IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.151181

Multi-Purpose Tourist Safety Monitoring and Incident Response System using AI, GRU Autoencoder, Geo-fencing and Blockchain Technology

Arti Suryavanshi[1], Jivan Divate[2], Jayshree Gaikwad[3], Yashashree Chakane[4]

Professor, Computer Department, HSBPVT's GOI Faculty of Engineering, Kashti, India[1]

Student, Computer Department, HSBPVT's GOI Faculty of Engineering, Kashti, India[2]

Student, Computer Department, HSBPVT's GOI Faculty of Engineering, Kashti, India[3]

Student, Computer Department, HSBPVT's GOI Faculty of Engineering, Kashti, India[4]

Abstract: Ensuring tourist safety in unfamiliar, high-risk, or remote environments requires continuous spatial-temporal profiling, real-time anomaly detection, and tamper-proof incident logging. Traditional safety frameworks rely heavily on static geographic perimeters and manual emergency reporting, leaving monitoring systems highly susceptible to high false-alarm rates, delayed emergency response times, and single-point data tampering. This paper introduces a multi-purpose, end-to-end tourist safety monitoring and incident response architecture that bridges client-side geodetic filtering, deep sequence-to-sequence anomaly detection, automated AI risk reasoning, and decentralized blockchain audit logging. The client application incorporates an accuracy-aware 2D Geodetic Kalman Filter operating on strict 5-second sampling intervals to eliminate multipath GPS sensor noise and atmospheric drift, storing smoothed telemetry within a 31-sample sliding window buffer. A client-side geofence state-machine dynamically manages movement relative to custom danger zones and automated +50-meter geodesic caution buffers, eliminating redundant API server polling and minimizing gas overhead. At the core of the backend engine, 31 raw coordinate points are transformed into 30-step, 4-dimensional kinematic feature matrices $X \in \mathbb{R}30 \times 4$—encompassing speed, acceleration, compass bearing deviation, and stationary duration flags—and evaluated by a PyTorch Gated Recurrent Unit (GRU) Autoencoder structured with a 32-dimensional latent bottleneck ($z \in \mathbb{R}32$). Trained on nominal human mobility data, the GRU autoencoder calculates sequence reconstruction Mean Squared Error (MSE), flagging movement irregularities, distress patterns, or erratic behavior when reconstruction loss breaches a calibrated decision threshold of $\tau = 0.55$. Upon a geofence state transition or trajectory anomaly, the incident triggers a cascading multi-model LangChain LLM fallback chain to generate structured, schema-validated JSON threat evaluation reports. Simultaneously, SHA-256 hashes of breach telemetry are immutably signed and broadcast as zero-value transactions to an Ethereum/EVM blockchain ledger for forensic non-repudiation. Evaluated on real-world spatial trajectory datasets, the GRU autoencoder achieves an initial training loss descent from 1.78 to 1.18, a fine-tuned training MSE of 0.094, an out-of-sample test reconstruction loss of 0.140, and an average pipeline execution latency of 6.8 ms. The proposed architecture demonstrates a highly scalable, robust, and tamper-proof solution for real-time emergency dispatching and proactive tourism security management.

Keywords: Anomaly Detection, Gated Recurrent Unit (GRU), Autoencoders, Spatial-Temporal Trajectories, Geofencing, Blockchain, Tourist Safety Systems.

I. INTRODUCTION

Ensuring the physical safety and operational monitoring of tourists navigating unfamiliar, hazardous, or high-risk environments represents a critical imperative for modern smart tourism frameworks and civil emergency dispatch services. As global travel expands across dense urban centers and isolated natural topographies, tourists face heightened vulnerabilities including physical disorientation, sudden medical distress, unauthorized entry into hazardous zones, and targeted criminal threats. Traditional safety architectures rely heavily on manual emergency reporting mechanisms, static geographic perimeters, and raw mobile Global Positioning System (GPS) telemetry feeds. However, unrefined mobile GPS data streams are inherently compromised by multipath signal degradation, atmospheric satellite drift, and urban canyon signal reflection, which generate high positioning noise and artificial telemetry spikes. When raw coordinate sequences are evaluated without kinematic filtering, transient positioning jumps produce false velocity spikes, corrupting spatial trajectory profiling and triggering high rates of false emergency alarms. Furthermore, conventional monitoring

 This work is licensed under a Creative Commons Attribution 4.0 International License



International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026
DOI: 10.17148/IJARCCE.2026.151181

models depend on centralized database infrastructure, leaving historical location logs and emergency audit trails vulnerable to single-point operational failures and post-incident data tampering.

