traffic disruptions, and large sporting activities, motivating a set of representative case studies in the Paris context. Since the Netmob 2023 dataset does not include labeled anomalies and our analysis is conducted in a fully unsupervised setting, we assign labels to the clusters obtained after the clustering and noise detection step of our methodology by cross-referencing major French news sources and official communication channels (such as France 24¹ and the Élysée website²), as well as the DBUE dataset. All identified events are reported in Table 1. For instance, on March 16, 2019, the model highlights clusters tracing the route of the Marche du Siècle, a large climate and social justice demonstration that coincided with a major Yellow Vests protest. Similarly, significant clusters appear on April 25, coinciding with President Macron’s speech concluding the Grand Débat National, and on May 16 during the inauguration of VivaTech at Porte de Versailles, where the presence of international leaders and large crowds produced distinct spatial signals.

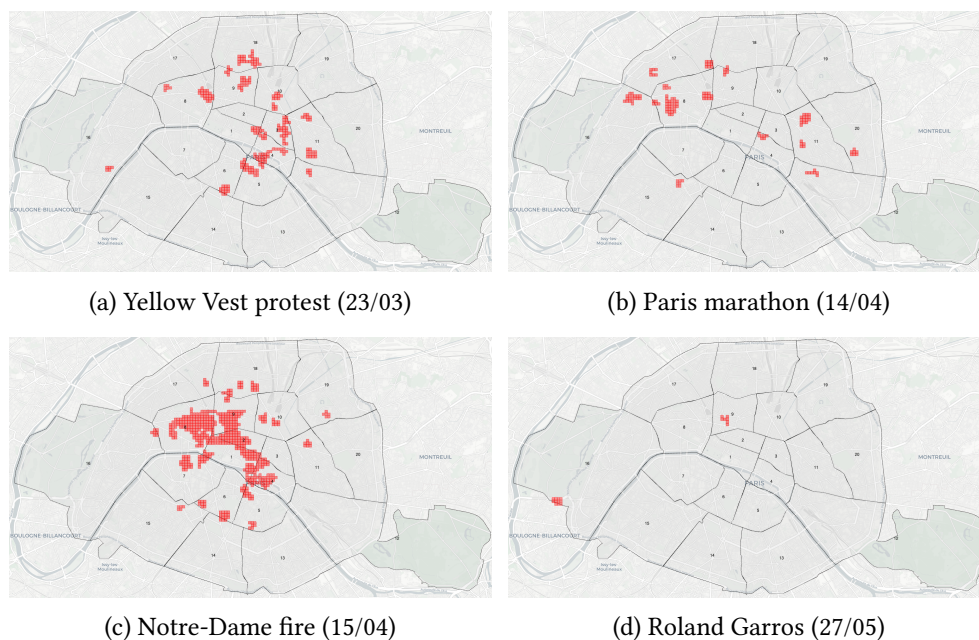


Figure 5: Four key anomalous and labeled days.

Overall, our model is particularly effective in detecting large-scale protests (e.g., the Yellow Vests) and prolonged events with broad urban impact. Notice that many of these events are not reported in the DBUE dataset, which primarily catalogs shorter occurrences such as concerts or localized gatherings. For instance, major Yellow Vests demonstrations between mid-March and late April 2019, as well as institutional events such as the presidential press conference concluding the National Grand Debate,

¹<https://www.france24.com/en/>

²<https://www.elysee.fr/en/>

are not documented in the DBUE dataset despite their impact on the city. As summarized in Table 1, the events identified by our approach span a diverse range of typologies: large collective movements (e.g., protests and demonstrations), exceptional disruptions (e.g., the Notre-Dame cathedral fire), and long-lasting festivals (e.g., the Roland-Garros), often occurring in areas that usually display low or regular traffic activity.

Event name	Dates	Day(s)	Source(s)	DBUE
Yellow Vests protests	16/03/2019	Saturday	France 24	Yes
Climate march	16/03/2019	Saturday	France 24	
EcoTrail Paris 2019	16/03/2019	Saturday	EcoTrail Paris ³	Yes
Inauguration Ceremony of the Newly appointed Paris Police Prefect	21/03/2019	Thursday	Ministère de l'Intérieur ⁴	
Yellow Vests protests	23/03/2019	Saturday	France 24	
Yellow Vests protests	06/04/2019	Saturday	France 24	
Yellow Vests protests	13/04/2019	Saturday	France 24	
Paris Marathon	14/04/2019	Sunday	France 24	Yes
Notre-Dame fire	15–16/04/2019	Mon–Tue	France 24	Yes
Yellow Vests protests	20/04/2019	Saturday	France 24	
Press conference: end of National Grand Debate	25/04/2019	Thursday	Elysée	
Yellow Vests protests	27/04/2019	Saturday	France 24	
May Day protests	01/05/2019	Wednesday	France 24	Yes
VivaTech inauguration	16/05/2019	Thursday	France 24	
Roland-Garros Kids' Day	25/05/2019	Saturday	Roland-Garros ⁵	
Roland-Garros	26–31/05/2019	Sun–Fri	Roland-Garros	Yes

Table 1

Events emerged and labeled during our analysis compared with events reported in the DBUE dataset.

