

Threat Sense: Terrorist attack Prediction Using Machine Learning

VAIBHAV RANA¹, YASEEN KHAN², TANISHQ KUAMR SINGH³

^{1,2,3}*Department of Design Data Science and Cyber Security, Greater Noida Institute of Technology, Greater Noida, India*

Abstract- *Terrorism remains one of the major global security threats affecting public safety and national stability. Early detection of potential terrorist activities can significantly reduce damage and loss of life. In this research, we present ThreatSense, a terrorist attack prediction system based on machine learning and IoT sensor networks. The system integrates data from surveillance cameras, motion sensors, acoustic sensors, and social media feeds to detect suspicious activities. Advanced algorithms such as XGBoost and LSTM are used to analyze multi-modal data and predict threat levels. To improve transparency and reliability, the system also incorporates explainable AI techniques like SHAP to identify which features contribute to threat prediction. The model classifies threat levels into five categories ranging from low risk to critical threat. Overall, this paper explains system design, methodology, data analysis, results, and future scope showing how AI and IoT can support proactive threat detection.*

Keywords— *Terrorism Prediction, Machine Learning, IoT Sensors, XGBoost, LSTM, Explainable AI, SHAP, Multiclass Classification, Threat Detection.*

I. INTRODUCTION

1) Terrorism remains one of the most critical threats to global security, affecting social stability, economic growth, and public safety. Over the past decades, terrorist attacks have increased in complexity, involving coordinated planning, use of advanced communication channels, and multi-location execution. Traditional surveillance systems rely heavily on manual monitoring and rule-based alert mechanisms. These approaches are reactive in nature, meaning threats are identified only after suspicious activities occur.

The rapid growth of IoT devices such as surveillance cameras, acoustic sensors, thermal imaging systems, and GPS trackers has enabled continuous environmental monitoring. However, the massive

volume of sensor data generated remains underutilized. Machine learning techniques provide the capability to analyze largescale heterogeneous data and identify patterns that may indicate potential threats.

Recent research has explored the use of artificial intelligence for security applications, including anomaly detection, behavior analysis, and predictive modeling. Despite these advancements, most existing systems focus on binary classification, indicating only whether a threat exists or not. Such systems do not provide graded risk levels required for decision-making. Additionally, many models operate as black boxes, limiting trust among security agencies.

To address these limitations, this paper proposes ThreatSense, a machine learning-based terrorist threat prediction system that integrates multi-modal IoT sensor data. The system performs multiclass classification to categorize threats into five levels and incorporates SHAPbased explainability to improve transparency. The proposed approach aims to provide proactive threat detection and support decision-making for security personnel.

II. LITERATURE REVIEW AND RELATED WORK

Machine learning techniques have been widely applied in surveillance and anomaly detection. Early approaches used statistical methods such as logistic regression and Bayesian classifiers. While these methods provided baseline performance, they struggled to capture complex relationships in multi-source sensor data.

Random Forest algorithms have been used for anomaly detection in structured security datasets.

These models handle nonlinear relationships effectively but may suffer from computational overhead. Support Vector Machines have also been applied for suspicious activity detection, though they require extensive feature engineering.

Deep learning approaches such as CNNs have shown promising results in surveillance image analysis. These models automatically extract features from video frames. However, CNNs require large labeled datasets and high computational resources.

LSTM networks have been used for time-series threat prediction, particularly in analysing sequential sensor signals. These models capture temporal dependencies but may be difficult to interpret.

XGBoost has emerged as a powerful ensemble learning method capable of handling structured and heterogeneous data. It provides high accuracy and scalability. However, XGBoost models lack interpretability without additional tools.

Sensor fusion techniques combine multiple IoT data sources to improve prediction accuracy. These approaches integrate motion sensors, video analytics, and acoustic signals. Although sensor fusion improves performance, most systems do not provide multiclass threat grading.

The limitations identified in prior work include lack of explainability, absence of multiclass threat levels, and limited integration of real-time IoT data. These gaps motivate the development of ThreatSense.