To establish an effective, real-time tourist monitoring infrastructure capable of immediate emergency dispatching, automated safety systems must overcome three primary technical challenges:

1. *Raw GPS Telemetry Noise and False Anomaly Triggers:* Mobile GPS coordinates suffer from receiver jitter and indoor multipath jumps, which artificially distort speed and acceleration metrics unless filtered by pedestrian kinematic process models.
2. *Backend API Endpoint Overhead and Blockchain Gas Waste:* Unthrottled perimeter pinging and continuous spatial boundary checking saturate microservice API endpoints and generate unsustainable transaction gas costs when writing continuous telemetry feeds to public decentralized ledgers during extended stays inside safe or caution zones.
3. *Centralized Log Vulnerabilities and Unstructured Emergency Context:* Standard relational logging lacks cryptographic non-repudiation, while raw coordinate pairs fail to deliver the actionable, structured risk assessment required for immediate dispatch by first responders.

To address these systemic operational limitations, this paper presents an end-to-end, multi-purpose Tourist Safety Monitoring and Incident Response System. The software architecture combines a vector-map frontend built on React 19 with a microservice backend powered by FastAPI. On the client side, raw geolocation feeds pass through an accuracy-aware 2D Geodetic Kalman Filter operating on strict 5-second sampling intervals. The filter models pedestrian kinematics by continuously balancing device-reported measurement noise covariance against process noise, smoothing incoming spatial telemetry into a stabilized 31-sample sliding window buffer. To prevent network endpoint saturation, a client-side spatial state-machine (SAFE -> CAUTION -> DANGER) monitors user position against dynamic circular or polygon danger zones drawn directly by municipal authorities or users, generating an automated 50-meter geodesic caution buffer. Telemetry pings are transmitted to the backend exclusively upon state transitions rather than continuous polling, suppressing redundant network bandwidth and eliminating unnecessary smart contract executions.

At the core of the backend analytical pipeline, spatial coordinate sequences are transformed into 4-dimensional kinematic feature matrices encompassing velocity, linear acceleration, compass bearing deviation, and stationary state flags calculated via vectorized Haversine transformations. Trajectory anomaly evaluation is performed by a deep sequence-to-sequence Gated Recurrent Unit (GRU) Autoencoder constructed with a 32-dimensional latent bottleneck. Trained exclusively on nominal pedestrian movement patterns, the autoencoder attempts to reconstruct input sequences; when a user exhibits panic behaviors such as erratic directional turns, sudden running, or abrupt stops, the reconstruction Mean Squared Error (MSELoss) spikes beyond a calibrated empirical decision threshold of 0.55, triggering an automated anomaly alert.

Upon a spatial geofence breach or trajectory anomaly detection, the incident is routed to a cascading multi-model LangChain AI reasoning pipeline (openai/gpt-oss-120b -> qwen/qwen3.8-27b -> openai/gpt-oss-20b -> openai/gpt-oss-safeguard-20b). The LLM engine evaluates spatial-temporal context, anomaly severity scores, and localized geofence parameters to output schema-validated JSON threat evaluations (LOW, MEDIUM, HIGH, CRITICAL) alongside recommended emergency protocols. Concurrently, SHA-256 cryptographic hashes of the incident telemetry are signed and broadcast to an Ethereum/EVM-compatible blockchain network via Alchemy RPC using zero-value transaction data payloads. This architecture establishes an immutable, non-repudiable forensic audit trail while maintaining an optimized PostgreSQL sliding-window storage policy capped at 100 entries per user.