In what follows, we focus on four key events within the observation period: the Yellow Vest protests (particularly the one that occurred on 23 March) [24, 25], the Paris Marathon (14 April), the Notre-Dame fire (15 April), and the Roland Garros tennis tournament (25 May–9 June). These events were selected both for their significance during the NetMob dataset timeframe and for the model's ability to capture their dynamics, despite their heterogeneous nature.

For instance, during the Paris Marathon, the reconstruction-based model identifies anomalous geolocated time series in the areas corresponding to the race's starting and finishing points (Avenue Champs-Élysées and Avenue Foch), in the left-top of Figure 5b. In contrast, during the Yellow Vest protests, anomalies appear around key demonstration sites, overlapping with the main locations of gatherings and disruptions, such as Place Denfert-Rochereau and Montmartre, which served as the starting and ending points of the March 23 protest, as well as along major central streets where clashes occurred, notably Boulevard de Strasbourg, as shown in Figure 5a.

As for the Notre-Dame fire, although the event itself was localized to a specific site, its impact resonated far beyond the immediate area. On that day, anomalous mobile traffic was detected almost uniformly across the entire Paris metropolitan area, reflecting the widespread attention and disruption associated with the event, as we observed in Figure 5c.

For the Roland Garros tennis event (26 May–9 June; our analysis covers 26–31 May), anomalous

clusters were mapped across all 20 Paris districts. Interestingly, a characteristic cluster was detected on each day within the 26–31 May period, as well as on 25 May, one day before the official start of the tournament, corresponding to the Roland Garros Kids’ Day. During the analyzed event days, the district where the event was held, i.e., district 16, which generally does not show anomalous dynamics, exhibited a single localized cluster precisely in the area of the Roland Garros Stadium, as shown in the left-bottom corner of Figure 5d.

In conclusion, the fact that our findings match events in the DBUE dataset, particularly those involving significant population movement, demonstrates the reliability of our intuition. Furthermore, the fact that some anomalies not listed in DBUE were successfully verified by news sources as major protests and gatherings reinforces this finding. This preliminary study confirms that service-specific mobile traffic data is indeed an effective proxy for capturing and understanding human mobility dynamics in the urban space.

5. Conclusion

In this paper, we propose a preliminary approach to demonstrate that mobile traffic data can serve as a good proxy for human mobility. We model the problem as a TSAD task by introducing an unsupervised methodology for detecting events that occur in an urban environment from the analysis of anomalies in specific-service mobile traffic data.

The proposed methodology employs a transformer-based autoencoder as the reconstruction model to learn compact representations of temporal dynamics, capturing complex nonlinear dependencies across services and locations. The reconstruction errors are then evaluated using an OC-SVM to isolate anomalous time series in a fully unsupervised manner. Finally, a spatial clustering algorithm aggregates anomalies and filters out noise, revealing coherent spatial patterns. By leveraging daily traffic volumes of four major platforms (Twitter, YouTube, Instagram, Facebook), the methodology identifies localized anomaly clusters across Paris’s central districts over 77 days (March–May 2019).

We demonstrated that our approach, operating at daily temporal resolution on social media traffic volumes, is able to capture anomalies that correspond to long and movement events, but also events that have a big resonance even if they occurred in a specific place, demonstrating that with this dataset and our methodology we are able to uncover additional events with respect to other works like [3]. Examples of these events are marches and demonstrations, such as the Yellow Vests, or catastrophic events, such as the Notre-Dame fire.

Finally, our study confirms the importance of exploiting multi-platform digital traces to capture the dynamics of disruptive urban events. In particular, our pipeline uses traffic volumes from four major platforms (Twitter, Facebook, YouTube, Instagram). Among these, Twitter and Facebook stand out as the most relevant sources in relation to prior research. Indeed, [26] developed an *atlas multi-plateforme* of the Yellow Vests movement, demonstrating how Facebook groups and Twitter interactions provided a privileged observation post to map the heterogeneity and dynamics of claims. This convergence with our findings supports the idea that combining cross-platform signals can enhance the robustness of urban anomaly detection.

