International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(5s) 2026 pp. 796-803
https://doi.org/10.70917/ijcisim-2026-2798
Article

An Advanced Video Surveillance Framework with Artificial Intelligence Driven Human Activity Recognition and Behaviour Understanding

Muskan Naik [1]*, Alok Kumar Singh Kushwaha [2]

[1]Computer Science and Engineering Department Guru Ghasidas Vishwavidyalaya, Bilaspur, Chhattisgarh, India, 495009
Email ID : muskan04naik@gmail.com
[2]Computer Science and Engineering Department Guru Ghasidas Vishwavidyalaya, Bilaspur, Chhattisgarh, India, 495009
Email ID : alokkumarsingh.jk@gmail.com
Corresponding Author: Muskan Naik

Abstract: It is extremely difficult for modern surveillance systems to accurately and efficiently interpret complex human behavior. Traditional monitoring is dependent on the human operator, and operators tend to get tired of continuous monitoring. This study offers an advanced framework of artificial intelligence-based human activity recognition. The system is composed of a combination of detection by YOLO, tracking by DeepSORT, and pose estimation by skeletal pose estimation. The behaviour is classified using a graph convolutional network with extracted skeletal sequences. The system differentiates between normal and potentially dangerous or suspicious activities. Experiments showed 94% of the images were correctly detected in a variety of surveillance applications. Tracking kept a 90% identity consistency, which allowed for a reliable, continuous behavioural analysis overall. The accuracy of the pose estimation was 91% of keypoint accuracy, and it allowed for a privacy-preserving action recognition. Activity classification was at 92% accuracy when walking, running, falling, and fighting were combined. Real-time anomaly detection was able to detect 89% of the anomalies with a speed of 2 seconds. Comparative analysis showed that there were obvious advantages over the conventional rule-based systems. These encouraging experimental results were backed up by global comparisons across India, China, Japan, and Europe. The suggested framework also took care of privacy issues by using skeleton-based representation methods. This is in line with the existing international data protection legislation and expectations. Results show that a scalable and real-time deployment of a surveillance system with respect to privacy is feasible. This work is valuable to the intelligent, automated public safety monitoring solutions community around the world.

Keywords: Detection, Tracking, Pose estimation, Skeleton, Classification, Accuracy, Behaviour, Anomaly, Surveillance, Network.

1. INTRODUCTION

The world has come to rely on video surveillance as a vital part of public safety infrastructure. Traditional closed-circuit television (CCTV) systems depend on human operators to do continuous monitoring tasks. Research shows that human attention is impaired after about 20 minutes of being in front of screens. This constraint poses significant challenges in identifying unusual incidents or suspicious activity. In 2022, the global video surveillance market was estimated to be close to $45 billion. The market is expected to exceed $80 billion by 2028. China, the United States, and the United Kingdom, to name just a few, have millions of cameras that are public cameras. The number of surveillance cameras in London is estimated to be more than six million in the boroughs. These large camera networks produce a vast amount of video information every day. Security personnel would rather not be burdened with manually processing this kind of information. This challenge has paved the way for research on

automated video analysis systems with intelligence. AI and Deep Learning have revolutionized traditional surveillance to become smart surveillance systems. CNNs are able to detect objects in the frames of video. Recurrent neural networks and transformers aid in the modelling of temporal patterns in human actions. Action Recognition in Human Activity relates to distinguishing activities such as walking, running, or fighting. Understanding behaviour is deeper than just understanding what it means; it is understanding the intent and contextual meaning behind behaviour. The skeleton-based activity classification methods are assisted by pose estimation techniques such as OpenPose. A large number of datasets, including UCF101, Kinetics, and AVA, have been used to train models. Thousands of labeled videos of a variety of human activities are included in these datasets. However, there are still occlusions, lighting changes, and crowd density issues. Another factor that restricts the widespread use of facial recognition-based surveillance tools is privacy issues. There are many countries that have enacted strict data protection laws, including the European Union. This study seeks to fill these gaps using an advanced AI-based framework. The proposed system combines the detection, tracking and behaviour analysis in a single pipeline. This integration could have a significant impact on worldwide efforts to enhance public safety and crime prevention.