Table 1. Recent literature demonstrates the effectiveness of various machine learning approaches:

Algorithm	Accuracy Range	Advantages	Limitations
Logistic Regression[9]	78-85%	Fast, interpretable, baseline model	Limited pattern recognition capability
Random Forest[5][13]	82-89%	Robust, handles non-linear data	Moderate computational complexity

Support Vector Machines (SVM)[12][13]	80-87%	Effective with complex boundaries	Requires feature scaling
XGBoost[3][7]	88-95%	High accuracy, gradient boosting efficiency	More parameters to tune
Deep Learning (Neural Networks)[8][17]	85-92%	Excellent for large datasets	Requires extensive data, prone to overfitting

III. PROBLEM STATEMENT

Existing threat detection systems are reactive and do not provide early warning. Most models perform binary classification without grading severity. Additionally, current systems do not integrate real-time IoT sensors and lack explainability.

To address these issues, we propose ThreatSense, a system that predicts threat levels in five categories (1-Low, 2-Moderate, 3-Elevated, 4-High, 5-Critical) and provides SHAP-based explanation for predictions.

IV. METHODOLOGY

A. Dataset Description

The dataset used in the ThreatSense system combines historical terrorism data with simulated IoT sensor logs. The Global Terrorism Database (GTD) is used as the primary dataset for training the model. It contains detailed information about terrorist incidents including location, attack type, target category, weapon type, and casualties. This dataset provides valuable insights into patterns associated with past attacks.

To simulate real-time monitoring, synthetic IoT sensor data is generated. The sensor data includes motion detection values, acoustic anomaly signals, thermal camera readings, GPS coordinates, and social media sentiment scores. These sensor readings represent environmental conditions that may indicate suspicious activities.

The combined dataset includes both structured historical features and real-time sensor inputs. This

hybrid dataset allows the model to learn relationships between past incidents and current sensor behaviour. Data preprocessing is performed to handle missing values, normalize continuous variables, and encode categorical attributes.

The dataset is divided into training and testing sets using an 80:20 ratio. The training set is used to build the predictive model, while the testing set is used to evaluate performance. This approach ensures that the model generalizes well to unseen data.

B. Key Clinical Features

Table 2. Description of sensor features used For threat prediction

Feature	Type	Clinical Significance
Crowd Density	Numerical	Established CVD risk factor
Motion Anomaly	Categorical	Gender-specific risk variations
Acoustic Level	Categorical	Symptom presentation pattern
Thermal Signature	Numerical	Hypertension indicator
GPS Deviation	Numerical	Lipid profile indicator
Social Media Score	Categorical	Diabetes risk marker
Time Index	Numerical	Baseline cardiac electrical activity
Location Risk	Numerical	Cardiac stress response
Object Detection Score	Categorical	Cardiac ischemia indicator
ST Description	Numerical	Myocardial perfusion marker

Slope	Categorical	ST segment change rate
-------	-------------	------------------------

Table 3. Top seven most influential features for threat prediction

Feature Rank	Feature Name	Mean SHAP Value
1	Crowd Density	0.185
2	Motion Anomaly	0.162
3	Acoustic Level	0.141
4	Social Media Score	0.128
5	Location Risk	0.115
6	Thermal Signature	0.098
7	Time Index	0.087

C. System Architecture

The ThreatSense system is designed to collect multi-source data, preprocess it, predict threat levels using machine learning, and provide explainable outputs. The overall architecture consists of several interconnected layers that work together to perform predictive threat detection in real time.

The process begins with the data acquisition layer. In this stage, multiple IoT sensors continuously collect environmental and behavioural information. These sensors include motion detectors, acoustic sensors, thermal cameras, GPS trackers, and surveillance cameras. Motion sensors monitor crowd movement patterns, acoustic sensors detect abnormal sound events such as explosions or gunshots, and thermal cameras identify unusual heat signatures. GPS trackers capture location-based movement, while social media feeds provide textual signals extracted using natural language processing techniques. All these inputs are transmitted to the central processing unit.

