iJRASET
International Journal For Research in Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** IV **Month of publication:** April 2026

DOI: https://doi.org/10.22214/ijraset.2026.81606

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue IV Apr 2026- Available at www.ijraset.com

Human Pose Estimation and Action Recognition in Video Using KeyPoint Based CNNs

M.Sailaja[1], E.Madhavi[2], P. Lakshmi Rajyam[3], Y. Anusha[4]
[1, 2, 3]*BTech,* [4]*Assistant Professor, Department of Computer Science & Engineering (AIML), Bapatla Women's Engineering College, Bapatla ,Andhra Pradesh*

Abstract: ***The Human Activity Recognition (HAR) from video remains a challenging problem in computer vision due to the complexity of human motion and the need to model both spatial appearance and temporal dynamics simultaneously. This paper proposes a hybrid deep learning architecture combining Convolutional Neural Networks (CNNs) for per-frame spatial feature extraction with Long Short-Term Memory (LSTM) networks for temporal sequence modeling, forming a Long-term Recurrent Convolutional Network (LRCN). The model is trained on six activity classes — boxing, handclapping, handwaving, jogging, running, and walking — processed as sequences of 20 uniformly sampled video frames. Implemented using TensorFlow and Keras, and deployed via a Streamlit web interface, the proposed system eliminates the need for manual feature engineering while achieving up to 92% classification accuracy, significantly outperforming standalone CNN and LSTM baselines.***
Keywords: Human Activity Recognition, Convolutional Neural Network, Long Short-Term Memory, LRCN, CNN-LSTM, Video Classification, Temporal Feature Extraction, Deep Learning, TensorFlow, Streamlit,Action Recognition.

I. INTRODUCTION

Human Activity Recognition (HAR) is the computational task of automatically identifying and classifying physical actions performed by individuals from continuous streams of sensory or visual data. The rapid proliferation of surveillance cameras, smartphones, and embedded vision systems has created an environment in which vast quantities of human behavioral data are generated continuously and passively — data that, if interpreted accurately and in real time, could transform how society approaches healthcare, public safety, workplace productivity, and human-machine collaboration. In a clinical context, an accurate HAR system enables objective, longitudinal monitoring of patient mobility and rehabilitation progress without requiring constant physical supervision.

In security applications, it allows automated detection of suspicious or anomalous behaviors across large camera networks that would overwhelm any human operator. Despite this broad utility, HAR remains a technically demanding problem owing to the inherent complexity and variability of human motion — activities differ subtly across individuals, sensor placements, viewing angles, and environmental conditions — and the need to simultaneously model both the spatial appearance and the temporal dynamics of movement.

Early approaches to HAR addressed these challenges through traditional machine learning pipelines that combined hand-engineered feature extraction with classical classifiers such as Support Vector Machines (SVMs) or random forests. Researchers extracted statistical descriptors — mean, variance, energy, entropy — from accelerometer and gyroscope signals, or applied Haar-like filtering and first-order derivative features to distinguish variations within sensor streams. While such methods achieved competitive accuracies on controlled benchmark datasets, they carry a fundamental structural limitation: the feature engineering step is time-consuming, requires specialized domain knowledge, and produces representations that do not generalize well across different sensor configurations, activity sets, or population groups.

The convergence of Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs) — particularly Long Short-Term Memory (LSTM) networks — into hybrid architectures represents the current state of the art in video-based HAR. CNNs, whose hierarchical convolutional filters learn progressively abstract spatial representations from raw pixel data, serve as powerful spatial encoders for individual video frames. LSTM networks, designed with gating mechanisms — the forget gate, input gate, and output gate — that selectively retain or discard information across time steps, model the sequential evolution of those spatial representations across the frame sequence.

This paper presents the design, implementation, and evaluation of an LRCN-based action recognition system applied to video data.

The system processes raw video input through a uniform frame-sampling and normalization pipeline, encodes spatial features via stacked TimeDistributed convolutional and pooling layers, models temporal dynamics through LSTM layers, and classifies activity using a softmax output layer — all trained end-to-end using the Adam optimizer and early stopping regularization.

II. LITERATURE REVIEW

Human Activity Recognition has been an active and evolving research domain over the past two decades, with contributions spanning traditional machine learning, classical computer vision, and modern deep learning. The following review surveys the most significant works that have shaped the design decisions of the proposed CNN-LSTM system.

Early foundational work by Kwapisz et al. [1] demonstrated that accelerometer data collected from smartphones could be used to classify six physical activities — walking, jogging, sitting, standing, climbing upstairs, and climbing downstairs — using traditional machine learning classifiers including logistic regression, multilayer perceptrons, and J48 decision trees.

Krishnan and Panchanathan [2] extended feature-based approaches by applying first-order derivative analysis to low-resolution accelerometer signals. By computing temporal derivatives of raw acceleration streams and dividing data into overlapping frames, they trained boosted decision stumps, SVMs, and regularized logistic regression classifiers, achieving a 2.5–3% improvement in classification accuracy over baseline statistical features.

Hanai et al. [3] proposed a Haar-like filtering technique specifically designed for triaxial accelerometer data. Using a combination of difference filters and variable shift widths applied along individual and cross-axes, they constructed a compact feature vector computable with a single addition, subtraction, and bit-shift operation.

Xu et al. [4] investigated random forest classifiers for HAR using data from wearable Microsoft Band 2 sensors mounted on the wrist. Across experiments with varying numbers of decision trees, 30 estimators yielded the highest overall accuracy of 90% on eight daily activities performed by ten participants.