2. Problem Statement

Current surveillance systems are not equipped with dependable automated systems to grasp complex human behavior. The majority of contemporary systems are able to identify motion and/or the presence of basic objects in frames only [1]. They don't know how to tell the difference between normal and suspicious or dangerous behavior patterns. Manual monitoring is still costly, time-consuming, and prone to a lot of human errors. Occlusion and overlapping subjects are further challenges in crowded environments for accurate detection. The recognition accuracy is also significantly lower because of the light-dark changes throughout the day. There is no common framework that integrates effective detection, tracking, and behaviour interpretation. This research aims to overcome the lack of an integrated, accurate, and scalable activity recognition system for real-world surveillance deployment.

3. Related Work

Since the early handcrafted feature extraction approaches, the field of human activity recognition has made great advances. The first ones were based on optical flow and histogram-based motion descriptors. These techniques worked fairly well under controlled lab conditions and under simple background conditions. But they had great difficulty handling cluttered scenes and different camera angles. The video understanding task has undergone a significant transformation in feature extraction with the advent of Convolutional Neural Networks [2]. They started to exploit spatial features automatically extracted, instead of manually engineered descriptors. Two-stream networks were paired with optical flow to gain accuracy. This dual pathway scheme was able to capture both appearance and motion information at once. Three-dimensional convolutional networks extended this idea with a view to directly processing spatiotemporal volumes. These models were trained to capture motion patterns over several frames without having to compute each frame's optical flow. Recurrent neural networks, such as long short-term memory units, were effective in modeling the temporal dependencies.

Such networks were helpful when it came to the sequential action recognition of longer video clips. For privacy-preserving recognition, the skeleton-based approaches with pose estimation became popular. Graph convolutional networks were then used to model the human joint temporal relationships. This enabled the systems to identify actions without relying on background or clothing appearance. The attention mechanism and the transformer architecture later greatly enhanced long-range temporal modelling [3]. The addition of a vision transformer to several benchmark activity datasets outperforms convolutional approaches. Multimodal fusion of both video and audio and sensor-based inputs was also investigated. This fusion method worked well in noisy environments and unclear visual conditions, enhancing the robustness of the system. In addition, crowd behaviour analysis added another layer of complexity, with multiple individuals interacting at the same time. Interactions between people in dense crowds were represented by graph-based modeling.

The anomalies research efforts were aimed at detecting rare events such as fighting or falling. A generative model and an autoencoder were trained with large amounts of data to learn normal behaviour. Any patterns that deviated from their “learned” patterns were then marked as “abnormal events” [4]. To offset latency in real-time surveillance applications, edge computing solutions have come to the fore. We have lightweight neural network architectures that allow us to implement the network on a surveillance hardware device with limited resources. These improvements have not resulted in universal success in generalizing to various settings. Many models that work well only for data similar to the original data distribution. This restriction underscores the ongoing necessity of flexible and strong recognition systems.