Key Technical Contributions

The key technical contributions presented in this work are summarized as follows:

- **2D Geodetic Kalman Filtering & State-Machine Triage:** We deploy a client-side geodetic Kalman filter balancing measurement noise covariance against pedestrian kinematic process noise, integrated with a 3-state perimeter engine (SAFE -> CAUTION -> DANGER) that eliminates duplicate network pings and minimizes blockchain transaction overhead.
- **Kinematic Feature Engineering & Deep GRU Anomaly Engine:** We design a PyTorch GRU Autoencoder (64 -> 32 -> 64) operating on 31-sample sequence windows transformed into 4D kinematic vectors, evaluating trajectory irregularities against a mathematically calibrated reconstruction threshold (MSE >= 0.55).
- **Cascading Multi-Model LLM Risk Reasoning Pipeline:** We implement an automated LangChain fallback chain spanning four open-weight foundation models to convert raw spatial metrics into structured, schema-validated JSON threat reports for first responders.
- **Decentralized Forensic Ledger & Storage Optimization:** We integrate zero-value EVM blockchain transactions carrying SHA-256 telemetry digests to guarantee cryptographic non-repudiation, backed by a PostgreSQL sliding-window retention mechanism capping historical pings to 100 entries per user.

 This work is licensed under a Creative Commons Attribution 4.0 International License

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.151181

II. RELATED WORK

The development of a robust real-time tourist monitoring and incident response architecture requires synthesizing critical advances across four primary domain pillars: spatial-temporal trajectory mining, deep sequential autoencoders for trajectory anomaly detection, perimeter-based geofencing systems, and decentralized forensic ledger integration. This section categorizes and analyzes existing literature across these domain pillars to highlight structural limitations in current state-of-the-art implementations and establish the motivation for our proposed system.

A. Spatial-Temporal Trajectory Mining and Kinematic Profiling

Trajectory data mining traditionally focuses on processing sequences of geographical coordinate pairs to model human mobility behaviors. Zheng established foundational frameworks for spatial trajectory analytics, dividing trajectory processing into three major subfields: mobility pattern mining, trajectory outlier detection, and location prediction. However, traditional spatial analysis techniques operate on raw, unfiltered GPS coordinate strings, making them highly susceptible to positioning noise, indoor multipath jumps, and satellite signal loss.

To reduce false alarms caused by positioning noise, Kennedy and Züfle demonstrated that human spatial movements are more accurately classified by constructing dynamic kinematic profiles—evaluating velocity distributions, acceleration limits, and heading changes—rather than relying solely on raw spatial coordinates. Expanding on kinematic feature modeling, Tang et al. proposed a Gated Recurrent Unit (GRU) framework for abnormal trajectory classification. Their evaluation proved that recurrent neural architectures achieve far higher feature extraction accuracy on sequential spatiotemporal data compared to classical density-based clustering models such as k-Nearest Neighbors (k-NN) or DBSCAN. Nevertheless, while these kinematic feature models excel at offline trajectory classification, they lack real-time integration with dynamic client-side geofencing perimeters and automated emergency response workflows.

B. Deep Sequence-to-Sequence Autoencoders for Anomaly Detection

Unsupervised sequence-to-sequence autoencoder architectures have become the standard paradigm for temporal sequence reconstruction and anomaly evaluation. Cho et al. introduced the Gated Recurrent Unit (GRU), proving that its dual-gating mechanism—comprising reset and update gates—effectively captures long-term temporal dependencies while requiring significantly lower parameter and computational complexity compared to standard Long Short-Term Memory (LSTM) networks.

Applying recurrent units to sequence reconstruction, Malhotra et al. established that sequence-to-sequence autoencoders trained exclusively on nominal sensor streams yield elevated reconstruction errors when processing anomalous or out-of-distribution patterns. Zhou and Paffenroth further validated deep robust autoencoders for high-dimensional anomaly detection, demonstrating that forcing hidden representations through a low-dimensional latent space bottleneck successfully filters transient sensor noise while preserving core structural sequence dynamics. In industrial IoT and edge-computing contexts, Konatham et al. and De Vita et al. benchmarked deep recurrent autoencoders against cyber-physical fault streams, confirming that sequence reconstruction models achieve exceptional precision and ultra-low false alarm rates. However, current deep learning implementations operate almost exclusively as standalone anomaly classifiers, failing to integrate contextual AI reasoning engines that translate anomaly scores into actionable risk intelligence.

C. Geofencing Architectures and Decentralized Forensic Logging

Perimeter-based spatial monitoring systems play a key role in location-based services and regional security. Gavade et al. introduced a tourist safety framework incorporating static geofencing perimeters and digital identity authentication. Although their architecture laid valuable groundwork for tourism security, relying on continuous backend perimeter polling introduces high network latency and endpoint saturation when users remain stationary inside designated perimeters.