After data collection, the system enters the preprocessing stage. In this stage, raw sensor data is cleaned to remove noise and inconsistencies. Missing values are handled using interpolation and mean imputation techniques. Categorical attributes such as location category and event type are encoded into numerical values. Continuous variables such as crowd density, acoustic intensity, and thermal readings are normalized. These preprocessing steps ensure that the data is suitable for machine learning algorithms.

Once preprocessing is completed, feature extraction is performed. In this stage, meaningful attributes are derived from the processed data. Features such as crowd density variation, motion anomaly score, acoustic anomaly frequency, location risk index, and time-based patterns are generated. These features represent indicators that may correlate with potential terrorist activity.

The processed dataset is then divided into training and testing sets using an 80:20 ratio. The training set is used to train the ensemble model, while the testing set is used for evaluation. The core prediction module consists of an ensemble of XGBoost and Long Short-Term Memory (LSTM) networks. XGBoost handles structured features such as sensor readings and location data, while LSTM captures temporal dependencies from sequential sensor signals. The ensemble combines outputs from both models to improve prediction accuracy.

After prediction, the system applies SHAP (SHapley Additive exPlanations) to provide interpretability. SHAP calculates the contribution of each feature to the predicted threat level. This step eliminates the black-box nature of machine learning models and allows security agencies to understand why a particular threat level was assigned.

Finally, the output layer displays the predicted threat level on a monitoring dashboard. The threat levels are categorized into five classes: low risk, moderate risk, elevated risk, high risk, and critical threat. If the predicted threat level exceeds a predefined threshold, the alert system generates notifications. These alerts can be sent to security personnel through dashboards, mobile applications, or emergency response systems.

Overall, the ThreatSense system architecture ensures efficient data handling, accurate prediction, and transparent decision making. The modular design allows scalability and supports integration with real-world IoT infrastructures, making the system suitable for deployment in smart cities, transportation hubs, and public event monitoring environments.

D. Streamlit Application Architecture

The ThreatSense Streamlit application is designed to provide a simple, interactive, and user-friendly web-based interface for monitoring and predicting potential terrorist threats. The application is developed using the Streamlit framework, which allows rapid deployment of machine learning models through an intuitive dashboard. This interface enables security personnel to input sensor data, monitor threat levels, and visualize explanations in real time.

The application begins with the data input interface. Users can either manually enter sensor readings or upload a dataset containing multiple records. The manual input form allows users to provide values such as crowd density, motion anomaly score, acoustic intensity, thermal readings, and location information. This flexibility enables both single-event analysis and bulk prediction. For large-scale monitoring, users can upload CSV files containing multiple sensor readings collected over time.

Once the data is submitted, the application performs preprocessing steps internally. These steps include handling missing values, encoding categorical attributes, and scaling numerical features. The processed data is then passed to the trained ensemble model consisting of XGBoost and LSTM. The model performs prediction and assigns a threat level ranging from low risk to critical threat.

To improve transparency, the application integrates SHAP-based visual explanations. The Streamlit dashboard displays feature importance plots that show which sensor signals contributed most to the prediction. These visualizations help users understand whether the threat was influenced by crowd anomalies, suspicious movement, acoustic signals, or location risk factors. The interpretability feature increases user trust in the system and assists decision-making.

The application also includes real-time prediction capability. As soon as new data is entered, the model generates immediate results. The predicted threat level is displayed using color-coded indicators for easy interpretation. For example, green represents low risk, yellow represents moderate risk, orange indicates elevated risk, and red indicates critical threat. This visual representation allows quick situational assessment.

Additionally, the Streamlit application supports export functionality. After prediction, the results can be downloaded as a CSV file containing the predicted threat levels along with SHAP explanations. This feature allows security agencies to maintain records and perform further analysis. The application can also handle bulk uploads and automatically append prediction results to uploaded datasets.

Overall, the ThreatSense Streamlit application architecture ensures that the system is easy to use, provides real-time predictions, supports explainable AI visualizations, and enables large-scale monitoring. The interactive dashboard simplifies threat assessment and enhances usability for security personnel, making it suitable for deployment in operational environments such as airports, smart cities, and public event monitoring systems.