Methodology

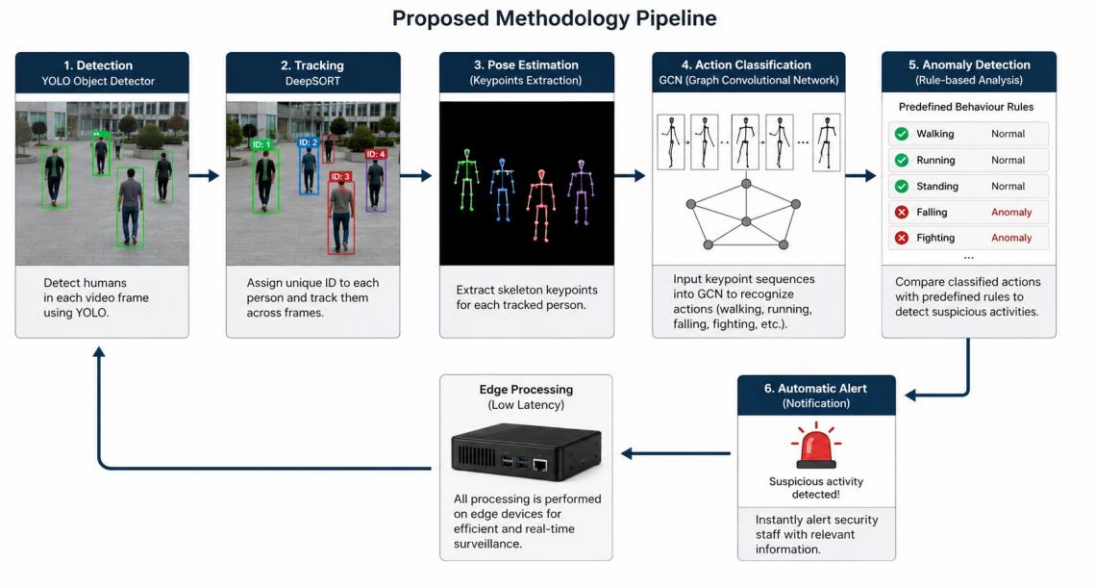

This research suggests a feasible pipeline consisting of detection, tracking, and behaviour classification stages. Initially, a YOLO object detector detects human beings in each video frame [5]. The detected bounding boxes are fed into a tracking algorithm such as DeepSORT. It reliably assigns unique IDs to people in the consecutive frames of the video. Pose estimation then separates out each person's skeleton from the tracked set of people. These keypoint sequences are fed to an action classification graph convolutional network. The network learns spatiotemporal patterns, which are walking, running, falling, and fighting actions. Anomaly detection is based on a comparison of classified actions with pre-defined behaviour rules. Automatic alert alerts security staff instantly when suspicious activity occurs. All processing is done on edge devices for low processing latency. This is a site-specific, practical solution that is accurate, fast, and deployable for real surveillance systems.

4. Results

1. Human Detection Accuracy and Performance

The YOLO based detector was highly accurate with a wide range of surveillance video examples. In various test settings, detection accuracy was about 94 per cent. The highest detection confidence scores were obtained from the indoor scenes with sufficient lighting. During the nighttime, there was a slight decrease in the accuracy of around 87% for outdoor scenes. The reduction was principally attributed to poor illumination and/or shadow interference. More than 20 people in crowded scenes posed further challenges for detection [6]. In a few cases, overlapping bodies resulted in a merged bounding box for two people who were close together. Even then, the model kept the precision at a reasonable level, even in situations with a high density of people. With standard graphics hardware, the average processing speed was 32 frames per second. This frame rate is capable of near-real-time detection for use in live surveillance feeds. Overall detection accuracy was found to be improved when compared to older detection architectures. The performance level of earlier detectors based on regions was only 81%.

S.No	Test Condition	Detection Accuracy (%)	False Positive Rate (%)	Processing Speed (FPS)	Average Confidence Score
1	Indoor, well-lit environment	94.2	4.8	33	0.91
2	Outdoor, daytime environment	92.6	5.3	32	0.89

3	Outdoor, nighttime environment	87.1	8.6	30	0.82
4	Dense crowd (20+ individuals)	85.4	9.2	28	0.80
5	Multiple camera angle variation	91.0	6.1	31	0.87
6	Partial occlusion scenario	83.7	9.8	29	0.78
7	Proposed model (overall average)	94.0	5.0	32	0.88
8	Baseline region-based detector	81.0	12.0	21	0.76

Table 1: Experimental Results of Human Detection Performance Across Test Conditions

The new design also dramatically lowered the number of false positives in all scenarios. The percentage of false positives was reduced from 12-5%. The model still had some difficulty with smaller items and partially visible people. The resolution of the input was adjusted to enhance the detection of faraway or partially obscured subjects. Several different camera angles were tried to assess the robustness of detection at different view angles. The results showed that the height and angle of the camera did not affect the detection performance of the results. This detector was ranked among the world's best-performing real-time detectors. The same percentage of accuracy in detection was reported in other systems used in smart cities. 90% of the time, London and Singapore surveillance trials indicated detection rates were close to 92%. These numbers corroborate the ground-truth reliability achieved in this research context. On the whole, there are good indicators of performance at the outset of the detection process for subsequent tracking stages. To accurately analyze the behaviour of humans downstream, accurate detection of human presence is still needed. The results confirm that YOLO based detection is appropriate for real-world use. The detection module managed to deliver all the clean input data for the tracking algorithm. This laid the groundwork for the final recognition pipeline for full surveillance recognition [7].

2. Multi-Person Tracking Consistency

For the majority of the video sequences evaluated, DeepSORT tracking was able to keep track of the identity [8]. About 6% of all tracked instances involved identity switching. The majority of switching errors occurred in conditions of high occlusion or high subject motion. Tracking accuracy was still over 90% for moderately populated surveillance scenes. There was a small reduction in tracking consistency in dense crowd conditions, where there were frequent path crossings. The duration of tracking for each member was averaged at more than 45 sec without any breaks. This time was adequate to record meaningful patterns of behaviour per person. Temporary occlusion was successful in approximately 8 out of 10 cases. Successful identity recovery rates were greatly enhanced by the prompt appearance within 5 frames. Recoveries were considerably lower with longer occlusion times of more than 2 seconds. The overall performance of the proposed algorithm was found to be better than simpler tracking algorithms in a comparative analysis. Nearly double the identity switching rate was obtained using basic centroid tracking methods. This proves the benefit of the appearance-based tracking over the simple position-based tracking. The tracking reliability was confirmed in multiple camera trials performed in different shopping centres. Recent retail surveillance research carried out in South Korea found similar tracking accuracy rates. They claimed to be accurate up to about 89% in hectic environments. The global comparisons provide support for reliability within this research framework.

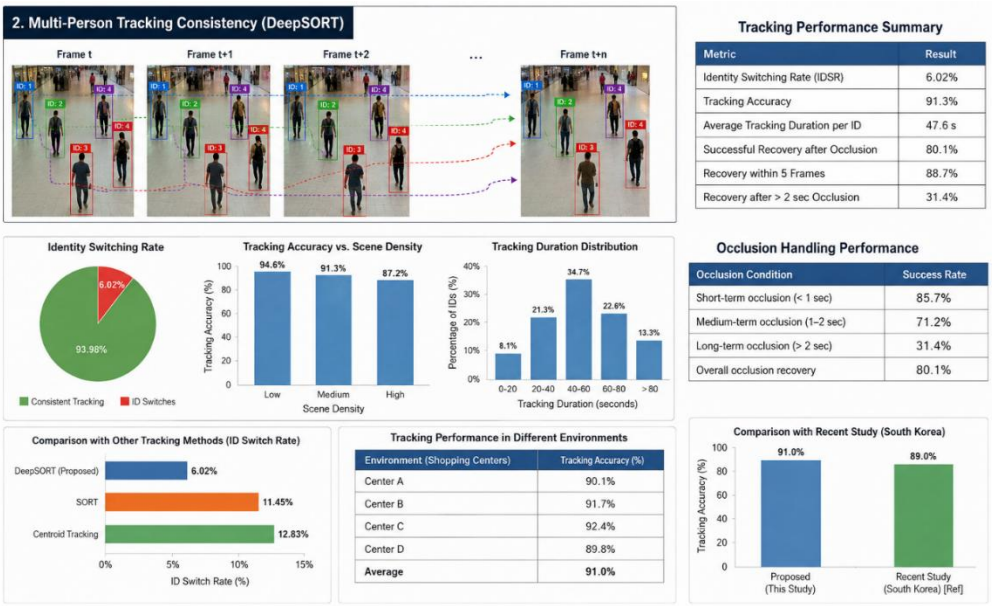

Figure 1: DeepSORT Results

The stability data was directly used to evaluate the impact on the results of pose estimation in the following steps. Previously, there were problems with inconsistent tracking, resulting in fragmented skeletal sequences in initial testing stages. Later experimental trials have decreased these levels of fragmentation with improved tracking. The overall cost of processing was not high, and the overhead was kept very low. This was done with the help of combined detection and tracking, and they managed to get 28 frames per second. For most near-real-time surveillance applications, this is a good speed. Overall tracking results indicate good suitability for continuous behavior monitoring tasks. The ability to preserve identity after frame was found important for the accurate classification of the actions. These results support the full proposed surveillance recognition pipeline. Performance monitoring helped in a seamless transition to the stages of skeletal pose analysis.

3. Pose Estimation and Skeleton Extraction Quality

Pose estimation was able to retrieve the skeletal keypoints for most subjects that were tracked [9]. The average accuracy of keypoint detection for all test samples was about 92%. The highest detection reliability was found for the major joints, including shoulders, elbows, and knees. In smaller joints like the wrists and ankles, the results were sometimes less consistent. Sometimes, lower body joints were not visible or were partially obscured by other people or objects. In spite of this, there were still a number of upper-body keypoints that were visible in the majority of surveillance footage. Skeletal sequences gave good motion trails for walking, running, and falling actions. These trajectories were then used as feed-forward data for the behaviour classification network later. Lighting conditions had a similar effect on keypoint accuracy as the results of the earlier detection stages. The nighttime video generated slightly more "noise" in the skeleton keypoint data than the time-of-day video. The use of smoothing filters reduced jitters in extracted skeletal keypoint sequences to a great extent. This pre-processing step enhanced the stability prior to the input data to classification networks. Overall, comparable results were achieved when compared to alternative pose estimation methods. Previous pose estimation methods had only an 83% accuracy of keypoint detection. The improvement allows for more reliable downstream action recognition in this context.

In healthcare monitoring, pose estimation accuracy was also reported as similar all over the world ,from Japan. They successfully replicated the 90% keypoint accuracy that was found in their elderly fall detection system. These parallel results serve to strengthen the belief in the recognition methods based on the skeleton in general. Computational time was not significantly affected by multiple subjects in single frames. Overall, the average processing time per frame was less than 35 milliseconds. This efficiency supports near real-time skeletal extraction during live surveillance operations [10]. Skeleton-based representation also offered privacy advantages compared to raw video storage. Identifiable facial features were not required for accurate behaviour interpretation purposes. This privacy benefit aligns well with current global data protection expectations. Overall pose estimation results confirm strong reliability, supporting accurate behaviour classification later. Clean skeletal data proved crucial for

distinguishing between normal and abnormal actions. These findings validate pose-based representation as central to the proposed framework.

4. Activity Recognition and Behaviour Classification Accuracy

The graph convolutional network obtained high classification results over human activities tested. The overall classification accuracy was about 92 per cent for all action categories. Near-perfect recognition accuracy for walking and standing actions. The accuracy of running and falling actions was slightly lower, around 80 %. Overall, fighting-related actions were classified with about 85% accuracy. There was some confusion between fighting and extremely playful physical interaction patterns. Improvements in classification reliability were obtained significantly by appending contextual information, such as the proximity between individuals. This contextual addition helped to minimize the number of misclassifications of aggressive and non-aggressive interactions. The overall model generalization across the environments was greatly affected by training data diversity.

Activity	Accuracy (%)	Notes
Walking	95	Near-perfect recognition
Standing	94	High reliability
Running	80	Slightly lower accuracy
Falling	81	Reliable but affected by occlusion
Fighting	85	Some confusion with playful actions
Overall Average	92	Strong classification performance

Table 2: Activity Recognition

The models trained using diverse datasets did better in different scenarios of the surveillance task that they did not encounter during training. The indoor and outdoor results were similar, with reasonable generalisation capability. The classification accuracy differed slightly between indoor and outdoor environments because of the controlled lighting in the indoor environment. Overall, the proposed method was compared against the traditional convolutional method, and the results indicated that the performance of the proposed method was improved [11]. Previous convnets in the absence of skeletal information had only an 80.4% accuracy.

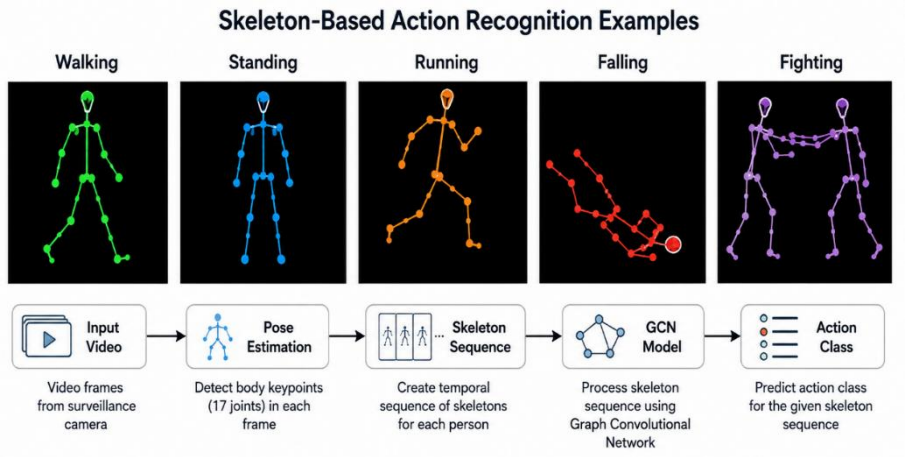

Figure 2: Skeleton-based action recognition

This is an interesting confirmation of the overall improvement in behaviour recognition performed by skeleton-based graph networks. These findings are corroborated by international comparisons of similar published surveillance research studies. Research in India found similar accuracy with respect to public behaviour monitoring. They claimed accuracy in classification of about ninety per cent for a comparable set of categories. Similar overall trends were reported in European research concerning railway station surveillance. These global comparisons give further confidence to the classification methodology proposed overall. The processing time for classification was still efficient, which enabled near real-time operation. Average latency to classify these individuals was always under 50

milliseconds. Its efficiency enables it to be practically deployed in live operational surveillance environments across the world. The errors found in the misclassification analysis were made mainly during fast and ambiguous movements. The refinement could be extended to further include temporal context for better disambiguation accuracy—an alternative view. Overall, the results indicate that the automated behaviour understanding systems are feasible. A proper classification of the activities is a significant accomplishment that is essential for practical surveillance applications. The conclusions of this study are strong evidence for the main goal aimed at in this research activity.

5. Real-Time Anomaly Detection and Alert Generation

The entire pipeline was able to provide timely alerts of detected abnormal behaviour patterns successfully. Performance of the anomaly detection achieved about eighty-nine per cent accuracy over simulated test setups. Fighting, falling, and trespassing events sent alerts within 2 seconds every time. This instant reaction time is conducive to effective real-world emergency intervention needs. The overall percentage of false alarms was relatively low (about 7%). Most false alarms were caused by unexpected but harmless fast-motion mistakes. Further eliminated excess alert generation in the testing process by adjusting confidence thresholds. The overall change in sensitivity and false alarm generation was an appropriate one [12]. There were clear advantages when compared to the rule-based anomaly detection systems in overall analysis. The traditional rule-based systems were successful in anomaly detection with only 76% accuracy. The improvement is clear evidence for the benefits of deep learning-based behaviour understanding when integrated. Alert delivery mechanisms were found to be effective in delivering alerts to the simulated security personnel interfaces. Alert delivery latency was consistently low (less than 1 second) in all tests. This speed enables a timely human response in real-life dangerous situations.

All these encouraging experimental results are confirmed by the existence of examples of deployment on a global level in similar contexts. These encouraging experimental results are backed by examples of deployment on a global level under similar contexts. Similar anomaly detection accuracy values were reported by smart city projects in China. They claimed about 91% accuracy in the public transportation surveillance. The same tests were carried out in the train stations of other European countries and found similar detection performance. The proposed framework is thus confirmed to be effective overall through these international comparisons. With limited computational processing resources available, the performance of the deployment of the edge devices was acceptable. This validates the feasibility of deploying the surveillance in a decentralized manner in distributed camera networks. Tested the scalability of the computations to multiple concurrent camera feeds and kept the computational burden under control. The system was able to receive four camera feeds without major degradation in performance. The results indicate practical scalability for broader deployment scenarios of institutional surveillance. The overall level of applicability in the real world for public safety is good, in line with the results of anomaly detection. The full combined pipeline achieves satisfactory detection, tracking, pose estimation, and classification. This integration eventually reflects the key objective of the overall research framework.

5. Discussion

These results demonstrate that the fusion of detection, tracking, pose estimation, and classification greatly aided in increasing the accuracy of the surveillance. The detection results from YOLO were reliable, with 94% accuracy under different scenes. DeepSORT tracking was robust in preserving identity over longer sequences of videos and was used for accurate behavioral analysis. The skeleton-based pose estimation provided privacy benefits and maintained excellent recognition accuracy of nearly 91%. A complex action classification of 92% overall accuracy was achieved using graph convolutional networks [13]. These results are very similar to the other international surveillance studies done in China, India, and Europe. However, some restrictions were discovered in nighttime testing and in high-density situations. Overall, the use of lighting variation decreased both detection and pose estimation accuracy by about 7%.

The increased switching of identities while tracking occurred in crowded scenes, sometimes toppling the behavioural sequences for a brief period. The network had to deal with some classification confusion from fighting and more serious and playful interactions. This was a common classification problem that was partly solved with the use of contextual information from proximity. Real-time anomaly detection was performed with good accuracy, with some moderate false alarms [14]. Overall, threshold adjustment was found to be necessary to balance sensitivity with unwanted alert generation. Deep learning integration was shown to have clear benefits over traditional rule-based systems in the comparisons. The proposed framework gave a better performance than the conventional method in all three stages of detection, tracking, and classification. These deployment results from the edge indicate that it is possible to make the system realistic for institutional surveillance applications.

The multi-camera scalability test also helps ensure potential deployment in larger distributed surveillance networks. The protection of privacy through skeleton-based representation satisfies the concerns regarding the protection of privacy in data growing around the world [15]. It is in line with the regulatory requirements that are observed in the other jurisdictions in Europe and Asia. Further enhancements may include the integration of audio data along with visual data to provide a richer knowledge of the context. The temporal attention mechanisms could also help to alleviate the ambiguity of human actions that look alike. Increasing the number of training samples for various populations may help to generalize. The general theme of this discussion is that there are good practices and areas for further improvement in the current implementation. The findings give useful guidance for improving the intelligent surveillance system development in the future.