Future work will focus on refining and automating the labeling and identification of detected anomalies to reduce manual intervention and improve scalability. We also plan to explore alternative modeling and detection strategies, for instance, experimenting with different reconstruction architectures or anomaly scoring mechanisms, to assess whether these adjustments can improve sensitivity to shorter, less pervasive, or semantically different types of events compared to those identified in this study.

To conclude, our results demonstrate that multivariate time series anomaly detection on social media traffic can uncover both event-driven anomalies and systematic spatiotemporal patterns in urban dynamics. This positions our pipeline as a complementary tool to high-resolution network signaling analysis and text-based social media event detection, offering new perspectives for monitoring urban events and their effects on collective behavior.

Declaration on Generative AI

During the preparation of this work, the author(s) used GPT-5 in order to: Grammar and spelling check. After using these tool(s)/service(s), the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content.

Acknowledgments

The NetMob 2023 dataset has been made available by Orange within the framework of the NetMob 2023 Challenge [4]. This research was supported by the “Resilienza Economica e Digitale” project (CUP D67G23000060001) funded by the Italian Ministry of University and Research (MUR) as “Department of Excellence” (Dipartimenti di Eccellenza 2023-2027, Ministerial Decree no. 230/2022) and by the European Union under the Italian National Recovery and Resilience Plan (NRRP) of NextGenerationEU, partnership on “Telecommunications of the Future” (PE00000001 - program “RESTART”).

References

- [1] Y. Dong, F. Pinelli, Y. Gkoufas, Z. Nabi, F. Calabrese, N. V. Chawla, Inferring unusual crowd events from mobile phone call detail records, in: A. Appice, P. P. Rodrigues, V. Santos Costa, J. Gama, A. Jorge, C. Soares (Eds.), *Machine Learning and Knowledge Discovery in Databases*, Springer International Publishing, Cham, 2015, pp. 474–492.
- [2] Y. Jiang, A. A. Popov, Z. Li, M. E. Hodgson, B. Huang, A sensor-based simulation method for spatiotemporal event detection, *ISPRS International Journal of Geo-Information* 13 (2024). URL: <https://www.mdpi.com/2220-9964/13/5/141>. doi:10.3390/ijgi13050141.
- [3] P. Lemaire, A. Furno, S. Rubrichi, A. Bondu, Z. Smoreda, C. Ziemlicki, et al., Early detection of critical urban events using mobile phone network data, *PLoS ONE* 19 (2024) e0309093. URL: <https://doi.org/10.1371/journal.pone.0309093>. doi:10.1371/journal.pone.0309093.
- [4] O. E. Martínez-Durive, S. Mishra, C. Ziemlicki, S. Rubrichi, Z. Smoreda, M. Fiore, The netmob23 dataset: A high-resolution multi-region service-level mobile data traffic cartography, 2023. URL: <https://arxiv.org/abs/2305.06933>. arXiv:2305.06933.
- [5] Z. Zamanzadeh Darban, G. I. Webb, S. Pan, C. Aggarwal, M. Salehi, Deep learning for time series anomaly detection: A survey, *ACM Computing Surveys* 57 (2024) 1–42. URL: <http://dx.doi.org/10.1145/3691338>. doi:10.1145/3691338.
- [6] F. Wang, Y. Jiang, R. Zhang, A. Wei, J. Xie, X. Pang, A survey of deep anomaly detection in multivariate time series: Taxonomy, applications, and directions, *Sensors* 25 (2025) 190. doi:10.3390/s25010190.
- [7] H. Torabi, S. Mirtaheri, S. Greco, Practical autoencoder based anomaly detection by using vector reconstruction error, *Cybersecurity* 6 (2023). doi:10.1186/s42400-022-00134-9.
- [8] G. Trappolini, A. Purificato, F. Siciliano, L. D’Addona, A. Spagnolo, D. Dato, F. Silvestri, Quantized auto encoder-based anomaly detection for multivariate time series data in 5g networks, *IEEE Access* PP (2025) 1–1. doi:10.1109/ACCESS.2025.3568133.
- [9] S. Guan, B. Zhao, Z. Dong, M. Gao, Z. He, Gtad: Graph and temporal neural network for multivariate time series anomaly detection, *Entropy* 24 (2022). URL: <https://www.mdpi.com/1099-4300/24/6/759>. doi:10.3390/e24060759.
- [10] Z. Chen, D. Chen, X. Zhang, Z. Yuan, X. Cheng, Learning graph structures with transformer for multivariate time-series anomaly detection in iot, *IEEE Internet of Things Journal* 9 (2022) 9179–9189. URL: <http://dx.doi.org/10.1109/JIOT.2021.3100509>. doi:10.1109/jiot.2021.3100509.
- [11] H. Zhao, Y. Wang, J. Duan, C. Huang, D. Cao, Y. Tong, B. Xu, J. Bai, J. Tong, Q. Zhang, Multivariate time-series anomaly detection via graph attention network, 2020. URL: <https://arxiv.org/abs/2009.02040>. arXiv:2009.02040.