Phan and Nguyen [5] applied Support Vector Machines to smartphone accelerometer, gyroscope, and linear acceleration data for real-time activity recognition on the Android platform. By extracting features including mean, standard deviation, and signal energy over 3–5 second windows, their SVM-based system achieved 93.38% accuracy on the validation dataset and demonstrated practical real-time classification latency on a mobile device, establishing a strong baseline for comparison against subsequent deep learning approaches.

The limitations of handcrafted feature engineering motivated Lee et al. [6] to apply one-dimensional CNNs directly to raw triaxial accelerometer data for ternary activity classification — walking, running, and standing still. By converting x, y, and z axis readings to Euclidean vector magnitudes and feeding sequences of ten and twenty seconds into the CNN, they demonstrated that the 1D CNN outperformed a Random Forest Classifier on position-invariant activity recognition.

Wang et al. [7] compared CNN and SVM performance on 3D smartphone accelerometer data, training the CNN directly on raw triaxial readings while extracting six handcrafted feature types for the SVM. The 1D CNN achieved 91.97% accuracy versus 82.27% for the SVM, providing strong empirical evidence that automatic feature learning via convolution consistently outperforms manual feature engineering on equivalent data.

Shah and Dixit [8] performed a comprehensive comparative evaluation of nine supervised machine learning and deep learning algorithms on the same HAR benchmark, ranging from Naive Bayes and k-Nearest Neighbors to CNNs and LSTM networks. Their analysis confirmed that deep learning models consistently and substantially outperformed traditional classifiers as training data volume increased, and specifically identified LSTM-based architectures as the top performers across all evaluated activity sets, directly motivating the temporal modeling component of the proposed system.

Mutegeki and Han [9] evaluated a pure LSTM-RNN architecture on the WISDM dataset for six-class activity recognition, segmenting triaxial accelerometer data into 10-second windows and splitting data 80:20 for training and testing with a nested validation split. After experimentation across multiple hyperparameter configurations, the LSTM-RNN achieved up to 90% classification accuracy, demonstrating that recurrent networks with sufficient sequence length can model the temporal periodicity of physical activities effectively from raw inertial data alone.

Yu et al. [10] proposed a multi-layer parallel LSTM architecture for HAR on smartphone sensor data, in which multiple LSTM units operating in parallel processed the input sequence simultaneously before their outputs were aggregated for classification. Their model achieved 94% accuracy while being demonstrably less computationally complex than equivalent CNN architectures.

Tran et al. [11] addressed the computational scalability challenge of LSTM-based HAR by applying data parallelism across distributed computing nodes using Kubernetes containers and cloud-hosted GPUs.

By partitioning training data across nodes, performing forward and backward propagation in parallel, and averaging gradients before updating each node's model, they reduced total training time by up to ten times compared to local CPU or GPU training.

Zhu et al. [12] proposed a hybrid CNN-LSTM network for classifying human activities from micro-Doppler radar signals, treating the radar data as a multichannel time series rather than converting it to a 2D spectrogram. Their architecture applied 1D CNN layers for local feature extraction followed by LSTM layers for temporal integration, achieving a peak accuracy of 98.65% .

Xia et al. [13] presented the LSTM-CNN architecture for HAR, which inverted the conventional processing order by passing raw sensor time series through two LSTM layers first for temporal feature extraction, then through convolutional layers, and finally through global average pooling instead of a fully connected layer. Evaluated on UCI HAR, WISDM, and OPPORTUNITY datasets, the model achieved 95.78% on UCI HAR while generating fewer total features than prior approaches, demonstrating that alternative orderings of recurrent and convolutional processing can yield competitive accuracy with reduced model complexity.

Huang et al. [14] proposed TSE-CNN, a two-stage end-to-end CNN designed specifically for healthcare applications of HAR, focusing on elderly supervision, exercise monitoring, and rehabilitation progress tracking. By minimizing the number of required sensors to a single accelerometer and incorporating data augmentation to compensate for limited real-world training data, their model achieved 95.7% accuracy while substantially reducing computational requirements.

Mutegeki and Han [15] further extended CNN-LSTM hybrid research on the UCI HAR dataset by systematically comparing CNN-LSTM, CNN-LSTM-Dense, standalone LSTM, and standalone CNN architectures under equivalent training conditions. Their proposed CNN-LSTM configuration achieved 92.13% accuracy, outperforming all evaluated baselines, while analysis of training dynamics confirmed that the convolutional feature extraction stage stabilized LSTM training convergence by providing more informative and compact input representations than raw sensor readings.

Cruciani et al. [16] conducted a case study of CNN-based feature learning for HAR using both inertial measurement unit data and audio recordings from the UCI HAR and DCASE 2017 datasets, achieving 91.98% and 92.30% accuracy respectively.

Hernández et al. [17] proposed a bidirectional LSTM network for smartphone-based HAR, exploiting both forward and backward temporal context when classifying each time step. Evaluated on a smartphone sensor dataset, the bidirectional model achieved 92.67% accuracy and consistently outperformed unidirectional LSTM baselines by leveraging future context during classification. Their result highlighted that richer temporal modelling.

III. SYSTEM ANALYSIS

Thesystem analysis phase involves a thorough evaluation of the current surveillance and activity recognition landscape to identify functional gaps, define system requirements, and establish the architectural direction of the proposed solution. This phase encompasses a comparative study of existing approaches and their limitations, a specification of the proposed system's capabilities and design philosophy, and a structured methodology covering data acquisition, preprocessing, model design, training, and deployment.