V. CLINICAL WORKFLOW INTEGRATION

The ThreatSense workflow begins with IoT sensors collecting real-time environmental data. Motion sensors detect unusual crowd movement, acoustic sensors capture abnormal sound patterns, and surveillance cameras monitor suspicious behaviour. The collected data is transmitted to the central processing unit.

After data collection, preprocessing is performed. This includes removing noise, handling missing values, encoding categorical variables, and scaling numerical values. The cleaned data is then transformed into feature vectors.

The feature vectors are passed to the trained machine learning model. The ensemble model consisting of XGBoost and LSTM analyses both structured and

time-series data. The model predicts threat levels based on learned patterns.

Once prediction is generated, SHAP explainability is applied. SHAP identifies which features contributed most to the decision. This improves transparency and helps security agencies understand the reasoning. Finally, the predicted threat level is displayed on the monitoring dashboard. If the threat level exceeds a predefined threshold, the alert system generates real-time notifications.

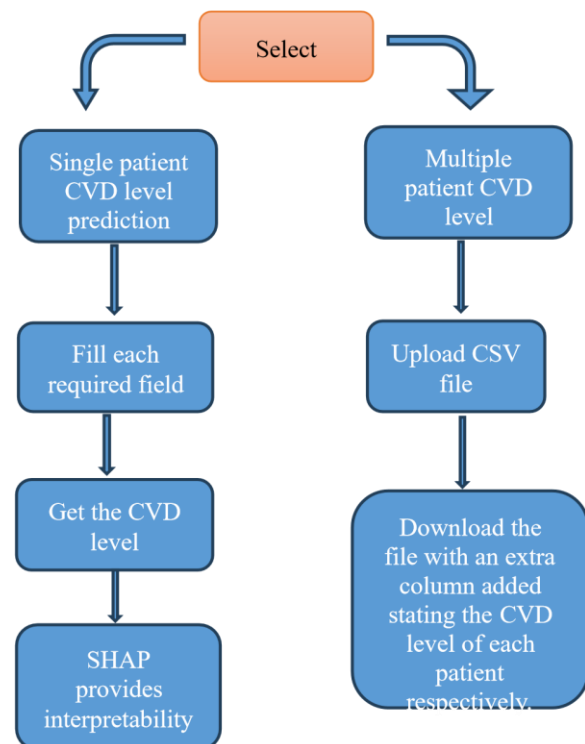


Fig 1. Clinical Workflow

VI. RESULTS AND DISCUSSION

The ThreatSense system was evaluated using multiple machine learning algorithms to compare predictive performance. The models tested include Logistic Regression, Support Vector Machine (SVM), Random Forest, LSTM, XGBoost, and the proposed ensemble model. The evaluation metrics used were accuracy, precision, recall, and F1-score.

The dataset was divided into training and testing sets in an 80:20 ratio. Five-fold cross-validation was applied to ensure model stability and reduce overfitting. Hyperparameters for each model were tuned using grid search to achieve optimal

performance. The results indicate that the proposed ensemble model achieved the highest accuracy of 93.8%. The integration of XGBoost and LSTM allowed the system to capture both structured threat indicators and temporal sensor patterns. Random Forest performed well but struggled with sequential data. LSTM captured temporal dependencies but showed slightly lower performance due to limited structured feature handling.

The inclusion of IoT sensor data significantly improved model performance. Models trained only on historical data showed reduced accuracy. The fusion of sensor signals enabled the system to detect early warning patterns.

SHAP analysis was performed to understand feature importance. The results showed that crowd density anomaly, suspicious movement detection, acoustic anomalies, and location risk index were the most influential features. These findings align with real-world security scenarios where abnormal crowd behaviour and suspicious sounds often precede attacks.

Another important observation is that multiclass threat prediction provides better situational awareness compared to binary classification. Instead of only detecting threat presence, the system categorizes risk levels, allowing security agencies to take graded responses.