- [12] J. Xu, H. Wu, J. Wang, M. Long, Anomaly transformer: Time series anomaly detection with association discrepancy, 2022. URL: <https://arxiv.org/abs/2110.02642>. arXiv:2110.02642.
- [13] S. Tuli, G. Casale, N. R. Jennings, Tranad: Deep transformer networks for anomaly detection in multivariate time series data, 2022. URL: <https://arxiv.org/abs/2201.07284>. arXiv:2201.07284.
- [14] S. Ahmed, I. Nielsen, A. Tripathi, S. Siddiqui, G. Rasool, R. Ramachandran, Transformers in time-series analysis: A tutorial, 2022. doi:10.48550/arXiv.2205.01138.
- [15] F. Rewicki, J. Denzler, J. Niebling, Is it worth it? comparing six deep and classical methods for unsupervised anomaly detection in time series, *Applied Sciences* 13 (2023) 1778. URL: <http://dx.doi.org/10.3390/app13031778>. doi:10.3390/app13031778.
- [16] Y. Jeong, E. Yang, J. H. Ryu, I. Park, M. Kang, Anomalybert: Self-supervised transformer for time series anomaly detection using data degradation scheme, 2023. URL: <https://arxiv.org/abs/2305.04468>. arXiv:2305.04468.
- [17] X. Peng, Y. Lin, Q. Cao, Y. Cen, H. Zhuang, Z. Lin, Traffic anomaly detection in intelligent transport applications with time series data using informer, in: *Conference: 2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*, 2022, pp. 3309–3314. doi:10.1109/ITSC55140.2022.9922142.
- [18] J. Guan, Y. Liu, Q. Kong, F. Xiao, Q. Zhu, J. Tian, Transformer-based autoencoder with id constraint for unsupervised anomalous sound detection, *EURASIP Journal on Audio, Speech, and Music Processing* 2023 (2023). doi:10.1186/s13636-023-00308-4.
- [19] D. Yuxiao, F. Pinelli, Y. Gkoufas, Z. Nabi, F. Calabrese, N. Chawla, Inferring unusual crowd events from mobile phone call detail records, in: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 2015. doi:10.1007/978-3-319-23525-7_29.
- [20] E. Akopyan, N.-E. El Faouzi, E. Gaume, Unsupervised real-time anomaly detection for multivariate mobile phone traffic series, in: *ESANN 2021 - European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning*, 2021, pp. 469–474. doi:10.14428/esann/2021.ES2021-113.
- [21] J. Lin, J. Fukuyama, Calibrating dimension reduction hyperparameters in the presence of noise, *PLOS Computational Biology* 20 (2024) 1–19. URL: <https://doi.org/10.1371/journal.pcbi.1012427>. doi:10.1371/journal.pcbi.1012427.
- [22] L. M. Manevitz, M. Yousef, One-class svms for document classification, *J. Mach. Learn. Res.* 2 (2002) 139–154.
- [23] S. Jung, T. Dagobert, J.-M. Morel, G. Facciolo, A review of t-sne, *Image Processing On Line* 14 (2024) 250–270. doi:10.5201/ipol.2024.528.
- [24] É. Walker, P. Blavier, Shedding light on the Yellow Vests movement through its spatial dimension, from the gathering places to the origins of the participants, *French Politics* 20 (2022) 318–345. URL: <https://shs.hal.science/halshs-03832223>. doi:10.1057/s41253-022-00187-9.
- [25] M. D. Sudda, E. Reungoat, Understanding the French Yellow Vests Movement through the lens of mixed methods, *French Politics* 20 (2022) 303–317. URL: <https://shs.hal.science/halshs-03808821>. doi:10.1057/s41253-022-00188-8.
- [26] P. Ramaciotti Morales, J.-P. Cointet, B. Benbouzid, D. Cardon, C. Froio, O. F. Metin, B. Ooghe, G. Plique, Atlas multi-plateforme d'un mouvement social : le cas des Gilets jaunes, 2021. URL: <https://hal.science/hal-03700086>.