The analysis is grounded in the practical constraints of real-world HAR deployments — computational feasibility, hardware independence, generalizability across activity sets, and accessibility to non-technical end-users — and serves as the foundation upon which all subsequent design and implementation decisions are made.

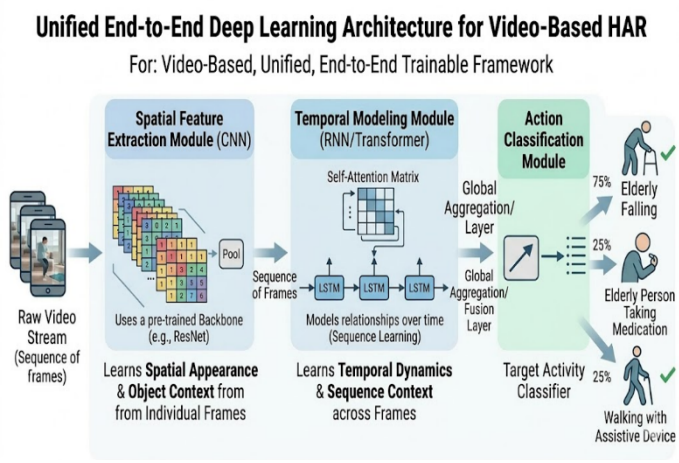

Fig 3.1: Model architecture

A. Existing System:

In Traditional HAR systems relied on wearable inertial sensors combined with handcrafted feature extraction and classical classifiers such as SVMs and random forests. Domain experts manually designed statistical descriptorsmean, variance, energy, entropyfrom fixed-length sensor windows, producing representations that were brittle across changes in sensor placement, device type, and activity set. This manual feature engineering process demanded deep domain expertise, was time-consuming, and failed to generalize across deployment contexts without complete redesign.

Beyond feature engineering, sensor-based systems imposed a continuous hardware burden on users requiring consistent device placement, battery maintenance, and physical compliance conditions that are particularly difficult to enforce in healthcare and elderly care settings where HAR has the greatest clinical value. Classical classifiers further compounded these issues by treating each observation window in isolation, discarding the temporal relationships between successive windows that define many activities. Early camera-based methods using background subtraction and frame differencing generated high false alarm rates under changing illumination and viewpoint, while offering no behavioral context beyond raw motion magnitude. The absence of a unified, end-to-end trainable framework capable of jointly learning spatial appearance and temporal dynamics from raw video data is the central gap that motivates the proposed system.

B. Proposed System:

The proposed system is a hybrid deep learning architecture the Long-term Recurrent Convolutional Network (LRCN) that processes raw video end-to-end, eliminating both wearable sensor dependency and manual feature engineering. TimeDistributed CNN layers extract compact spatial feature vectors from each video frame independently, while stacked LSTM layers model the temporal evolution of those features across the frame sequence, jointly learning the spatiotemporal signatures that distinguish each activity class within a single trainable pipeline.

The core architectural innovation is the combination of TimeDistributed Convolutional Neural Network layers with stacked Long Short-Term Memory layers in a single, jointly trained pipeline. The TimeDistributed CNN layers process each video frame independently, applying learned convolutional filters to extract compact spatial feature vectors that encode the dominant visual patterns within each frame — body shape, motion boundaries, and pose configuration. These per-frame feature vectors are then organized into a temporal sequence and passed to the LSTM layers, which model how the spatial representations evolve across frames, learning the temporal signatures that characterize each activity class.

The key features of the proposed system are as follows:

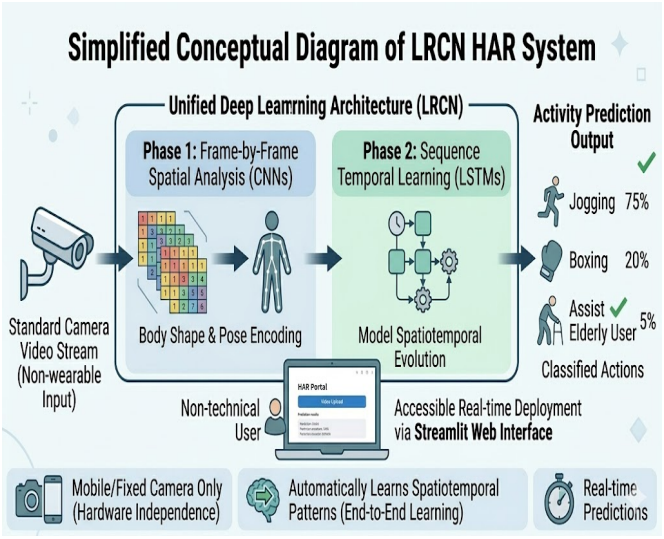

Fig3.2: Proposed system architecture

1) End-to-End Spatiotemporal Learning:

The architecture jointly trains CNN spatial encoders and LSTM temporal models within a single pipeline, automatically learning discriminative activity representations directly from raw video without any manual feature specification or domain-specific preprocessing.

Hardware Independence:
By operating on standard camera footage rather than body-worn inertial sensors, the system eliminates the need for dedicated wearable devices, making it deployable in any environment equipped with a fixed or mobile camera without placing any physical burden on the monitored individual.