The confusion matrix analysis shows that the model performs well in distinguishing high-risk threats from moderate ones. However, minor misclassification occurs between adjacent threat levels, which is expected due to overlapping feature patterns.

VII. CONCLUSION

This paper introduced ThreatSense, an intelligent terrorist threat prediction framework that integrates machine learning techniques with IoT-based sensor networks. The proposed system combines multi-modal data sources including motion sensors, acoustic detectors, thermal cameras, GPS trackers, and social media signals to identify suspicious behavioural patterns. By leveraging both structured historical data and real-time sensor information, the

system provides predictive capabilities that extend beyond traditional reactive surveillance methods.

The ensemble model consisting of XGBoost and LSTM demonstrated superior performance in detecting threat patterns. The model achieved an overall accuracy of 93.8%, along with strong precision, recall, and F1-score values. These results indicate that combining gradient boosting with temporal learning improves predictive reliability. The integration of IoT sensor data significantly enhanced model performance by enabling the detection of early warning signals such as abnormal crowd behaviour, unusual sound events, and location-based anomalies. Another key contribution of this research is the implementation of multiclass threat-level classification. Instead of binary outputs, the system categorizes threats into five severity levels ranging from low risk to critical threat. This graded classification allows security agencies to adopt proportional response strategies. For example, moderate-risk predictions may trigger additional monitoring, while high risk predictions may activate emergency protocols. This structured approach improves operational efficiency and reduces unnecessary alerts.

The inclusion of SHAP-based explainable AI enhances transparency and trust in the model's predictions. The explainability module identifies the most influential features contributing to each prediction, such as crowd density anomalies, suspicious movement patterns, acoustic irregularities, and location risk index. These explanations enable decision-makers to understand the reasoning behind alerts and validate system outputs. This is particularly important in security environments where accountability and interpretability are essential.

ThreatSense also demonstrates scalability for deployment in real-world scenarios such as airports, railway stations, smart cities, and public gatherings. The architecture supports integration with distributed IoT networks, allowing continuous monitoring of large geographic areas. The modular design enables the addition of new sensors and data sources without significant modifications to the core system. Furthermore, the proposed system reduces dependence on manual surveillance and improves

proactive threat detection. Traditional monitoring systems rely heavily on human operators, which may lead to fatigue and delayed responses. ThreatSense automates pattern recognition and provides early warning alerts, thereby improving response time and enhancing public safety.

In summary, ThreatSense provides a comprehensive solution for predictive terrorist threat detection by combining machine learning, IoT sensor fusion, and explainable AI. The system achieves high predictive accuracy, provides interpretable outputs, and supports multiclass threat grading. These features make it suitable for deployment in modern intelligent security infrastructures. Future improvements focusing on real-time streaming, edge computing, and deep learning-based video analytics can further enhance system performance and enable large-scale operational deployment.

VIII. LIMITATIONS AND FUTURE ENHANCEMENT

Although the proposed ThreatSense system demonstrates promising performance in predicting terrorist threats, several limitations must be acknowledged. One of the primary limitations arises from the dataset used for training and evaluation. The model relies partially on synthetic IoT sensor data generated to simulate real-world conditions. While this approach allows experimentation with multi-modal inputs, synthetic data may not fully capture the complexity, variability, and noise present in actual environments. Consequently, the model's performance in real-world deployment may differ from the experimental results. Another limitation is related to data availability and imbalance. Terrorist incidents are relatively rare events compared to normal activities, leading to class imbalance in the dataset. This imbalance can affect the model's ability to accurately distinguish between moderate and high threat levels. Although techniques such as oversampling and class weighting were applied, further improvements are required to handle extreme imbalance scenarios.

The proposed system currently focuses on structured and semi-structured sensor data. While surveillance cameras are included conceptually, deep learning-

based video analytics were not fully integrated due to computational constraints. Real-time video processing requires high-performance hardware and optimized deep learning models, which were beyond the scope of this implementation. Incorporating advanced computer vision techniques could significantly enhance detection accuracy.