To ensure absolute auditability and data integrity during emergency incidents, decentralized blockchain ledgers have been adapted for immutable forensic recording. Wood and Buterin established the core principles of Ethereum state transitions, smart contract execution, and decentralized consensus. In public safety applications, publishing telemetry digests to EVM-compatible blockchains guarantees cryptographic non-repudiation, preventing post-incident data alteration by third parties. However, streaming high-frequency GPS coordinate pings directly to public blockchain ledgers incurs prohibitive network transaction gas fees.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.151181

D. Summary of Proposed Advances

Existing frameworks in literature address isolated aspects of tourist safety—such as kinematic profiling, deep sequence anomaly detection, or decentralized logging—in separate implementations. As outlined in Table I, our proposed framework unifies these domain pillars into a single end-to-end architecture. By combining client-side 2D Geodetic Kalman filtering, a 3-state geofence state-machine, a PyTorch GRU Autoencoder (64 -> 32 -> 64 bottleneck), a cascading multi-model LLM risk evaluation engine, and zero-value EVM cryptographic audit logging, the proposed system provides a complete, scalable, and tamper-proof solution for emergency dispatching.

TABLE I: COMPARATIVE ARCHITECTURAL ANALYSIS OF TOURIST SAFETY SYSTEMS

Framework	Kinematic Smoothness	GRU Engine	State Machine	LLM Chain	EVM Log
Zheng et al.	No	No	Static	No	No
Tang et al.	No	Yes	No	No	No
Kennedy and Zufle	Yes	Stat	No	No	No
Konatham et al.	No	Hybrid	No	No	No
Gavade et al.	No	No	Static	No	ID only
Proposed System	**2D Kalman**	**GRU - AE**	**3 - State**	**Multi - LLM**	**SHA – 256**

III. SYSTEM ARCHITECTURE AND METHODOLOGY

A. Client-Side Kinematic Filtering and Geofence State Machine

Raw mobile GPS sensors exhibit significant variance due to atmospheric interference and indoor multipath reflection. To stabilize trajectory plotting, the client application incorporates an accuracy-aware 2D Geodetic Kalman Filter operating on strict 5-second sampling intervals. Let the system state vector be defined as:

$$x_k = \left[\text{lat}_k, \text{lon}_k, v_{\text{lat},k}, v_{\text{lon},k}\right]^T$$

The filter updates the location estimate by continuously balancing device-reported measurement noise covariance (R, derived from HTML5 Geolocation API accuracy) against an assumed pedestrian kinematic process noise (Q).
In parallel, a spatial state-machine triages user movement relative to custom circular or polygon geofences drawn by authorities or users. The client engine automatically computes a +50-meter geodesic caution buffer surrounding danger zones. To eliminate redundant server pings and optimize blockchain gas fees, coordinates are transmitted to the backend microservice exclusively upon state transitions (SAFE → CAUTION → DANGER), rather than continuously pinging the server while a user remains stationary inside a zone.

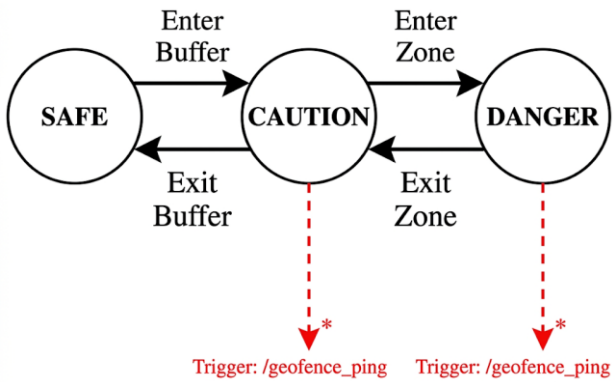


Fig. 1. Geofence State Machine transition logic utilized to suppress redundant API telemetry and minimize blockchain transaction overhead.

B. Spatial-Temporal Feature Engineering

The FastAPI microservice backend maintains a sliding temporal window of 31 consecutive GPS coordinate samples for each active user. Because kinematic transition metrics are calculated between adjacent temporal points, the 31 raw GPS

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.151181

coordinates yield a sequence of exactly $T = 30$ interval steps. To standardize raw location inputs for deep learning inference, geographical coordinates are transformed into a 4-dimensional kinematic feature matrix $\mathbf{X} \in \mathbb{R}^{30\times4}$.