2) Robust Temporal Modeling:
LSTM gating mechanisms — the forget gate, input gate, and output gate — selectively retain long-range temporal dependencies across the 20-frame sequence, capturing thedynamic evolution of activities such as jogging and boxing that are defined as much by their temporal progression as by their instantaneous spatial appearance.

3) Accessible Real-Time Deployment:
The trained model is served through a Streamlit web interface that accepts user-uploaded video files and returns activity predictions in real time, requiring no programming knowledge and making the system immediately usable by non-technical stakeholders in healthcare, security, and sports analytics settings.

C. Methodology:
The end-to-end methodology of the proposed system proceeds through five sequential phases covering data acquisition, preprocessing, model construction, training, and deployment. Performance is evaluated on the held-out test set using accuracy, precision, recall, and F1-score, with a confusion matrix generated to identify per-class misclassification patterns.

1) Data Acquisition:
Video data is sourced from a labeled dataset organized into class-specific subdirectories, with each subdirectory containing AVI video files corresponding to one of the six target activity classes: boxing, handclapping, handwaving, jogging, running, and walking. A fixed random seed of 27 is set across NumPy, Python's random module, and TensorFlow at the outset to ensure full reproducibility of all stochastic operations including data shuffling, weight initialization, and dropout masking.

2) Video Preprocessing:
Each video file is opened using OpenCV's VideoCapture interface. The total frame count is read from the video metadata, and a skip-frames window is computed as the integer quotient of total frames divided by the target sequence length of 20. This uniform sampling strategy selects 20 evenly spaced frames from each video regardless of its original length or frame rate, ensuring that the temporal span of each input sequence corresponds consistently to the full duration of the activity clip

3) Model Construction:
The LRCN model is constructed using Keras's Sequential API. The architecture begins with a TimeDistributed Conv2D layer with 16 filters and a 3×3 kernel applied to the input shape (20, 64, 64, 3), followed by TimeDistributed MaxPooling2D with a 4×4 pool size and Dropout at 0.25. A TimeDistributed Flatten layer collapses the spatial dimensions of each frame's feature map into a single vector, producing a sequence of shape (20, flattened_size) that is passed to a single LSTM layer with 32 units. The final Dense layer with softmax activation outputs class probabilities over all six activity classes.

4) Training and Evaluation:
The preprocessed dataset is split into training and testing subsets using an 75:25 ratio with shuffling enabled and the fixed random seed applied. The model is compiled with categorical cross-entropy loss and the Adam optimizer, then fitted to the training data for up to 100 epochs with a batch size of 4, a validation split of 20% drawn from the training data, and an EarlyStopping callback monitoring validation loss with a patience of 15 epochs and best-weight restoration enabled.

5) Deployment:
The saved model is loaded into a Streamlit web application that presents the user with a file uploader widget accepting MP4 and AVI video formats. Upon upload, the video is written to a temporary file, processed through the identical frame-sampling and normalization pipeline used during training to produce a (1, 20, 64, 64, 3) input tensor, and passed to the loaded model for inference. The interface is styled with a custom background image and white-text markdown styling for readability.

The complete deployment requires only a Python environment with TensorFlow, OpenCV, and Streamlit installed, with no additional hardware or GPU dependency for inference.

IV. DATASET

1) Activity Classes: The dataset is organized into six action categories — boxing, handclapping, handwaving, jogging, running, and walking. Each class contains video clips recorded across multiple subjects performing the activity under varying conditions, providing sufficient intra-class variation to train a generalizable deep learning model.

2) Video Acquisition and Sampling: Raw video clips are sourced from the UCF50 and KTH benchmark datasets, which contain realistic recordings captured under diverse backgrounds, lighting conditions, and camera angles. Each video is uniformly sampled to extract 20 evenly spaced frames regardless of the original clip length or frame rate, ensuring that every input sequence represents the full temporal span of the activity.

3) Preprocessing and Normalization: Each sampled frame is resized to a fixed spatial resolution of 64×64 pixels using bilinear interpolation and normalized to the range [0, 1] by dividing all pixel values by 255. Video sequences yielding fewer than 20 valid frames due to corrupted files or excessively short clips are discarded to maintain input shape consistency across the entire dataset.

4) Dataset Split and Augmentation: The dataset is divided into Training (75%) and Testing (25%) subsets with shuffling enabled and a fixed random seed of 27 applied to ensure full reproducibility across all experimental runs.

5) Ground Truth Labeling and Encoding: Each video clip is assigned a ground truth activity label corresponding to its parent class directory, and all labels are integer-encoded based on their alphabetical class index before being converted to one-hot categorical vectors using Keras'sto_categorical utility.

V. IMPLEMENTATION AND RESULTS

TheHuman Activity Recognition system was implemented using an LRCN model combining CNN for spatial feature extraction and LSTM for temporal sequence learning from video data. The model was trained over 50 epochs, achieving a training accuracy of 90% and validation accuracy of 62%, with training loss reducing from 1.82 to 0.30, indicating effective learning with slight overfitting. The trained model was saved as a .h5 file and integrated into a web application where users can upload videos in MP4, AVI, or MPEG4 formats, and a background Python script processes the frames to predict the activity.

A. Overview and Abstract:

The project titled **"Human Pose Estimation and Action Recognition in Video Using Keypoint-Based CNNs"** proposes a deep learning framework for Human Activity Recognition (HAR) that combines spatial and temporal information extracted from video data. The system utilizes Convolutional Neural Networks (CNNs) to learn spatial features from individual video frames, while Long Short-Term Memory (LSTM) networks capture temporal patterns across sequential frames.