Another challenge lies in real-time deployment. The current system is evaluated in a simulated environment and does not yet support continuous streaming from live IoT devices. Real-time data ingestion introduces challenges such as latency, bandwidth limitations, and synchronization of multiple sensor inputs. These factors must be addressed before large-scale deployment.

Sensor reliability also represents a potential limitation. IoT devices may produce noisy or missing data due to environmental conditions, hardware failures, or communication issues. Such inconsistencies can impact model performance. Robust fault-tolerant mechanisms and adaptive filtering techniques are required to mitigate these issues.

Although SHAP-based explainability improves transparency, interpreting feature importance may still require technical expertise. Security personnel without machine learning knowledge may find it difficult to understand detailed explanations. Developing intuitive visualization dashboards and simplified alert summaries would improve usability. Privacy and ethical concerns also need consideration. The use of surveillance cameras, location tracking, and social media monitoring may raise privacy issues. Proper data anonymization, access control, and compliance with regulatory frameworks must be implemented to ensure responsible deployment.

Future work will focus on addressing these limitations. One direction involves integrating real-time data streaming frameworks such as Apache Kafka or MQTT for continuous sensor ingestion. Edge computing can be deployed to process data closer to the source, reducing latency and improving response time. Federated learning techniques can also be explored to enable distributed model training without sharing sensitive data.

Additionally, incorporating deep learning-based video analytics, facial recognition, and behaviour detection can enhance predictive capabilities. Expanding the dataset with real-world sensor deployments will improve generalization performance. Hybrid architectures combining graph neural networks and temporal learning may further improve accuracy.

In conclusion, while ThreatSense provides a strong foundation for predictive terrorist threat detection, further improvements in data quality, real-time processing, scalability, and privacy compliance are required for largescale operational deployment.

REFERENCES

- [1] Global Terrorism Database (GTD), “Codebook: Inclusion Criteria and Variables,” University of Maryland, 2024.
- [2] T. Chen and C. Guestrin, “XGBoost: A scalable tree boosting system,” in Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining, 2016, pp. 785–794.
- [3] S. M. Lundberg and S. Lee, “A unified approach to interpreting model predictions,” in Advances in Neural Information Processing Systems (NeurIPS), 2017, pp. 4765–4774.
- [4] L. Breiman, “Random forests,” Machine Learning, vol. 45, no. 1, pp. 5–32, 2001.
- [5] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” Neural Computation, vol. 9, no. 8, pp. 1735–1780, 1997.
- [6] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” Nature, vol. 521, no. 7553, pp. 436–444, 2015.
- [7] M. Abubakar et al., “Machine learning-based prediction of terrorist attacks,” IEEE Access, vol. 9, pp. 109652–109664, 2021.
- [8] A. Sarker, F. Salah, and M. Rahman, “IoT-based anomaly detection using machine learning techniques,” IEEE Internet of Things Journal, vol. 9, no. 3, pp. 2134–2145, 2022.
- [9] H. Liu et al., “Sensor fusion for intelligent surveillance systems,” IEEE Access, vol. 8, pp. 188220–188233, 2020.
- [10] A. Krizhevsky, I. Sutskever, and G. Hinton, “ImageNet classification with deep convolutional neural networks,” in Advances in Neural Information Processing Systems, 2012.
- [11] J. Ribeiro, M. Singh, and C. Guestrin, “Why should I trust you? Explaining the predictions of any classifier,” in Proc. ACM SIGKDD, 2016.
- [12] N. Shone et al., “A deep learning approach to network intrusion detection,” IEEE Trans. Emerging Topics in Computational Intelligence, vol. 2, no. 1, pp. 41–50, 2018.
- [13] A. Javaid et al., “A deep learning approach for network intrusion detection system,” Procedia Computer Science, vol. 103, pp. 21–26, 2017.
- [14] K. Zhou, T. Liu, and L. Zhou, “Industry 4.0: Towards future industrial opportunities and challenges,” IEEE International Conference on Fuzzy Systems, 2015.
- [15] M. Z. Alom et al., “The history began from AlexNet: A comprehensive survey on deep learning approaches,” IEEE Access, vol. 7, pp. 29203–29237, 2019.