For two consecutive points $(\text{lat}_{i-1}, \text{lon}_{i-1})$ and $(\text{lat}_i, \text{lon}_i)$ separated by sample interval Δt_i (where $i = 1,2,\ldots,30$), the physical distance d_i is calculated using the vectorized Haversine formula. From this, four core kinematic features are extracted per interval step:

1. **Speed (v_i):** Calculated as $v_i = \frac{d_i}{\Delta t_i}$.
2. **Acceleration (a_i):** Calculated as linear speed change $a_i = \frac{v_i - v_{i-1}}{\Delta t_i}$.
3. **Directional Deviation ($\Delta\theta_i$):** The absolute compass bearing deviation between consecutive segments, defined as $\Delta\theta_i = \min(|\theta_i - \theta_{i-1}|, 360^\circ - |\theta_i - \theta_{i-1}|)$.
4. **Stationary State Flag (S_i):** A binary indicator where $S_i = 1$ if velocity $v_i \leq 0.3$ m/s for a sustained duration of ≥ 60 seconds.

This feature matrix is normalized using a pre-fitted StandardScaler to ensure uniform gradient propagation during network inference.

C. Deep GRU Autoencoder Dynamics

At the core of the trajectory anomaly detection pipeline is a PyTorch Gated Recurrent Unit (GRU) Autoencoder structured as a bottleneck sequence-to-sequence model. The architecture compresses the $T \times d$in input tensor down to a rigid latent space before attempting full trajectory reconstruction.

The encoder GRU processes the input sequence $\boldsymbol{X} \in \mathbb{R}^{B\times30\times4}$, updating its hidden state $\boldsymbol{h}_t \in \mathbb{R}^{64}$ at each timestep via standard GRU gating mechanisms (reset gate $\boldsymbol{r}_t$ and update gate $\boldsymbol{z}_t$). The final hidden state $\boldsymbol{h}_T$ encapsulates the entire 30-step temporal context.

To enforce structural generalization, $\boldsymbol{h}_T$ is linearly projected into a lower-dimensional bottleneck:

$$\boldsymbol{z}_{latent} = \boldsymbol{W}_{enc}\boldsymbol{h}_T + \boldsymbol{b}_{enc} \in \mathbb{R}^{32}$$

This 32-dimensional vector $\boldsymbol{z}_{latent}$ is then expanded back to the decoder's hidden dimension space ($\mathbb{R}^{64}$). To facilitate sequence reconstruction without autoregressive teacher-forcing, the 64-dimensional context vector is broadcast across all $T = 30$ sequence steps:

$$DecoderInput = Repeat(\boldsymbol{W}_{dec}\boldsymbol{z}_{latent} + \boldsymbol{b}_{dec}, T) \in \mathbb{R}^{B\times30\times64}$$

The decoder GRU processes this repeated sequence, and a final linear layer maps the 64-dimensional outputs back to the original 4-dimensional kinematic feature space, producing the reconstructed sequence $\hat{\boldsymbol{X}} \in \mathbb{R}^{B\times30\times4}$.

During inference, the system calculates the Mean Squared Error (MSE) between original input matrix $\mathbf{X}$ and reconstructed matrix $\hat{\mathbf{X}}$. Because the autoencoder is exclusively trained on nominal, safe trajectory data, irregular kinematic profiles—such as sudden running, erratic direction changes, or panic stops—produce high reconstruction loss. An anomaly is officially flagged if:

$\text{MSELoss}(\mathbf{X}, \hat{\mathbf{X}}) \geq 0.55$

 This work is licensed under a Creative Commons Attribution 4.0 International License

