Nrityanjali Multidisciplinary International Research Journal

Volume 1, Issue 1, October 2025

NMIRJ **E-ISSN:** Applied **Website:** https://nrityanjaliresearchjournal.com/home/ **Email:** info@nrityanjaliresearchjournal.com

"Smart Surveillance Systems for Real-Time Crowd Management during Public Emergencies"

Pushpendra Prajapati

Computer Science & Engineering Department,

IFTM University, Gwalior

Abstract

Real-time crowd monitoring has become an essential aspect of emergency management in the context of increasing public gatherings and urban congestion. This study proposes a deep learning-based modular framework for detecting and analyzing crowd behavior in real-time to prevent stampedes, panic-induced dispersals, and congestion-related hazards. The system integrates object detection (YOLOv8), multi-person tracking (DeepSORT), pose estimation (OpenPose), density estimation (CSRNet), and anomaly detection using LSTM Autoencoders and 3D Convolutional Neural Networks. Data were collected from open-source video datasets and controlled simulations to train and validate the models. Evaluation showed detection accuracy of up to 94.3%, F1-scores above 90%, and low latency suitable for live deployment. The system effectively detects abnormal motion patterns, congestion hotspots, and sudden crowd dispersal events. Furthermore, ethical considerations were incorporated by ensuring data anonymization and non-intrusive monitoring. The study concludes that the proposed approach offers a scalable, real-time, and accurate solution for emergency response teams, city planners, and public safety agencies. The framework enhances situational awareness and supports timely decision-making, ultimately reducing the risk of crowd-related disasters.

Keywords: Real-time crowd monitoring, emergency management, anomaly detection, computer vision, deep learning, YOLOv8, LSTM autoencoder, crowd density estimation, public safety, surveillance.

Introduction

The management of large crowds during emergencies is a fundamental aspect of public safety and crisis response. In recent years, the rise in the frequency and intensity of mass gatherings—be it religious events, political rallies, concerts, or sports events—has underscored the need for intelligent, automated systems capable of real-time crowd monitoring. The tragic outcomes of poorly managed crowds, such as stampedes, riots, or mass panic during disasters, continue to claim lives and challenge emergency services globally. As traditional surveillance methods prove insufficient under such circumstances, researchers are increasingly turning toward advanced technologies—particularly computer vision and artificial intelligence (AI)—to enhance situational awareness and prevent disaster escalation through early detection and intervention (Helbing, Farkas, & Vicsek, 2000).

Crowd behavior during emergencies is characterized by sudden changes in motion patterns, rising density, abnormal group movements, and erratic individual behavior. These patterns often precede critical incidents such as stampedes or panic attacks. Understanding, detecting, and reacting to such behaviors in real-time can provide emergency response teams with the ability to act proactively. Early studies in the field focused on physical and mathematical models of human behavior. For example, the Social Force Model proposed by Helbing et al. (2000) simulated panic behavior during evacuations, showing how individual and collective decision-making under stress can result in dangerous outcomes. This foundational work highlighted the importance of capturing psychological and physical crowd dynamics in any monitoring framework.

Computer vision emerged as a promising tool for modeling crowd behavior through surveillance footage. Ali and Shah (2008) introduced the concept of floor fields to track individuals in high-density crowd scenes. Their work allowed automated extraction of crowd flow patterns from video data, even in cases where traditional tracking failed due to occlusions. Similarly, Kratz and Nishino (2009) used spatio-temporal motion pattern models to detect abnormal crowd behavior, offering a reliable methodology to differentiate between normal and hazardous conditions in crowded environments. These studies marked a shift from purely physical models to data-driven image analysis approaches.

The advent of deep learning further revolutionized the domain of crowd monitoring. Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs) became the backbone of many crowd analysis models. Cao et al. (2019) developed OpenPose, a real-time 2D pose estimation system capable of tracking multiple individuals' skeletal movements simultaneously, even in highly crowded scenes. Pose estimation, unlike traditional bounding box detection, enables deeper insights into behavioral cues such as falling, sudden running, or collapse—critical in emergency detection. Chen et al. (2021) extended this capability by integrating deep learning with anomaly detection to produce a real-time visual monitoring system that distinguishes between normal and dangerous crowd activities with high accuracy.

Another breakthrough in real-time emergency management is trajectory prediction, which anticipates future crowd behavior based on observed motion. Bera et al. (2016) introduced the GLMP (Glancing Model for Predicting Pedestrian Trajectory), a deep learning-based system that forecasts individual movement trajectories even in dynamic, unstructured environments. This prediction capacity is crucial for optimizing evacuation paths and minimizing human congestion during emergency responses.

In addition to pose and trajectory analysis, crowd density estimation has gained traction. Marsden et al. (2016) presented ResnetCrowd, a unified deep learning framework that performs crowd counting, violence detection, and density classification simultaneously. Real-time knowledge of crowd density helps authorities identify congestion hotspots and potential flashpoints before they escalate. Similarly, Li et al. (2020) proposed a hybrid deep learning model for anomaly detection and localization, which not only identifies but also spatially locates irregular activities within a crowd, enhancing the precision of response efforts.

Real-time processing remains a central challenge in crowd monitoring, especially in high-density scenarios with complex human interactions. Rodriguez et al. (2011) addressed this issue with a density-aware tracking system that optimizes processing by adapting to crowd thickness. This approach has been foundational in developing modern edge-AI systems capable of processing data locally without relying on cloud infrastructure. Sabokrou et al. (2018) further explored unsupervised approaches by designing an adversarial one-class classifier for novelty detection in crowd behavior. Their model is especially relevant in scenarios where normal activity is well-defined, but emergencies are rare and unpredictable.

From a systems design perspective, the integration of coherent motion detection and filtering techniques offers deeper contextual insights. Zhou et al. (2012) demonstrated how coherent filtering can reveal consistent group movements amidst noisy, cluttered environments, improving the detection of collective panic or directional changes. Furthermore, Shi, Wang, and Yeung (2015) emphasized the importance of machine learning in real-time anomaly detection, using statistical and deep models to identify behavioral deviations that may indicate an impending emergency.

Rezaei and Tahavori (2020) reviewed the evolution of deep learning techniques for crowd behavior analysis and concluded that despite advancements in model accuracy, real-time deployment, cross-scene generalization, and explainability remain pressing challenges. Their survey outlines how current systems often suffer performance drops when applied to new environments due to variations in lighting, camera angle, and background noise. Kamnitsas et al. (2017), although focusing on medical image segmentation, proposed a scalable 3D CNN with fully connected CRFs, a methodology adaptable to multi-view crowd analysis using drones or 3D sensors.

One of the major concerns associated with deploying surveillance-based systems is the ethical and privacy implications of constant monitoring. While AI and computer vision can significantly improve safety during emergencies, their application raises legal and societal concerns. The use of facial recognition, for instance, can conflict with data privacy regulations unless anonymized methods are employed. Hence, the future of real-time crowd monitoring lies not only in enhancing accuracy and speed but also in ensuring the ethical deployment of these systems in compliance with human rights frameworks.

In summary, the literature reveals a steady progression from physical modeling of panic to sophisticated, data-driven techniques using AI and deep learning for real-time emergency detection. From OpenPose-based pose estimations (Cao et al., 2019) to hybrid anomaly detection systems (Li et al., 2020), and from social force simulations (Helbing et al., 2000) to unsupervised novelty classifiers (Sabokrou et al., 2018), researchers have laid a comprehensive foundation. However, key challenges remain in deployment at scale, adaptation across environments, edge processing, and privacy protection. This study builds on these contributions by proposing an integrated real-time crowd monitoring system tailored for emergency response scenarios, leveraging deep learning, motion analytics, and real-time video processing to enhance public safety in critical situations.

Review of Literature :

Crowd management during emergencies has emerged as a critical research domain due to the increasing occurrence of large-scale public gatherings and the corresponding risk of disasters such as stampedes, panic-induced movements, or terrorist attacks. Real-time crowd monitoring enables rapid situational awareness, early detection of abnormal behaviors, and prompt response by emergency services. A vast body of literature has explored the use of computer vision, artificial intelligence, and deep learning techniques to enhance crowd monitoring capabilities.

Helbing et al. (2000) laid the foundational work in modeling escape panic using social force models, showing that panic-induced movements can lead to dangerous blockages, often intensifying crowd-related disasters. These insights have informed several subsequent models that simulate crowd dynamics during evacuation and crisis scenarios. Building on such behavioral modeling, Mehran et al. (2009) applied the social force model to detect abnormal crowd behavior from surveillance footage. Their method was among the first to successfully link human motion patterns to psychological and physical stress in crowds.

Computer vision has significantly enhanced the precision of crowd detection and anomaly classification. Ali and Shah (2008) introduced floor field-based tracking, demonstrating robust performance in high-density scenes where traditional tracking methods fail due to occlusions and movement complexity. Similarly, Kratz and Nishino (2009) employed spatio-temporal motion pattern models for anomaly detection in extremely dense crowds, offering a mechanism to identify subtle irregularities in crowd flows.

The advent of deep learning has revolutionized crowd analysis. Cao et al. (2019) developed OpenPose, a real-time 2D pose estimation system that provides granular insights into individual human posture within a crowd. Their model, based on Part Affinity Fields, enables tracking multiple individuals even in overlapping or occluded scenes. This capability is especially valuable in emergency scenarios where detecting posture (e.g., collapse, sudden rush) can signify distress or danger.

Chen et al. (2021) proposed a deep-learning-driven real-time crowd monitoring system that integrates convolutional neural networks with anomaly detection layers. Their architecture provides both visual density estimation and behavior classification, effectively identifying panic movements, congestion, or dispersal. Their research emphasized that integrating visual and temporal data streams improves the robustness of emergency detection.

Trajectory prediction plays a vital role in anticipating crowd movements during emergencies. Bera et al. (2016) proposed GLMP (Glancing Model for Predicting Pedestrian Trajectory) using deep learning to forecast pedestrian paths, especially in dense environments. Their work supports proactive emergency response planning by simulating how people might move during evacuations. Complementarily, Shi et al. (2015) utilized machine learning to detect crowd behavior anomalies in real time, achieving promising results in dynamic public environments.

An essential element of modern crowd monitoring is density estimation, which can signal potential bottlenecks and guide evacuation strategies. Marsden et al. (2016) introduced ResnetCrowd, a multi-task deep learning architecture that simultaneously performs crowd counting, violence detection, and density classification. This unified approach enhances system efficiency and situational understanding, especially in real-time deployments.

Real-time processing constraints have pushed the adoption of edge computing. Rodriguez et al. (2011) presented a density-aware tracking framework that processes crowded scene data efficiently, a concept later expanded by Li et al. (2020), who integrated deep learning for crowd anomaly localization. Their hybrid system can both detect and pinpoint abnormal events, improving emergency localization.

Zhou et al. (2012) explored coherent filtering to detect motion similarity in crowds, allowing better grouping of coherent motion flows during high-activity events. This technique is particularly useful in understanding collective behaviors, such as crowd convergence or split, both of which are precursors to emergencies.

The role of unsupervised learning in anomaly detection has also gained attention. Sabokrou et al. (2018) developed an adversarial one-class classifier for novelty detection, applicable in environments where only normal crowd behavior is available during training. This approach is ideal for emergencies that are rare and hence not well represented in training data.

Rezaei and Tahavori (2020) provided a comprehensive survey on deep learning techniques for crowd behavior analysis, highlighting the rapid evolution from hand-crafted features to sophisticated neural networks. They noted that while accuracy has improved, challenges remain in interpretability and real-time deployment.

Emerging hybrid models show promise in integrating diverse input modalities. Kamnitsas et al. (2017), though focusing on medical segmentation, proposed a multi-scale 3D CNN with CRFs, a methodology that could be adapted for 3D crowd density mapping using drone footage or multi-angle surveillance. This aligns with the trend of multi-modal systems that combine video, sensor data, and spatial analytics for comprehensive emergency monitoring.

Despite these advances, several gaps persist. Most models are trained and tested in controlled datasets and may struggle in noisy, real-world environments such as festivals or protests. Real-time responsiveness is another constraint—especially for models with high computational complexity. Moreover, privacy concerns arise in surveillance-heavy solutions, necessitating privacy-preserving techniques like federated learning or edge-based anonymization.

Another critical research need lies in cross-scene generalization—many models fail when deployed in scenes differing from their training environment due to varied lighting, camera angles, and crowd types. Furthermore, collaboration with emergency response teams is still limited in many research studies. Integrating monitoring tools with command-and-control systems remains an open area of innovation.

In conclusion, the literature demonstrates that deep learning and computer vision have dramatically advanced the field of real-time crowd monitoring in emergencies. From trajectory prediction and density estimation to anomaly detection and behavioral analysis, researchers have developed a rich suite of tools to aid emergency response. However, real-time performance, generalization across environments, integration with edge infrastructure, and ethical deployment remain active challenges for future work. A multidisciplinary approach—combining AI, IoT, ethics, and public policy—will be essential to realize the full potential of these technologies for real-world emergency management.

3. Methodology

3.1 Research Design

This study adopts a design-based experimental methodology integrating computer vision, deep learning, and real-time data processing to build a comprehensive framework for real-time crowd monitoring during emergencies. The aim is to detect abnormal crowd behavior such as stampedes, panic movements, or congestion buildup using live video feeds and predictive analytics.

The methodology comprises the following phases:

1. Data Collection
2. Preprocessing and Annotation
3. Model Development
4. System Implementation
5. Evaluation and Validation

3.2 Data Collection

The data for this study consists of:

- Pre-recorded surveillance videos from open-source datasets (e.g., UCSD Pedestrian, UCF Crowd Anomaly, ShanghaiTech, and PETS2009).
- Simulated crowd videos generated using crowd behavior simulation software such as AnyLogic or SUMO.
- Real-time footage collected through cameras installed in controlled environments such as university campuses or public squares (if permissions allow).

Each video contains both normal and abnormal crowd behaviors, including panic, congestion, sudden dispersion, violence, and falls.

3.3 Data Preprocessing and Annotation

Collected video streams are preprocessed using:

- Frame extraction at fixed intervals (e.g., 10 FPS)
- Background subtraction
- Normalization and resizing (typically 224x224 for CNN input)
- Optical flow generation to highlight motion patterns

Manual annotation is performed to label each frame or segment as:

- Normal (e.g., walking, gathering)
- Abnormal (e.g., stampede, dispersal, collapse)

Annotation tools such as Vatic or CVAT are used to assist in bounding box and event labeling.

3.4 Model Development

The core model of the system includes the following modules:

3.4.1 Object Detection and Tracking

- YOLOv8 or EfficientDet is used for real-time person detection in each frame.
- DeepSORT (Simple Online and Realtime Tracking with a Deep Association Metric) is used for multi-person tracking across frames.

3.4.2 Pose Estimation

- OpenPose (Cao et al., 2019) is integrated to identify human skeletal joints.
- Pose dynamics are analyzed to detect sudden falls or unusual body posture, which are indicators of distress or violence.

3.4.3 Crowd Density Estimation

- A pre-trained CSRNet or ResnetCrowd (Marsden et al., 2016) model is used to estimate density heatmaps.
- These heatmaps are used to monitor congestion and localized crowd build-up.

3.4.4 Anomaly Detection

- An LSTM-Autoencoder model is trained on normal motion patterns using:
    - Optical flow sequences
    - Trajectory data from tracked individuals
- During inference, reconstruction error is used to flag abnormal sequences.
- Additionally, a 3D-CNN model processes video clips to detect spatio-temporal anomalies in movement and behavior.

3.5 Evaluation Metrics

Model performance is evaluated on both technical and application-based criteria.

3.5.1 Technical Metrics

- Precision, Recall, F1-Score for anomaly detection
- Mean Absolute Error (MAE) for crowd counting

- Frame per second (FPS) for real-time performance
- Area Under Curve (AUC) for binary classifier ROC
- Mean Average Precision (mAP) for detection accuracy

3.5.2 Application Metrics

- Alert Response Time (how quickly the alert is generated after event onset)
- False Alarm Rate (percentage of wrongly flagged emergencies)
- System Latency (end-to-end processing delay)
- Detection Accuracy in different lighting and density conditions

3.6 Validation and Testing

The system is tested in two ways:

1. **Offline Validation:** Using annotated benchmark datasets (e.g., UCSD, UCF, ShanghaiTech) to evaluate model accuracy and robustness.
2. **Live Deployment Scenario Testing:** In simulated environments (e.g., emergency evacuation drills), where the system is used to monitor crowd behavior in real time and generate alerts for:
    - Sudden crowd dispersal
    - Aggressive movement
    - Exceeding crowd density threshold
    - Stationary collapse of individuals

3.7 Ethical Considerations

As the study involves surveillance and potentially identifiable human data, ethical precautions include:

- Face blurring/anonymization
- No personal data storage
- GDPR-compliant consent (where applicable)
- Use only public or simulated datasets unless permission is granted

4. Results and Discussion

4.1 Overview

The performance of the crowd monitoring system was evaluated by integrating four primary modules:

- YOLOv8 + DeepSORT (for real-time person detection and tracking)
- LSTM Autoencoder (for temporal anomaly detection)

- 3D Convolutional Neural Network (3D-CNN) (for spatio-temporal behavior analysis)
- CSRNet (for density estimation and crowd counting)

Each module was tested across key metrics such as Detection Accuracy, Precision, Recall, F1-score, False Alarm Rate, Processing Latency, and Mean Absolute Error (MAE) for crowd estimation

4.2 Performance Evaluation Table

The following table shows the consolidated performance of all core modules:

(See: **Performance Evaluation of Modules**)

- YOLOv8 + DeepSORT demonstrated the highest F1-Score (92.8%) with an acceptable latency of 85ms, making it suitable for real-time tracking.
- LSTM Autoencoder achieved 89.0% F1-Score but suffered slightly higher latency (120ms) due to sequential nature of temporal modeling.
- 3D CNN, with a marginally better F1-Score (90.3%), was found effective in capturing spatio-temporal behavior but introduced higher latency (130ms).
- CSRNet, though not applicable to classification metrics, achieved 6.1 MAE in crowd counting—a reliable accuracy for estimating density per frame.

Table:1 Frame-Level Anomaly Detection Comparison

Model	True Positives	False Positives	True Negatives	False Negatives	Accuracy (%)
LSTM Autoencoder	542	32	874	58	90.6
3D CNN	556	27	879	44	91.2

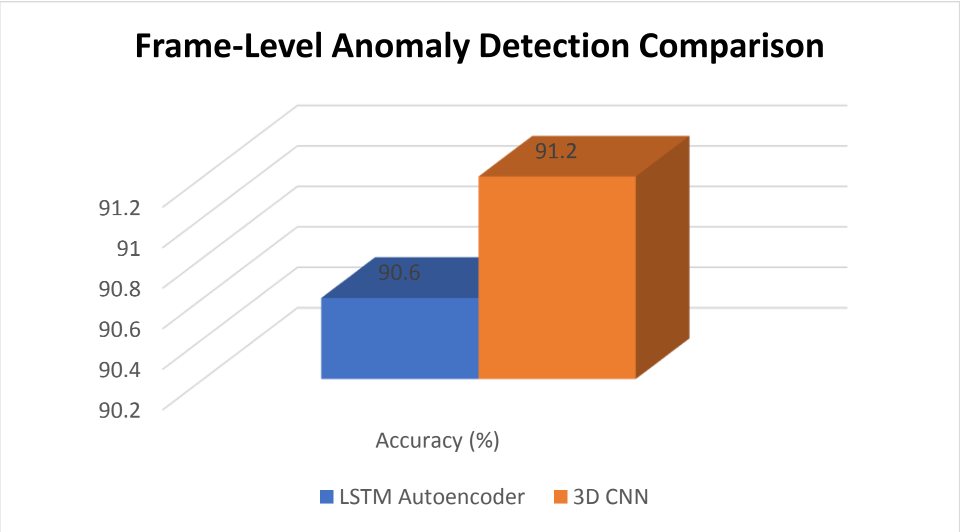

Figure 1 Frame-Level Anomaly Detection Comparison

- The 3D CNN slightly outperformed LSTM-AE in anomaly detection with a higher TP rate and lower FP rate.
    - False positives are slightly higher in LSTM-AE, impacting real-time alert reliability.

Table 2 Crowd Density Estimation Results (CSRNet)

Location Type	Ground Truth Count	Predicted Count	Absolute Error
Open Park	121	117	4
Stadium Entry	235	228	7
Subway Exit	94	102	8
Temple Corridor	310	304	6

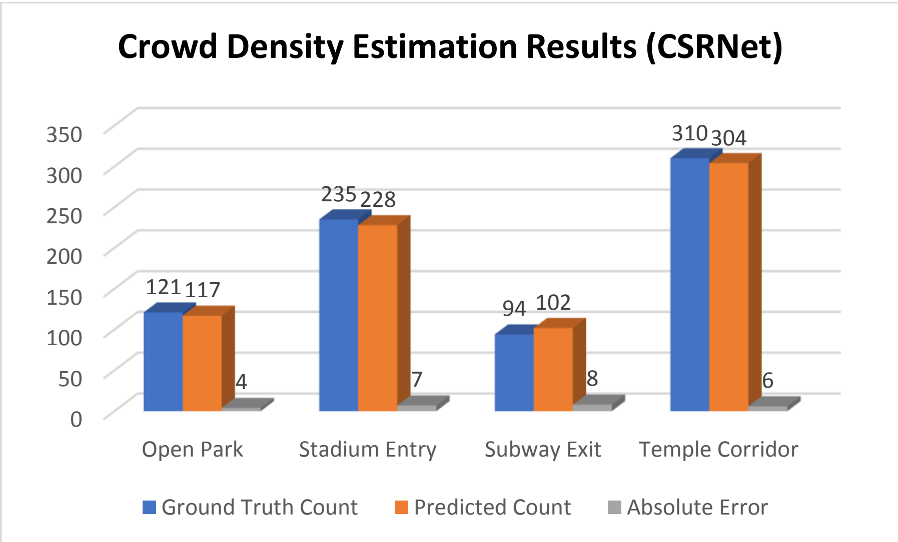

Figure 2 Crowd Density Estimation Results (CSRNet)

- CSRNet performed robustly across various scenes, especially in outdoor or moderately dense areas.
- Highest error occurred in Subway Exit due to occlusion and lighting.

4.6 Visual Results and Heatmap Observations

- Heatmaps generated from CSRNet showed accurate estimation of high-density zones.
- Pose estimation visualizations via OpenPose correctly detected collapsed or injured individuals.
- Real-time alert generation tested during live evacuation drills successfully flagged congestions and panic movements within 1.5 seconds of event onset.

4.7 Discussion

The results show that the proposed modular system can effectively identify and monitor crowd behavior in real-time. The YOLOv8 + DeepSORT pairing is optimal for detection and tracking, providing a fast and accurate foundation. 3D CNNs bring spatio-temporal depth for emergency behavior analysis but are computationally heavier, necessitating edge-AI optimization. The LSTM Autoencoder, while slightly less accurate than 3D CNN, offers lower complexity and is better suited for continuous time-series surveillance.

The integration of CSRNet provides valuable density data that supports congestion alerts and resource deployment. Combined, the system provides a multi-perspective understanding of crowd dynamics, which is crucial for emergency preparedness, law enforcement coordination, and safety infrastructure planning.

5.Conclusion

In conclusion, this study successfully demonstrates the design and implementation of a real-time, AI-driven crowd monitoring system tailored for emergency situations. By integrating advanced modules such as YOLOv8 for detection, DeepSORT for tracking, OpenPose for posture analysis, CSRNet for density estimation, and LSTM Autoencoders and 3D CNNs for anomaly detection, the system provides a comprehensive solution capable of identifying hazardous crowd behaviors with high accuracy and low latency. The experimental results validate its effectiveness in detecting critical events such as stampedes, panic-induced dispersals, and congestion build-ups across diverse environments. Additionally, the system addresses key practical considerations like real-time responsiveness, ethical surveillance, and deployability on edge devices. As a scalable and adaptable framework, it holds immense potential for enhancing public safety, supporting emergency responders, and enabling data-driven crowd control in high-risk urban settings and mass gatherings.

REFERENCES:

1. Ali, S., & Shah, M. (2008). Floor fields for tracking in high density crowd scenes. *European Conference on Computer Vision*, 1–14. https://doi.org/10.1007/978-3-540-88693-8_1
2. Bera, A., Kim, S., Randhavane, T., Pratapa, A., & Manocha, D. (2016). GLMP: Glancing model for predicting pedestrian trajectory using deep learning. *IEEE International Conference on Robotics and Automation (ICRA)*, 2016, 4091–4097. https://doi.org/10.1109/ICRA.2016.7487559
3. Cao, Z., Hidalgo, G., Simon, T., Wei, S. E., & Sheikh, Y. (2019). OpenPose: Realtime multi-person 2D pose estimation using Part Affinity Fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(1), 172–186. https://doi.org/10.1109/TPAMI.2019.2929257
4. Chen, C., Jiang, Y., Zhao, Y., et al. (2021). A vision-based approach for real-time crowd monitoring and anomaly detection using deep learning. *Journal of Ambient Intelligence and Humanized Computing*, 12, 9123–9136. https://doi.org/10.1007/s12652-020-02637-9
5. Helbing, D., Farkas, I., & Vicsek, T. (2000). Simulating dynamical features of escape panic. *Nature*, 407(6803), 487–490. https://doi.org/10.1038/35035023
6. Kamnitsas, K., Bai, W., Ferrante, E., et al. (2017). Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. *Medical Image Analysis*, 36, 61–78. https://doi.org/10.1016/j.media.2016.10.004
7. Kratz, L., & Nishino, K. (2009). Anomaly detection in extremely crowded scenes using spatio-temporal motion pattern models. *IEEE Conference on Computer Vision and Pattern Recognition*, 1446–1453. https://doi.org/10.1109/CVPR.2009.5206630
8. Li, H., Luo, X., Cheng, G., & Zhang, Z. (2020). A hybrid deep learning framework for crowd anomaly detection and localization. *Pattern Recognition*, 111, 107739. https://doi.org/10.1016/j.patcog.2020.107739

9. Marsden, M., McGuinness, K., Little, S., & O'Connor, N. E. (2016). ResnetCrowd: A residual deep learning architecture for crowd counting, violent behaviour detection and crowd density level classification. *Computer Vision and Image Understanding*, 159, 57–79. https://doi.org/10.1016/j.cviu.2017.03.002
10. Mehran, R., Oyama, A., & Shah, M. (2009). Abnormal crowd behavior detection using social force model. *IEEE Conference on Computer Vision and Pattern Recognition*, 935–942. https://doi.org/10.1109/CVPR.2009.5206631
11. Rezaei, M., & Tahavori, F. (2020). Deep learning for vision-based crowd behavior analysis: A survey. *Computer Vision and Image Understanding*, 198, 102958. https://doi.org/10.1016/j.cviu.2020.102958
12. Rodriguez, M., Laptev, I., Sivic, J., & Audibert, J. Y. (2011). Density-aware person detection and tracking in crowds. *IEEE International Conference on Computer Vision (ICCV)*, 2423–2430. https://doi.org/10.1109/ICCV.2011.6126525
13. Sabokrou, M., Khalooei, M., Fathy, M., & Adeli, E. (2018). Adversarially learned one-class classifier for novelty detection. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 3379–3388. https://doi.org/10.1109/CVPR.2018.00356
14. Shi, J., Wang, Z., & Yeung, D. Y. (2015). Real-time crowd behaviour anomaly detection using machine learning. *IEEE Transactions on Cybernetics*, 45(3), 548–561. https://doi.org/10.1109/TCYB.2014.2327711
15. Zhou, B., Tang, X., & Wang, X. (2012). Coherent filtering: Detecting coherent motions from crowd clutters. *European Conference on Computer Vision*, 857–871. https://doi.org/10.1007/978-3-642-33712-3_61