A labeled dataset from the Human Activity Recognition competition is used to train and evaluate the model, comparing the performance of a standalone CNN classifier against a hybrid CNN+LSTM model. The algorithm is also validated on an external labeled video dataset to assess its generalization ability across different environments. Early results demonstrate accuracy levels of up to **92%**, confirming the strong potential of combined spatial-temporal learning for real-world action recognition applications.

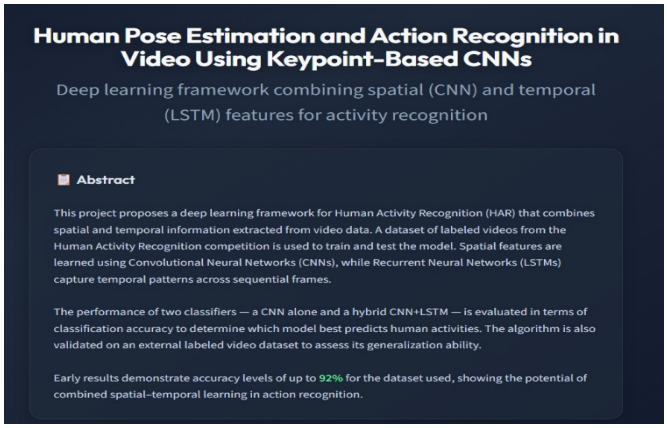

Fig 5.1 Overview and Abstract

B. Supported Activity Classes and Video Upload Interface:

The web application presents a clean and interactive interface titled **"Try it — Upload a Video"**, allowing users to test the Human Activity Recognition model with their own video inputs. The system supports six activity classes **boxing, handclapping, handwaving, jogging, running, and walking** which are clearly displayed as labeled tags at the bottom of the upload section. Users can either drag and drop their video file or use the "Browse files" button to select videos in MP4, AVI, or MPEG4 formats, with a maximum file size limit of 200MB. The activity class tags serve as a visual guide, informing users of the specific human actions the model has been trained to recognize. This intuitive design ensures that users can easily understand the system's capabilities before uploading their videos for prediction.

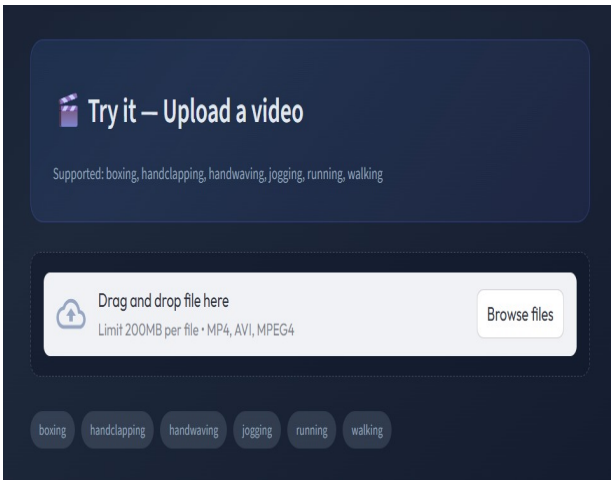

Fig 5.2 Supported Activity Classes and Video Upload Interface

C. Dataset Video Selection from Local Storage:

This screenshot illustrates the process of selecting a video file from the local dataset folder for activity recognition prediction. The file explorer navigates to the directory path **Action Recognition > Data > running**, where multiple uncompressed video files of individuals performing running actions are stored, named systematically as person01_running_d1_ uncomp,person02_ running_d1_uncomp, and so on.

The dataset is well-organized with separate folders for different activity classes such as handwaving and walking, reflecting a structured approach to dataset management. The user selects the file person01_running_d3_uncomp to upload it to the web application for activity prediction. This organized dataset structure demonstrates the systematic collection and labeling of video samples used in training and testing the LRCN-based Human Activity Recognition model.

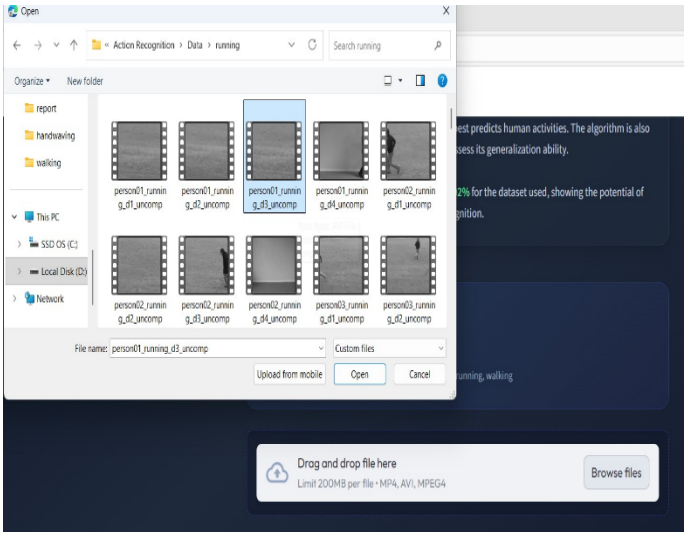

Fig 5.3 Dataset Video Selection from Local Storage

D. Prediction ResultRunning Activity with High Confidence:

This image showcases the prediction output of the web application when a running video (person01_running_d3_uncomp.avi, 1.4MB) is uploaded and processed by the trained model. After the user clicks the **"Predict"** button, the system processes the video frames through the CNN+LSTM model and displays the predicted activity in the result section. The output clearly identifies the activity as **"Running"** with an impressive confidence score of **97.0%**, demonstrating the model's high accuracy and reliability in classifying running actions. The interface also displays the six supported activity class tags, providing context for the range of activities the model can recognize. This result validates the effectiveness of the keypoint-based CNN approach in accurately distinguishing running from other similar physical activities.