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.151181

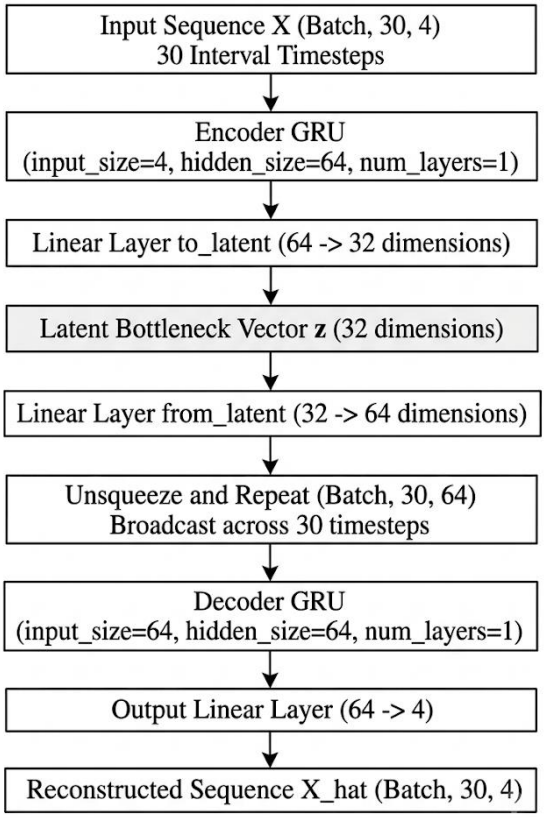


Fig. 2. Data flow and dimensionality bottleneck through the PyTorch GRU Autoencoder architecture.

D. LLM Risk Evaluation and Decentralized Logging

When a geofence state transition occurs or the GRU MSE threshold (≥ 0.55) is breached, the incident is routed to a multi-model LangChain fallback pipeline. The system queries openai/gpt-oss-120b as the primary engine, cascading to smaller Qwen and OSS models (qwen/qwen3.8-27b → openai/gpt-oss-20b → openai/gpt-oss-safeguard-20b) if rate limits or timeouts occur. The LLM consumes the trajectory metadata, anomaly scores, and zone coordinates to generate a strictly formatted JSON risk report detailing the threat level (LOW, MEDIUM, HIGH, CRITICAL) and actionable emergency protocols.

Simultaneously, the backend concatenates the user identity, latitude, longitude, and an ISO-8601 timestamp into a raw telemetry string. This string is hashed using SHA-256 and broadcast to an Ethereum/EVM-compatible network as a zero-value transaction payload via Alchemy RPC. The resulting transaction hash is stored in a highly optimized PostgreSQL relational database—utilizing a sliding-window retention script that caps location history at 100 entries per user—ensuring a computationally lightweight, tamper-proof forensic audit trail.

IV. Experimental Results and Performance Analysis

A. Model Training and Convergence Dynamics

The GRU Autoencoder minimizes the Mean Squared Error (MSE) reconstruction loss during training on nominal trajectory data. The network was trained using a phased approach to ensure stable weight convergence and prevent gradient instability within the recurrent hidden states.

As illustrated in Fig. 3, the initial training phase captures the broad structural dynamics of the 4D kinematic inputs. The sequence reconstruction error experiences a rapid initial descent, converging from an initial MSE of approximately 1.78 down to 1.18 by the tenth epoch.

 This work is licensed under a Creative Commons Attribution 4.0 International License

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026

DOI: 10.17148/IJARCCE.2026.151181



Fig. 3. Phase 1 Training: Initial loss convergence of the GRU Autoencoder across 10 epochs.

Following this initial structural alignment, the model underwent a secondary fine-tuning phase (Fig. 4). During this phase, the latent bottleneck ($\mathbf{z} \in \mathbb{R}^{32}$) optimized its representations of fine-grained pedestrian kinematic profiles. This phase demonstrates smooth asymptotic convergence, starting at an MSE of approximately 0.123 and stabilizing at a final training loss of 0.094.