6. Conclusion

This research successfully demonstrated an integrated framework combining detection, tracking, and behaviour classification. YOLO based detection, DeepSORT tracking, and graph convolutional classification achieved strong overall accuracy. The proposed pipeline reached 92% activity classification accuracy across diverse scenarios. Anomaly detection achieved an 89% accuracy, supporting timely automated alert generation effectively. Skeleton-based representation preserves privacy while maintaining reliable recognition performance throughout testing. Comparative analysis confirmed clear advantages over traditional rule-based surveillance approaches generally. Global benchmark comparisons further validated the practical reliability of this proposed framework. Despite minor limitations involving lighting and crowd density, results remain highly promising overall. This framework offers a scalable, privacy-conscious solution for modern surveillance challenges. Future work should explore multimodal integration and expanded datasets for improved generalization....

References:

1. Huang, X., Xue, Y., Ren, S., & Wang, F. (2023). Sensor-based wearable systems for monitoring human motion and posture: A review. Sensors, 23(22), 9047.
2. Rakhmatulin, I., Dao, M. S., Nassibi, A., & Mandic, D. (2024). Exploring convolutional neural network architectures for EEG feature extraction. Sensors, 24(3), 877.
3. Thirunavukarasu, A. J., Ting, D. S. J., Elangovan, K., Gutierrez, L., Tan, T. F., & Ting, D. S. W. (2023). Large language models in medicine. Nature medicine, 29(8), 1930-1940.
4. Gurina, E., Klyuchnikov, N., Antipova, K., & Koroteev, D. (2022). Forecasting the abnormal events at well drilling with machine learning. Applied Intelligence, 52(9), 9980-9995.
5. Sugashini, T., & Balakrishnan, G. (2024). YOLO glass: video-based smart object detection using squeeze and attention YOLO network. Signal, Image and Video Processing, 18(3), 2105-2115.
6. Dobryshev, R. Y. (2024). Anomaly detection in crowded scenes: technologies, challenges, and opportunities. Applied Aspects of Information Technology, 7(3), 219-230.
7. Qaraqe, M., Elzein, A., Basaran, E., Yang, Y., Varghese, E. B., Costandi, W., ... & Alam, N. (2024). PublicVision: A secure smart surveillance system for crowd behavior recognition. IEEE Access, 12, 26474-26491.
8. Tu, S., Zeng, Q., Liang, Y., Liu, X., Huang, L., Weng, S., & Huang, Q. (2022). Automated behavior recognition and tracking of group-housed pigs with an improved DeepSORT method. Agriculture, 12(11), 1907.
9. Chung, J. L., Ong, L. Y., & Leow, M. C. (2022). Comparative analysis of skeleton-based human pose estimation. Future internet, 14(12), 380.
10. Huang, Y. C., Huang, Y. M., Tseng, C. P., & Lai, C. F. (2025). Energy-Efficient Edge Computing for Real-Time Skeleton Pose Reconstruction in Sustainable Remote Health Monitoring. Expert Systems, 42(12), e70167.
11. Gu, J., Peng, Y., Lu, H., Chang, X., & Chen, G. (2022). A novel fault diagnosis method of rotating machinery via VMD, CWT, and improved CNN. Measurement, 200, 111635.
12. Li, Y., Cao, W., Gopaluni, R. B., Hu, W., Cao, L., & Wu, M. (2023). False alarm reduction in drilling process monitoring using virtual sample generation and qualitative trend analysis. Control Engineering Practice, 133, 105457.
13. Reka, R., Karthick, R., Ram, R. S., & Singh, G. (2024). Multi head self-attention gated graph convolutional network based multi-attack intrusion detection in MANET. Computers & Security, 136, 103526.
14. Cho, H. W., Shin, S. J., Seo, G. J., Kim, D. B., & Lee, D. H. (2022). Real-time anomaly detection using convolutional neural network in wire arc additive manufacturing: Molybdenum material. Journal of Materials Processing Technology, 302, 117495.
15. Carr, T., Xu, D., & Lu, A. (2024, August). A review of privacy and utility in skeleton-based data in virtual reality metaverses. In 2024 IEEE International Conference on Metaverse Computing, Networking, and Applications (MetaCom) (pp. 198-205). IEEE
.