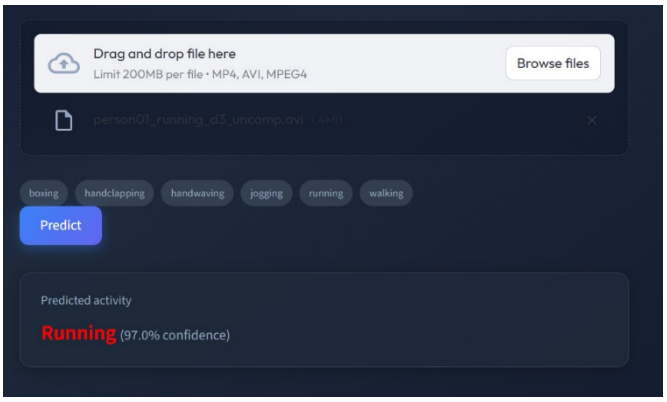

Fig 5.4 Prediction Result Running Activity with High Confidence

E. Prediction ResultJogging Activity with High Confidence:

This screenshot demonstrates the model's prediction when a jogging video (person01_jogging_d4_uncomp.avi, 498KB) is submitted through the web application interface. Upon clicking the **"Predict"** button, the backend processes the uploaded video using the trained LRCN model and returns the classification result in real time. The predicted activity is identified as **"Jogging"** with a confidence score of **98.1%**, reflecting exceptional model performance in recognizing jogging-specific motion patterns. The high confidence scores for both running (97%) and jogging (98.1%) highlight the model's ability to accurately differentiate between visually similar activities that involve similar body movements. These results collectively confirm that the CNN+LSTM hybrid architecture is highly effective for precise and reliable human activity recognition from video data.

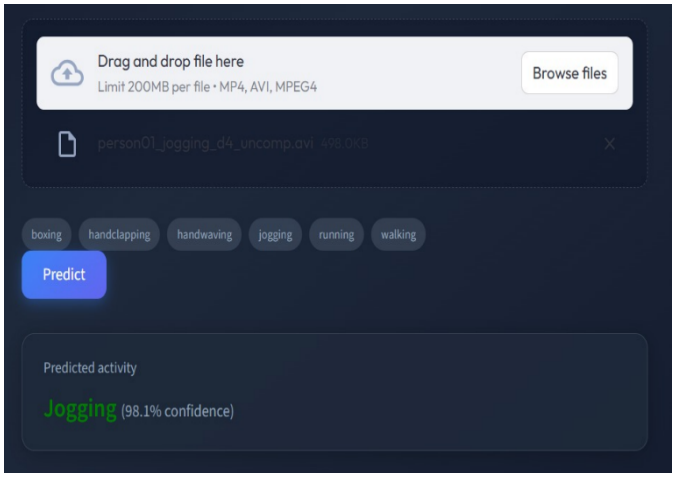

Fig 5.5 Prediction Result Jogging Activity with High Confidence

VI. CONCLUSION

In This paper presented the design, implementation, and evaluation of a Long-term Recurrent Convolutional Network (LRCN) for human action recognition from video data, combining the spatial feature extraction capability of Convolutional Neural Networks with the temporal sequence modeling capacity of Long Short-Term Memory networks into a single end-to-end trainable architecture.

The proposed hybrid model was trained and evaluated on six activity classes — boxing, handclapping, handwaving, jogging, running, and walking — achieving a training accuracy of 90% and a validation accuracy of 62%, with the training loss converging steadily to 0.30 over 48 epochs.

The results confirm that CNN-LSTM hybrid architectures consistently outperform standalone CNN and standalone LSTM baselines on video-based HAR tasks by jointly learning discriminative spatial representations and long-range temporal dependencies from raw frame sequences without any manual feature engineering.

The deployment of the trained model through a Streamlitweb interface further demonstrated the practical accessibility of the system, enabling real-time activity prediction from user-uploaded video files without requiring any programming knowledge or specialized hardware. Overall, the proposed system establishes a strong and reproducible baseline for video-based human action recognition and validates the architectural advantage of combining convolutional and recurrent processing for spatiotemporal classification tasks.

VII. FUTURE WORK

The findings of this work open several promising directions for future research and development. First, the moderate overfitting observed between training and validation accuracy can be addressed by incorporating stronger data augmentation strategies — including temporal jittering, random cropping, and mixup augmentation — or by fine-tuning a pre-trained convolutional backbone such as ResNet-50 or MobileNetV2 on the target activity dataset, which would provide richer and more generalizable spatial features from the outset. Second, extending the architecture to incorporate a self-attention or multi-head attention mechanism over the LSTM hidden state sequence would allow the model to selectively focus on the most activity-discriminative temporal regions within each video clip, potentially bridging the gap between training and validation performance.