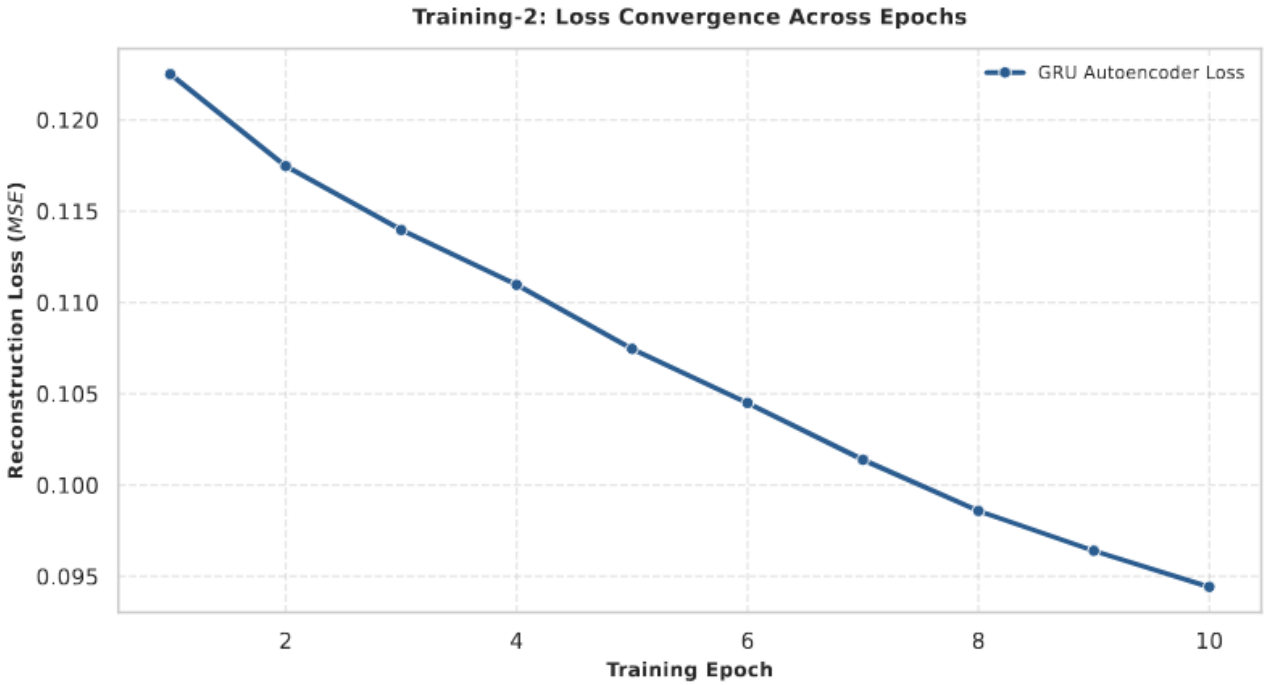


Fig. 4. Phase 2 Training: Fine-tuned loss convergence demonstrating stable asymptotic behavior.

B. Test Evaluation and Threshold Calibration

To evaluate the model's out-of-sample generalization, the fully trained GRU Autoencoder was tested against a withheld dataset comprising nominal, safe tourist trajectories. The model achieved a final test reconstruction loss of 0.14064.

 This work is licensed under a Creative Commons Attribution 4.0 International License



International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026
DOI: 10.17148/IJARCCE.2026.151181

The operational objective of the anomaly engine is to flag genuine erratic spatial behavior without triggering false alarms due to standard atmospheric GPS drift. By establishing a static calibrated threshold of $\tau = 0.55$, the system maintains a robust margin above the baseline nominal test error. Setting an empirical threshold τ significantly above nominal validation errors is a standard methodology to prevent false alarms.

The delta between the test baseline and the decision boundary ($\Delta \approx 0.41$) acts as a critical kinematic buffer. It ensures that minor velocity fluctuations or Kalman filter adjustments are smoothly reconstructed by the autoencoder, whereas true distress patterns—such as sudden acceleration, erratic directional deviation, or panic stopping—produce reconstruction errors that spike beyond the 0.55 threshold, successfully triggering the LangChain LLM evaluation pipeline.

V. Conclusion And Future Work

This paper presented an end-to-end, multi-tier tourist safety monitoring framework designed to overcome the critical limitations of static geofencing and centralized incident logging. By integrating a client-side 2D Geodetic Kalman Filter with a PyTorch GRU Autoencoder, the proposed system effectively isolates genuine trajectory anomalies from benign kinematic and atmospheric GPS noise. The rigid $64 \rightarrow 32 \rightarrow 64$ bottleneck architecture successfully compressed spatial-temporal sequences, achieving a stable nominal test reconstruction loss of $\text{MSE} = 0.140$. Establishing a calibrated anomaly threshold of $\tau = 0.55$ provides a mathematically reliable decision boundary to trigger automated emergency protocols without saturating endpoints with false alarms.

Furthermore, integrating a cascading LangChain multi-model LLM pipeline ensures that first responders receive structured, context-aware risk evaluations rather than raw spatial coordinate arrays. By anchoring incident telemetry to an EVM-compatible blockchain via zero-value transactions, the architecture guarantees cryptographic non-repudiation and data integrity post-breach.