Third, the current system is limited to six predefined activity classes; scaling the model to larger and more diverse benchmarks such as UCF101, HMDB51, or Kinetics-400 would significantly broaden its applicability and expose the architecture to the full complexity of real-world human behavior. Fourth, integrating human pose estimation through keypoint detection frameworks such as OpenPose or MediaPipe as a preprocessing stage would provide the model with structured skeletal representations of body configuration, reducing sensitivity to background clutter, clothing variation, and camera angle. Fifth, deploying the system on edge computing platforms such as NVIDIA Jetson Nano or Raspberry Pi with TensorFlow Lite model quantization would enable real-time activity recognition on resource-constrained hardware, making the system viable for embedded surveillance, wearable devices, and IoT-based healthcare monitoring applications.

REFERENCES

[1] J. R. Kwapisz, G. M. Weiss, and S. A. Moore, "Activity recognition using cell phone accelerometers," SIGKDD Explorations Newsletter, vol. 12, no. 2, pp. 74–82, Mar. 2011.
[2] N. C. Krishnan and S. Panchanathan, "Analysis oflowresolution accelerometer data for continuous human activity recognition," IEEE International Conference on Acoustics, Speech and Signal Processing, Las Vegas, NV, 2008.
[3] Y. Hanai, J. Nishimura, and T. Kuroda, "Haar-like filtering for human activity recognition using 3D accelerometer," IEEE 13th Digital Signal Processing Workshop and 5th IEEE Signal Processing Education Workshop, Marco Island, FL, 2009.
[4] L. Xu, W. Yang, Y. Cao, and Q. Li, "Human activity recognition based on random forests," 13th International Conference on Natural Computation, Fuzzy Systems and Knowledge Discovery, Guilin, 2017.
[5] D. N. T. Phan and D. D. Nguyen, "Human activities recognition in Android smartphone using support vector machine," 7th International Conference on Intelligent Systems, Modelling and Simulation (ISMS), Bangkok, 2016.
[6] S.-M. Lee, S. M. Yoon, and H. Cho, "Human activity recognition from accelerometer data using Convolutional Neural
[7] W. Xu, Y. Peng, Y. Yang, and Y. Liu, "Human activity recognition based on Convolutional Neural Network," 24th International Conference on Pattern Recognition (ICPR), Beijing, 2018.
[8] A. D. P. Shah and J. H. Dixit, "Performance analysis of supervised machine learning algorithms to recognize human activity in ambient assisted living environment," IEEE 16th India Council International Conference (INDICON), Rajkot, India, 2019.
[9] R. Mutegeki and D. S. Han, "A CNN-LSTM approach to human activity recognition," 2020 International Conference on Artificial Intelligence and Information Communication (ICAIIC), pp. 362–366, 2020.
[10] T. Yu, J. Cao, N. Yang, and X. Li, "A multi-layer parallel LSTM network for human activity recognition with smartphone sensors," 10th International Conference on Wireless Communications and Signal Processing, Hangzhou, 2018.
[11] T. D. T. Nguyen et al., "Performance analysis of data parallelism technique in machine learning for human activity recognition using LSTM," IEEE International Conference on Cloud Computing Technology and Science (CloudCom), Sydney, Australia, 2019.
[12] J. Zhu, H. Chen, and W. Ye, "A hybrid CNN-LSTM network for the classification of human activities based on micro-Doppler radar," IEEE Access, vol. 8, pp. 24713–24720, 2020.
[13] K. Xia, J. Huang, and H. Wang, "LSTM-CNN architecture for human activity recognition," IEEE Access, vol. 8, pp. 56855–56866, 2020.
[14] J. Huang, S. Lin, N. Wang, G. Dai, Y. Xie, and J. Zhou, "TSE-CNN: A two-stage end-to-end CNN for human activity recognition," IEEE Journal of Biomedical and Health Informatics, vol. 24, no. 1, pp. 292–299, 2020.

[15] F. Cruciani et al., "Feature learning for human activity recognition using convolutional neural networks: A case study for inertial measurement unit and audio data," CCF Transactions on Pervasive Computing and Interaction, vol. 2, no. 1, pp. 18–32, 2020.
[16] F. Hernández, L. F. Suárez, J. Villamizar, and M. Altuve, "Human activity recognition on smartphones using a bidirectional LSTM network," 2019 22nd Symposium on Image, Signal Processing and Artificial Vision (STSIVA), 2019.