Future work will focus on scaling the system from single-agent trajectory profiling to macro-level crowd handling. By cross-referencing individual GRU kinematic anomalies with spatial crowd density metrics within active geofences, the architecture will be extended to detect multi-agent emergency events—such as stampedes, mass panic, or synchronized perimeter breaches—enabling highly coordinated disaster response operations.

ACKNOWLEDGMENT

The authors express their sincere gratitude to the Department of Computer Engineering at Parikrama College of Engineering, Kashti, for providing the necessary computational infrastructure, research facilities, and academic guidance required to complete this work.

We also extend our heartfelt appreciation to our project guide, faculty members, and peer reviewers whose constructive feedback, insights, and continuous support greatly contributed to the refinement and successful implementation of this research paper.

REFERENCES

[1]. Y. Zheng, "Trajectory data mining: an overview", *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 6, no. 3, pp. 1-41, 2015.

[2]. R. L. Gavade, B. H. Ukarande, P. D. Kumbhar, R. S. Bharade, S. V. Redekar, and P. D. Patil, "Smart tourist safety monitoring and incident response system using artificial intelligence, geo-fencing, and blockchain-based digital identity", *International Journal of Innovative Research in Technology (IJIRT)*, vol. 12, no. 12, pp. 9670-9676, 2026.

[3]. R. E. Kalman, "A new approach to linear filtering and prediction problems", *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35-45, 1960.

[4]. G. Wood, "Ethereum: A secure decentralised generalised transaction ledger", *Ethereum Project Yellow Paper*, vol. 151, no. 2014, pp. 1-32, 2014.

[5]. B. Konatham, T. Simra, F. Amsaad, M. I. Ibrahem, and N. Z. Jhanjhi, "A secure hybrid deep learning technique for anomaly detection in iiot edge computing", *TechRxiv preprint*, vol. 1, no. 1, pp. 1-10, 2024.

[6]. L. Kennedy and A. Züfle, "Kinematic detection of anomalies in human trajectory data", *Proc. 32nd ACM SIGSPATIAL Int. Conf. Adv. Geogr. Inf. Syst.*, vol. 32, no. 1, pp. 1-4, 2024.

[7]. V. Buterin, "A next-generation smart contract and decentralized application platform", *Ethereum White Paper*, vol. 3, no. 37, pp. 1-32, 2014.

[8]. S. Ramírez, "FastAPI framework, high performance, easy to learn, fast to code, ready for production", *FastAPI Documentation*, vol. 1, no. 1, pp. 1-5, 2018.

 This work is licensed under a Creative Commons Attribution 4.0 International License



International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 11, November 2026
DOI: 10.17148/IJARCCE.2026.151181

[9]. K. Cho, B. van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using rnn encoder-decoder for statistical machine translation", *Proc. EMNLP*, vol. 14, no. 1, pp. 1724-1734, 2014.
[10]. P. Malhotra, A. Ramakrishnan, G. Anand, L. Vig, P. Agarwal, and G. Shroff, "LSTM-based encoder-decoder for multi-sensor anomaly detection", *ICML Workshop on Anomaly Detection*, vol. 1, no. 1, pp. 1-8, 2016.
[11]. G. Tang, H. Zhao, and B. Yu, "Low-cost and high-performance abnormal trajectory detection based on the GRU model with deep spatiotemporal sequence analysis in cloud computing", *Journal of Cloud Computing*, vol. 13, no. 1, pp. 53-65, 2024.
[12]. C. Zhou and R. C. Paffenroth, "Anomaly detection with robust deep autoencoders", *Proc. 23rd ACM SIGKDD Int. Conf. Knowl. Discov. Data Min.*, vol. 23, no. 1, pp. 665-674, 2017.
[13]. H. Chase, "LangChain: Building applications with LLMs through composability", *LangChain Framework Guide*, vol. 1, no. 1, pp. 1-12, 2022.
[14]. H. Touvron et al., "Llama 2: Open foundation and fine-tuned chat models", *arXiv preprint arXiv:2307.09288*, vol. 1, no. 1, pp. 1-77, 2023.
[15]. F. De Vita, G. Nocera, D. Bruneo, and S. K. Das, "A novel echo state network autoencoder for anomaly detection in industrial IoT systems", *IEEE Transactions on Industrial Informatics*, vol. 19, no. 8, pp. 8985-8994, 2023.

 This work is licensed under a Creative Commons Attribution 4.0 International License