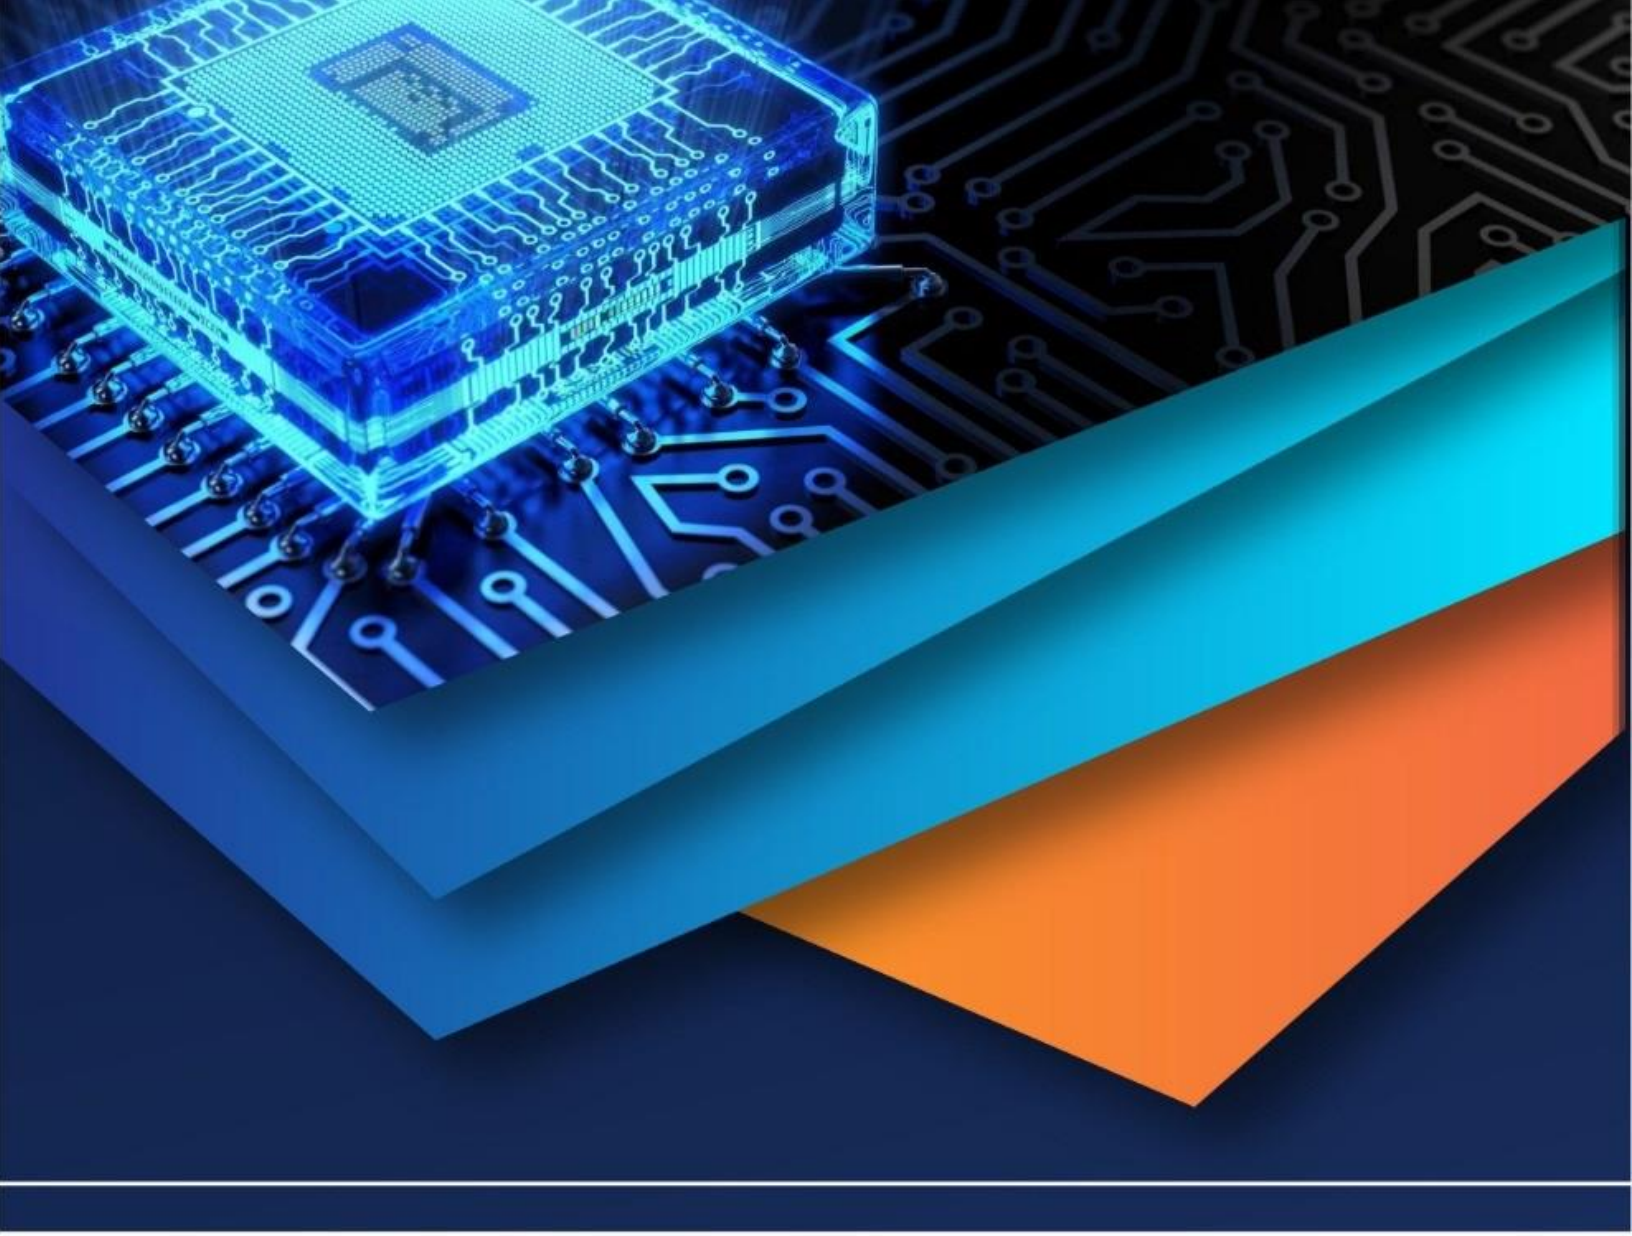

10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY