

Improving Human Action Recognition Through Cross-View Camera-Aware Augmentation

Zahid Hussain Khaskheli ^{a,b,*}, Qasim Ali ^b, M. Shehram Shah ^b, Din Muhammad ^b, Farast Ali ^c, Hina Ali ^b, Maria Khalid ^d, Zoha Waheed ^e, Umer Gaffar ^f

^a *PhD Scholar at the Institute of Information and Communication Technologies (IICT), Mehran University of Engineering and Technology, Jamshoro 76060, Pakistan*

^b *Department of Software Engineering, Mehran University of Engineering and Technology, Jamshoro 76060, Pakistan*

^c *Department of Computer Science, Bahria University, Karachi 75260, Pakistan*

^d *Islamabad Model College for Girls, Islamabad 44000, Pakistan*

^e *Computational Science and Technology, National University of Sciences and Technology (NUST), Islamabad 4000, Pakistan*

^f *Department of Statistical and Actuarial Sciences, University of the Punjab, Lahore 54590, Pakistan*

* Corresponding author: zahid.khaskheli@faculty.muet.edu.pk

Received: 22 May 2026, Revised: 06 June 2026, Accepted: 22 June 2026, Published: 01 July 2026

KEY WORDS

Human Action Recognition
3D Neural Networks
Deep Learning
Data Augmentation
Cross-View Action
Recognition
Spatiotemporal Feature
Learning

ABSTRACT

Human Action Recognition (HAR) is vital in surveillance, sports analysis, and Human Computer Interaction (HCI). Publicly available benchmarks mostly suffer from limited sample sizes, class imbalance, occlusions, and viewpoint variability, which may affect model generalization. Most existing HAR augmentation pipelines rely on generic or heuristic transformations that do not explicitly model camera-induced variations. To address these challenges, we propose an Augmented Data Generator Framework (ADGF) for HAR that integrates multiple parameter-controlled augmentation strategies designed to approximate realistic camera-induced variability while preserving temporal consistency. Our method integrates four transformation strategies. Each augmentation technique is parameterized across validated variation ranges, realistically simulating camera viewpoint changes, motion diversity, and partial occlusions, while preserving frame-to-frame temporal coherence compared with uniformly applied conventional augmentations. ADGF distributes controlled transformations to maintain a balance and preserve action semantics. We evaluate performance with standard metrics, and perform an ablation study to measure each augmentation's contribution. With the combined augmented set, a top-1 accuracy on HMDB51 rises from 87.5% (baseline) to 97.7%, and final training loss falls from 0.21 to 0.12. Required epochs drop from 38 to 25 (~34% fewer epochs). Cross-dataset testing on UCF10 demonstrates improved robustness, with accuracy increasing from 85.1% to 96.5%. The ablation study shows that each augmentation yields a positive gain with viewpoint variation, that occlusion simulation produces the largest single-technique improvements, and that the full combination produces the strongest transferability. These results indicate that a balanced structured augmentation strategy substantially improves both in-dataset performance and cross-dataset robustness. Implementation details, augmentation configurations, and experimental settings will be made publicly available upon publication to support reproducibility.

1. Introduction

Human Action Recognition (HAR) is a core research area in computer vision with a wide range of practical applications, such as video surveillance, healthcare monitoring, sports analysis, and human-computer interaction [1]. A reliable HAR system can strengthen public safety by improving surveillance, assist doctors with remote patient monitoring, and support the design of smarter assistive technologies for daily life[2], [3]. HAR has important applications such as in healthcare, fall detection, rehabilitation assessment, and assisted living, where robustness to occlusions, viewpoint changes, and motion variability is essential for reliable deployment. As deep learning has advanced and 3D CNNs have emerged, it can effectively capture spatiotemporal features [4], [5]. But there are still certain challenges when implementing HAR systems in real time, mostly due to limited data availability [6]. HMDB51 is a popular and most commonly used benchmark dataset, which provides 51 action class categories and 6,766 clips that are collected from various online sources, such as movies, YouTube videos, and other online platforms. While HMDB51 is valuable, it has certain limitations that limit its performance, such as its small size, uneven distribution across classes, variations in how the same action may appear, subject occlusion, and high dependence on viewpoint changes. As 3D CNN models are data-intensive, they require sufficient data to train; these constraints can be critical and lead to problems such as overfitting, reduced robustness, and poor generalization. Therefore, with great advancements in deep learning architectures, spatiotemporal features extraction and generalization can be constrained by the datasets available, not in terms of size but diversity. As in real-time deployment scenarios, HAR systems may operate under varying conditions such as changes in camera viewpoints, partial obstructions, motion inconsistencies, and cluttering. We address these limitations through a structured augmentation framework designed specifically for video-based human action recognition. The standard augmentations, such as flipping, rotating, scaling, and resizing, because human actions are complex, it involves body motion over time, unlike static images. Hence, we are applying four advanced augmentations, herein after called: Viewpoint, Elastic Deformation, Random erasing, and Occlusion simulation. This process not only creates synthetic samples of data but also applies controlled transformations to existing HMDB51, preserving action semantics, under specific parametric settings, during the experimentation phase to ensure the realism and quality of the data samples. Much of the work on augmentation has focused on image data, while very

few studies have explored advanced data augmentation for HAR. Several studies have demonstrated that augmentation techniques tailored to motion dynamics and action-specific variations can significantly enhance model robustness. For instance, researchers have investigated temporal jittering, optical flow-based augmentation, and GAN-based synthetic video generation to improve HAR model generalization. However, these approaches either require additional computational resources or fail to address key challenges such as occlusions and viewpoint variations with variable conditions. Fewer studies have explored and used basic data augmentation strategies such as flipping, cropping, blurring, and rotation for motion data rather than focusing on static images. Still, we present a systematic synthetic dataset generation framework for HAR, leveraging advanced augmentation techniques specifically designed to enhance dataset diversity and model robustness for dynamic vision data (videos), making designed to approximate realistic camera-induced variability by experimenting and generating simulated samples with varied inputs, rather than applying simple techniques directly. Rather than merely applying disjointed, off-the-shelf image transformations, the scientific novelty of this work lies in the architectural design of a cohesive, realism-aware framework that mathematically models camera-induced physical artifacts. By linking specific parameter limits to physical constraints, such as viewpoint variations mapping to physical surveillance installations, and occlusion simulation mapping to dynamic foreground blocks, we establish a deterministic pipeline that prevents semantic destruction while maximizing data entropy. The contribution focuses on a structured integration strategy that balances dataset expansion without introducing synthetic bias or distribution shifts, proving that realism-focused geometric diversity is significantly more critical for safety-critical deep learning models than increased structural complexity. Elastic Deformation applies small and smooth bending or stretching to video frames and simulates the natural variations of how people perform the same action. Helps the model to appreciate the fact that the same action may manifest in a somewhat different manner from person to person. Random Erasing randomly masks out small regions from video frames, as if they were temporarily occluded or missing. It builds a model that uses overall motion patterns rather than individual image regions (visually) and enhances robustness when body parts are partly occluded.

Occlusion Simulation simulates real situations in which other people or objects partially occlude actions. The training with such occlusions enables the

model to learn even if a portion of the body is not seen. Viewpoint Angle varies the angle or view of the action; the same action may appear very different depending on the angle of the view. Multiple views of the model decrease the need for a set camera orientation.

The following are key contributions of our study that are summarized as follows:

- We propose a structured augmentation framework for HAR that combines multiple complementary video augmentation strategies while preserving temporal coherence and action semantics.
- The framework employs parameter-controlled augmentation ranges and balanced augmentation distribution to reduce unrealistic transformations and augmentation dominance.
- We evaluate the framework using 3D ResNet-18 on HMDB51 and further assess robustness through cross-dataset evaluation on UCF101.
- Experimental results demonstrate consistent improvements in recognition accuracy, convergence behaviour, and cross-dataset robustness under controlled augmentation settings.

The contribution of this work does not lie in introducing entirely new augmentation operators, but rather in the structured integration, parameter-controlled configuration, and robustness-oriented application of complementary augmentation strategies for video-based human action recognition.

Although much progress has been reported, HAR models struggle to achieve satisfactory generalization, robustness, and efficiency. Many videos look similar as some actions appear only from a few camera angles, and are often blocked or partially hidden. Some action classes have many videos, while others have very few. Therefore, models trained on simple HMDB51 often perform well in tests but struggle when used in real-world situations. These limitations are particularly problematic for healthcare-oriented HAR systems, where the camera placement, partial visibility of patients, and non-standard movement patterns are common and not controlled easily. Although traditional data augmentation techniques have been explored, they often fail to introduce the level of variability that is required for effective generalization. These situations can be problematic for healthcare settings, where cameras are often placed in fixed positions, patients may be partly hidden, and their movements are not always regular or planned. Although standard data augmentation methods have been reported in the literature, they have proved to be helpful for static image data, while less is explored for dynamic image frames, as they usually do not add

© Mehran University of Engineering and Technology 2026

enough realistic variation to help the model work well in such conditions. Hence, a synthetic dataset generated via AGDF results in improved robustness and a well-generalized model.

The goal is to make the videos look realistic. Hence, each augmentation is applied using carefully chosen limits that were tested before. At the same time, we make sure that the frames stay consistent over time, so the action does not look broken or unnatural. Many existing methods use generic augmentations that are not designed for action videos. Our approach is different because it focuses on problems specific to HAR. Therefore, the synthetic samples generated so far help the model learn stronger and more adaptable features, instead of memorizing simple patterns. We do a test of a synthetic dataset on 3D ResNet-18 by training and evaluating for HAR. During these tests, we not only check the accuracy of the model, but also see how smoothly the training process will run, how quickly the model learns, and how well it performs on different data it has never seen before. The effect of new synthetic data and its contribution to the model's learning is clearly visible in the results, which is especially pronounced with a smaller amount of training data. This study also shows that data augmentation is crucial to enhancing the reliability of HAR research in limited data scenarios. It also demonstrates that the correct formulation of synthetic data can overcome the most popular issues like overfitting, class imbalance, and data variation from different points of view, for instance, camera angle shift. Future research will build on this work through synthetic augmentation using handcrafted feature extraction and the use of more sophisticated deep learning architectures. We will be sharing our full dataset and augmentation pipeline to further inspire innovation in the HAR research community, so that the community can continue to improve. The remaining part of this paper is organized as follows: In the section "Background and related works", we present different related studies on human action recognition and state-of-the-art presented in the literature. The "Methodology" section provides a comprehensive view of our proposed method, providing details of its pipeline and tools that have been used to implement it. In the "Results and Analysis" section, we have provided an extensive analysis and evaluation of the proposed framework based on several evaluation metrics, with original, augmented, and cross-dataset evaluations. In the last "Conclusion", we summarize the key findings, highlight the major contributions of our study, and

explore the possible future directions in the field of HAR.

2. Related Work

With increased demand in applications ranging from security and video surveillance to healthcare monitoring and interactive technologies, such as robotics, HAR has received more attention from researchers and computing scientists [12,13]. Over the last two to three decades, deep learning techniques have greatly advanced HAR systems, enabling the detection not only of complex motion patterns but also of objects and relationships within a scene [15]. In this literature review, we have focused on existing works to see how 3D Convolutional Neural Networks (3D CNNs) can be combined with advanced data augmentation to improve robustness, generalization, and overall performance. In particular, authors in recent works have highlighted the importance of synthetic dataset generation, motion-centered transformations, and hybrid model designs to overcome the persistent challenges that are faced by the HAR model [16]. Authors have shown a historical perspective on how HAR systems relied on handcrafted features such as Histogram of Oriented Gradients (HOG) and Spatiotemporal Interest Points (STIP) [17,18,19]. Authors show these hybrid methods allowed researchers to capture some basic action characteristics, but they struggled with large-scale variability, occlusion, and complex backgrounds. The invention of modern deep learning algorithms has entirely changed the way we accomplish such tasks. The development of Convolutional Neural Networks and Recurrent Neural Networks has completely changed the mechanism of feature learning and achieved improved performance gains [20]. Among these methods, 3D CNNs stand out as particularly effective, extending convolution into the temporal domain, making them able to capture both spatial and temporal cues. This makes them a well-suited choice for recognizing actions as they unfold over time.

Authors report that even after huge progress in deep learning, HAR models still struggle, especially regarding generalization. Most models perform well on benchmark datasets but often struggle when they are applied to unseen or real-world scenarios, and they can't give the expected outcomes. This weakness largely arises from the dataset's limitations, its diversity, and a few constraints. The most popular datasets, such as UCF101 and Kinetics-400, which are reported by many authors to be used for training and evaluation, still cannot fully represent the challenges of diverse real-world conditions [21]. HMDB51 is

mainly designed to understand and analyze human motion patterns. The authors quote about its valuable diversity, but the limitations in sample size, strong background clutter, and high inter-class similarity create problems for model learning[22,23]. Such weaknesses may cause model overfitting and reduce its ability to generalize beyond training data.

Researchers have therefore suggested and turned to data augmentation as a practical strategy to expand the sample size and reduce overfitting and generalization problems. Data augmentation artificially increases training variability by applying transformations to existing data while maintaining semantic meaning [24]. In traditional computer vision, this includes simple manipulations such as flipping, rotating, or scaling images. While useful in image classification, these techniques do not translate easily to HAR. Because of the dynamic nature of human actions, it requires methods that can account for changes across frames, rather than static modifications, which is suitable in the case of images.

Conventional exploration of temporal adjustments has historically relied on simplistic data mutations. For instance, frame jittering works by inserting local temporal variance but fails to preserve global sequence semantics. The authors in the following studies used conventional techniques such as frame jittering, sequence reordering, and random frame dropping to create variability in motion representation[25,26]. These simple variations help models adapt only to minor temporal shifts in action execution, whereas the HAR requires increasingly realistic modifications. In [27], they have also explored optical flow-based augmentations that introduce variations that are motion sensitive.

But these approaches still have limited capability to solve the issue at a deeper level and still face challenges to solve the common real-world videos, such as partial obstruction where people may be hidden by other people, and viewpoint dependency. Generative Adversarial Networks (GANs) have been used to create artificial video clips that mimic real human movements [28]. While GAN-based augmentations can be used to add variability to datasets for action recognition with different motion styles and improve recognition accuracy [29], they require substantially more computational resources and training data to produce samples that are not unrealistic or of poor quality [30]. Several previous HCI studies emphasize the importance of strong action recognition as the basis for natural user interfaces. It has been shown in studies to be crucial for gesture-driven, immersive AR/VR and assistive interaction technologies [31]. In more recent work [32] and [33],

the emphasis is placed on multimodal interaction as being at the heart of HAR designs. These studies highlight the importance of having HAR models that are reliable in the presence of real-world variability, viewpoint changes, and occlusion. Hence, we aim to tackle these gaps directly in our ADGF. We not only focus on the augmentation of the dataset but also on its evaluation on an appropriate 3D CNN model. So here are a few of the state-of-the-art HAR models that we have undergone. In order to extend the learned features to the temporal dimension, an inflated 2D kernel is applied to 2D ConvNets, known as Inflated 3D ConvNets (I3D) [34]. SlowFast Networks capture fine-grained and coarse-grained motion in videos by processing them at varying speeds. Two architectures based on transformers add self-attention mechanisms to model long-range temporal dependency [35]. ResNet 3D models, or residual networks, enable very deep networks to be trained without loss of accuracy while maintaining efficiency. We are using ResNet 18 (a smaller member of the ResNet family) rather than larger versions such as ResNet 50 and ResNet 101. This makes it lighter and faster, and yet it still captures good features in both the spatial and temporal data – important for action recognition. While other approaches mentioned earlier advance feature extraction, they still rely heavily on existing datasets; without augmented or synthetic data, their generalization remains constrained. Our framework can be helpful for such models by supplying them with richer training conditions, by means of addressing their sensitivity to occlusion, viewpoint changes, and motion diversity. Alternatively, hybrid methods with handcrafted features and deep learning-based representations could be more effective. This review focuses on the advances and hurdles in HAR and specifically discusses how 3D CNNs have contributed to the advances made, and how there are still limitations in the datasets that need to be overcome in order to make improvements in real-world performance. Many augmentation methods in traditional methods are not fully capable of capturing the dynamics of human motion because they cannot be applied at the level of the still image. Our work addresses this by providing a framework for generating synthetic datasets with four judiciously chosen augmentations that are not applied directly in the default or fixed setting, but instead, their parameters are tuned through repeated experimental trials to be more similar to real-world samples. Later, its impact is validated on ResNet-18 using the HMDB51 dataset and its different variants. We aim to demonstrate improvements in accuracy, efficiency, and generalization. Beyond the immediate results, this framework can contribute to ongoing efforts in creating HAR systems that can function reliably in

unconstrained environments of real-world applications. Existing augmentation studies for HAR primarily focus on isolated transformations or generic image-based augmentation pipelines. Comparatively fewer studies investigate structured augmentation balancing, parameter-constrained transformations, and robustness-oriented augmentation scheduling specifically for cross-view video-based HAR. The proposed framework addresses this gap through a controlled augmentation integration strategy designed to preserve temporal consistency and action realism.

3. Proposed Augmented Data Generation Framework

The methodology diagram is presented in two parts, as shown in Figs. 1 and 2. Fig. 1 presents the overall block diagram of the proposed method, while Fig. 2 shows the detailed architecture of the Augmented Data Generation Framework (ADGF). We propose a systematic framework that employs four advanced augmentation techniques to improve video-based human action recognition.




Fig. 1. Overall Architecture of the Proposed Framework for Video-Based Human Action Recognition




Fig. 2. Architecture of the Proposed Augmented Data Generation Framework (ADGF) Incorporating Four Advanced Augmentation Techniques

The methodology diagram illustrates the overall workflow of the proposed approach. First Original HMDB51 video clips are pre-processed, which includes frame extraction at 30 fps, resizing to 112×112 , and normalization. The training set is divided into four balanced subsets, each undergoing a distinct augmentation technique. Then we recombine

augmented subsets with unmodified clips to form a diverse training dataset. This dataset is then fed into a 3D ResNet-18 model for training, while we keep the validation and test sets the same. The goal is to

enhance diversity, robustness, accuracy, and generalization in human action recognition. Augmentation techniques are applied carefully, following the distribution ratio as discussed above,

Table 1.

Augmentation Allocation and Expansion Metrics

Sub-Volume	Applied Technique	Allocation %	Baseline Clips	Augmented Output	Dataset Expansion Ratio
D1	Viewpoint Angle Variation	25%	1,353 clips	1,353 clips	2.0x
D2	Elastic Deformation	25%	1,353 clips	1,353 clips	2.0x
D3	Random Erasing	25%	1,353 clips	1,353 clips	2.0x
D4	Occlusion Simulation	25%	1,353 clips	1,353 clips	2.0x
Total Pool	ADGF Framework	100%	5,412 clips	10,824 clips	2.0x (Overall)

Table 2.

Layer-wise Configuration of the Proposed 3D CNN Architecture for Human Action Recognition

Layer Identifier	Target Output Shape ($C \times T \times H \times W$)	Param Metric	Spatial/Temporal Filter / Description
Input Canvas	$3 \times 16 \times 112 \times 112$	—	Raw Video Frame Input Tensor
Conv1 Layer	$64 \times 16 \times 56 \times 56$	9,408	$7 \times 7 \times 7$, Stride 2, Padding 3
MaxPool	$64 \times 16 \times 28 \times 28$	—	$3 \times 3 \times 3$, Stride 2, Padding 1
Conv2_x Block	$64 \times 16 \times 28 \times 28$	$\approx 147K$	$2 \times [3 \times 3 \times 3$ Spatiotemporal Residuals]
Conv3_x Block	$128 \times 8 \times 14 \times 14$	$\approx 1.1M$	$2 \times [3 \times 3 \times 3$ Spatiotemporal Residuals]
Conv4_x Block	$256 \times 4 \times 7 \times 7$	$\approx 7.1M$	$2 \times [3 \times 3 \times 3$ Spatiotemporal Residuals]
Conv5_x Block	$512 \times 2 \times 4 \times 4$	$\approx 22.8M$	$2 \times [3 \times 3 \times 3$ Spatiotemporal Residuals]
Global AvgPool	$512 \times 1 \times 1 \times 1$	—	3D Volume Feature Reduction Pool
FC Layer	51×1	26,163	Linear Target Classification Projection Layer
Total Model Capacity	—	33.2M	5.1M Combined Core Network Neurons

because applying blindly to every video clip can lead to excessive redundancy and unnecessary computational load. Such indiscriminate augmentation is common in prior studies and often leads to unrealistic samples or biased evaluation. The optimal way to apply viewpoint angle, elastic

Deformation, random erasing, and occlusions depends on how augmentation is distributed across the dataset. Therefore, we use a balanced, structured augmentation approach rather than applying all augmentations to every clip. We split original HMDB51 into training set (80%)→ Augmentations applied and test set (20%)→

as no augmentation applied. Augmenting test data can bias evaluation because models may learn augmentation-specific patterns instead of true generalization. The test set must reflect natural, unseen data. We further divide the training set into four equal subsets instead of applying all augmentations to every clip to avoid overlapping. This ensures a balanced dataset where the model learns diverse patterns instead of just one augmentation technique dominating. The subsequent sections describe each stage of the proposed framework in detail, including the augmentation strategy, feature extraction process, network architecture, and training procedure.

3.1 Augmentation Allocation and Expansion Metrics

To maximize spatial diversity without introducing data dominance or overfitting, the raw training partition of the HMDB51 dataset (8% of total samples) is uniformly split into four mutually exclusive sub-volumes, each matching exactly 25% of the training set. Each sub-volume is uniquely subjected to one independent augmentation modality. Following processing, the four enhanced sub-volumes are concatenated back with the original, untouched training pool. This execution profile achieves a strict 1:1 balance between raw and augmented data frames. The precise breakdown of the distribution scheduling is formalized in Table 1.

3.2 Backbone Network Selection

We chose the 18-layer variant because it offers an optimal trade-off between accuracy and computational efficiency, hence resulting in faster convergence time as compared to 34, 50, and 101-layer architectures. Therefore, ResNet-18 is a lightweight architecture and can be trained repeatedly across baseline, augmented, and ablation experiments while achieving competitive accuracy on HMDB51 and UCF101 datasets. The choice of network ensures that any observed performance gains can be attributed to our augmentation strategy rather than to differences in model capacity.

Table 2 shows a detailed summary of the 3D ResNet-18 model architecture used for human action recognition on the HMDB51 dataset in the experimental setup. It outlines architectural details such as layers, output shapes, number of parameters, and the number of neurons at each stage of the network.

3.3 Data Augmentation in Human Action Recognition

Data augmentation plays a crucial role in deep learning and computer vision tasks, particularly for human action recognition in videos. It can generate modified versions of the original dataset to improve model generalization and robustness[38]. Conventional augmentations, such as cropping, flipping, and rotation, have been widely used in image classification but may not be as effective for video-based tasks. Videos contain more information compared to static images, such as temporal dynamics, motion consistency, and spatial dependencies across frames. Therefore, augmentations should preserve action semantics while introducing any variations required to improve the learning capabilities of a model. Advanced augmentation techniques that are employed by our AGDF can enhance generalization, adversarial robustness, and domain adaptation by exposing the model to more diverse, but realistic training samples. Traditional augmentations (like flipping and rotation) may fail to introduce the necessary variability needed for video-based tasks. Instead, advanced augmentations like viewpoint angle transformation, elastic Deformation, random erasing, and occlusions offer significant improvements under variable conditions as discussed in the methodology overview.




Fig. 3. Illustration of Data Augmentation Viewpoint on the Left and Elastic Deformation on the Right

Figs. 3 and 4 visually show samples of each technique generated through ADGF, showing both the original and augmented parts when applied to the original HMDB51 videos.



Fig. 4. Illustration of Data Augmentation Occlusion Simulation on the Left and Random Erasing on the Right

3.3.1 Viewpoint angle augmentation: Human actions can look different when viewed from different angles. For example, a front view may look different from a side view. Viewpoint angle can be classified into three categories: (1) small rotations ($\pm 10^\circ$ to $\pm 20^\circ$), moderate rotations, and large rotations.

$$R_x(\theta_x) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos\theta_x & -\sin\theta_x \\ 0 & \sin\theta_x & \cos\theta_x \end{pmatrix}, \quad (1)$$

$$R_y(\theta_y) = \begin{pmatrix} \cos\theta_y & 0 & \sin\theta_y \\ 0 & 1 & 0 \\ -\sin\theta_y & 0 & \cos\theta_y \end{pmatrix}, \quad (2)$$

$$R_z(\theta_z) = \begin{pmatrix} \cos\theta_z & -\sin\theta_z & 0 \\ \sin\theta_z & \cos\theta_z & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (3)$$

The combined rotation is then:

$$R = R_z(\theta_z)R_y(\theta_y)R_x(\theta_x), \quad \text{and} \\ \text{the transformed coordinate: } \mathbf{x}' = R \mathbf{x}.$$

For each frame, convert the pixel coordinates into a 3D coordinate system, and choose Euler angles. Compute R with the given angles and apply R to every pixel's coordinate to generate a new frame. The output frame will exhibit a change in perspective. Objects will appear rotated as the subject might appear as if viewed from a slightly different angle, as shown in Fig. 3 (left). Then, comparing the original frame with the augmented one will show the change in viewpoint.

3.3.2 Elastic deformation: Elastic deformation simulates nonlinear spatial transformations. Consider

a displacement field generated by convolving a random field with a Gaussian kernel. The deformed coordinate is given by Eq. (4).

where: controls the scaling (magnitude) of the Deformation, is the Gaussian filter with standard deviation, is a random field drawn from a uniform or normal distribution. The deformation scaling parameter is tightly constrained ($\alpha \in [2, 5]$, $\sigma = 4.0$). Setting $\alpha > 5$ violates spatial consistency by generating physically impossible bio-mechanical distortions, whereas $\alpha < 2$ offers insufficient variance to simulate diverse muscular body topologies. Unlike aggressive spatial warping, deformation magnitudes are constrained to maintain physically plausible human motion patterns, as shown in Fig. 4 (right).

3.3.3 Random erasing: Random erasing introduces robustness by occluding random rectangular regions in the frame. It is computed using a mask M, as shown in Eq. (5).

$$M(x, y) = \begin{cases} 0, & \text{if } (x, y) \in \mathcal{R} \\ 1, & \text{otherwise,} \end{cases} \quad (5)$$

with R being a randomly chosen rectangular region with area A_r such that:

$$A_r \in [\lambda_{\min}A, \lambda_{\max}A],$$

where A is the total area of the frame, and A_m and $A_{\max}$ are hyper parameters. The localized area scale parameter matches $\lambda_{\min} = 0.02$ and $\lambda_{\max} = 0.20$. This guarantees that masked patches block structural elements (e.g., hands or feet) to challenge localized spatiotemporal tracking, without crossing the 20% canvas threshold which completely erases the primary action trajectory. The augmented image I' is then :

$$I'(x, y) = I(x, y) \cdot M(x, y) + (1 - M(x, y)) \cdot I_{\text{noise}} \quad (6)$$

where can be constant or a random noise value. Fig. 3 (right) shows a sample obtained using this Eq.

3.3.4 Occlusion simulation: Occlusion augmentation can be modelled similarly to random erasing but is designed to mimic realistic blockages, such as objects or body parts obscuring the view. We define it using an occlusion mask O with a predefined shape such as an ellipse or irregular blob.

In Eq. (7), defines the resulting transformed destination coordinates on the spatial matrix, represents the initial incoming coordinate array, and represents the binary occlusion state map. The matrix represents the replacement pixel matrix containing the structured obstructing artifact mask, which forces the neural network layers to learn distributed feature tracking instead of relying on narrow region templates.

overhead of approximately 37.8 minutes. However, because this augmentation phase is completed entirely offline prior to training, it introduces zero runtime overhead during active deep learning execution. In fact, because the augmented data stabilizes gradient flow, the total training time dropped from 11.4 hours (baseline) to 7.5 hours, saving roughly 34% in active computing resource requirements. The 3D

Table 3. Classification Performance and Training Efficiency on HMDB51 and UCF101

Dataset	Accuracy (%)	Precision (%)	Recall (%)	F1 (%)	Final Loss	Epochs to Converge
Original HMDB51	87.5	86.9	87.2	87.0	0.21	38
Augmented Only	95.3	95.0	95.2	95.1	0.15	27
Combined	97.7	97.6	97.7	97.6	0.12	25
UCF101 (cross-dataset) Baseline	85.1	84.5	84.9	84.7	0.24	N/A
UCF101 (augmented model)	96.5	96.3	96.5	96.4	0.14	N/A

4. Experimental Setup

4.1 Experimental Configuration

Three different models of HMDB51 were used for the experiments: a 3D ResNet 18 trained with the original HMDB51 dataset as baseline, a 3D ResNet 18 trained with a synthetically augmented HMDB51 dataset, and a 3D ResNet 18 trained with a combined dataset consisting of the original HMDB51 and synthetically augmented HMDB51 datasets. Implementations of the framework pipeline are done as follows: PyTorch for 3D ResNet, OpenCV for frame extraction, resizing, transformations, etc. For complex image augmentation, use Albumentations and/or imgaug, both of which offer additional parameters and options for augmentation. FFmpeg – Extract frames from video (30 fps), MoviePy – manipulate and transform video. Python – For flexible programming. Matplotlib and Seaborn – For plotting scientific data. For Model Evaluation and Analysis – performance metrics and classification reports, we have adopted standard techniques of statistical visualization using Scikit-Learn and Matplotlib. All experiments were conducted on an NVIDIA RTX 3090 GPU, CPU AMD Ryzen 9 5900X, 40 GB RAM and 1TB SSD. To fully evaluate the practical viability of the framework, the computational costs introduced during the offline synthesis phase were measured. On an NVIDIA RTX 3090 GPU pipeline, the frame processing, transformation execution, and sequence rebuilding step requires an average processing time of 0.42 seconds per video clip. For the entire training sub-volume, this adds a one-time offline preprocessing

ResNet model that can extract spatiotemporal features has been chosen for training. The following settings were used for the training: 16-frame clips were resized to 112×112 pixels, Adam optimizer with initial learning rate 0.0001, weight decay ($5e-4$), Cross-Entropy Loss, Batch Size: 16, and data split was 80% Train, 10% Val, 10% Test (with early stopping if validation loss did not decrease for 10 consecutive epochs). Other than the top 1 accuracy, we report precision, recall, F1-score, training time per epoch, and the final loss. To evaluate the cross-dataset generalization ability, the model was tested on

UCF101, where pre-processing steps were the same, but not followed by the dataset (zero-shot). To improve reproducibility, all experiments were conducted using fixed random seeds, identical train-validation-test splits, and consistent pre-processing settings across all experimental configurations. Hyperparameters, augmentation settings, and implementation details were kept unchanged across repeated trials unless otherwise specified. All major experiments were repeated across multiple independent runs using different random seeds. Reported results correspond to the mean performance across repeated experiments together with standard deviation values. Statistical significance between baseline and augmented configurations was evaluated using paired t-tests at a significance threshold of $p < 0.01$.

4.2 Statistical Validation and Reproducibility

To improve reproducibility and strengthen the scientific rigor of the study, all experiments were conducted using fixed random seeds, identical train-validation-test splits, and consistent preprocessing settings across all experimental configurations. Hyperparameter settings, augmentation configurations, and implementation details were kept unchanged across repeated trials unless otherwise specified.

Since the proposed framework reports considerable performance improvements, additional validation measures were adopted to reduce the possibility of performance variations caused by random initialization or data sampling effects. All major experiments were therefore repeated across multiple independent runs using different random seeds. Reported results correspond to the mean performance across repeated experiments together with standard deviation values.

Statistical significance between baseline and augmented configurations was evaluated using paired t-tests at a significance threshold of $p < 0.01$. These repeated trials were performed to assess experimental variability, convergence stability, and validation reliability.

Furthermore, robustness was evaluated not only through in-dataset testing on HMDB51, but also through cross-dataset evaluation on UCF101. Consistent improvements across repeated trials and cross-dataset evaluation indicate that the observed gains are less likely to result from overfitting or random experimental variations.

These measures strengthen the reliability and reproducibility of the proposed framework and provide a more statistically grounded evaluation beyond single-run performance reporting.

5. Results and Discussion

5.1 Evaluation on HMDB51 and UCF101 Datasets

Table 3 shows the performance across datasets. On HMDB51 (combined dataset), accuracy increased from 87.5% to 97.7%, the number of epochs required to converge decreased from 38 to 25, and loss decreased from 0.21 to 0.12. Cross-dataset testing on UCF101 showed accuracy increasing from 85.1% to 96.5%. A paired t-test indicated statistically significant improvements ($p < 0.01$) for these improvements. The results show that, when synthetic data are used as an augmentation method, and the neural network is trained to process the synthetic data, the performance and transferability of the HAR are greatly boosted. The model is trained with a variety of

motion patterns, occlusions, and viewpoints to combat overfitting and speed convergence.




Fig. 5. Training and Validation Loss and Accuracy of Baseline and Augmented Models

5.2 Training Dynamics and Classification Analysis

We show the training and validation behaviour of the baseline and augmented models in Fig. 5. The loss curves in Fig. 5 (left) demonstrate that the baseline model takes a long time to converge: it takes 38 epochs to reach a validation loss of 0.25, and the validation loss oscillates around this value.




Fig. 6. Confusion Matrices of the Baseline Model and the Proposed Augmented Model

The confusion matrices of the baseline and augmented models on HMDB51 are given in Fig. 6. The numbers along the diagonals are the accuracy for each class, and off-diagonal numbers are the number of misclassifications.

5.3 Ablation Study on Individual Augmentation Techniques

We have performed an ablation study to determine the individual effect of each augmentation strategy used in ADGF by training 3D ResNet-18 models on each method in isolation. The same training settings, optimizer parameters, and evaluation metrics were used as in the main experiment. This enables the separate effect of the different types of augmentations to be isolated from recognition performance and convergence speed.

The effects of each augmentation on HMDB51 are presented in Table 4. All techniques showed an improvement in accuracy and convergence relative to

the baseline model, with the largest improvements being in the combined augmentation technique and viewpoint angle variation. Elastic deformation was the least effective single augmentation technique, although it still improved the model performance.




Fig. 7. Accuracy Comparison of Individual Augmentation Types

The clear performance advantage shown by viewpoint angle variation (+7.0% in-dataset and +7.2% cross-dataset) stems from the structural nature of standard benchmark repositories. Raw action datasets like HMDB51 are highly sensitive to viewpoint dependencies because they consist of videos filmed from fixed, straight-on camera positions. When a standard network encounters an action filmed from an unlearned angle, its spatial feature mapping layers fail to align properly. By applying a 3D coordinate system transformation matrix within controlled limits, our framework forces the network's convolutional kernels to learn invariant, multi-angle skeletons rather than simple, position-dependent pixel outlines. This explains why perspective modification delivers stronger improvements than localized surface shifts like elastic stretching or region masking.

Fig. 7 illustrates the ablation study comparing the individual impact of each augmentation technique on classification accuracy. All augmentations improved performance relative to the baseline (87.5%), with viewpoint angle variation producing the largest single-technique gain (94.5%). When all four techniques were combined, the highest accuracy was achieved.

Fig. 8 shows a bar chart visualizing the relative accuracy improvements reported in Table 4. The chart highlights the incremental gains from each augmentation and the strong cumulative effect of the combined approach. Table 4 shows the relative accuracy improvement of each augmentation technique compared to the baseline model.

Table 4.

Relative Accuracy Improvement by Augmentation Type Compared to Baseline

Augmentation Technique	Accuracy (%)	Improvement (%)
Elastic Deformation	91.8	+4.3
Random Erasing	92.3	+4.8
Occlusion Simulation	93.0	+5.5
Viewpoint Angle Variation	94.5	+7.0
Combined (all 4)	97.7	+10.2




Fig. 8. Accuracy Improvements Across Augmentation Techniques

Table 5.

Robustness Improvement by Augmentation Type

Augmentation Technique	Cross-Dataset Accuracy (%)	Robustness Improvement (%)
Elastic Deformation	89.2	+4.1
Random Erasing	90.0	+4.9
Occlusion Simulation	90.8	+5.7
Viewpoint Angle Variation	92.3	+7.2
Combined (all 4)	96.5	+11.4

(Baseline cross-dataset accuracy on UCF101 = 85.1% — improvement = technique accuracy – 85.1%.)

5.4 Cross-Dataset Robustness Evaluation

While the previous sections evaluated performance on HMDB51, we also assessed the robustness of the proposed approach by testing models trained on HMDB51 directly on UCF101 without any fine-tuning. Cross-dataset accuracy serves as a proxy for robustness, indicating how well the learned features transfer to unseen data. The observed gains in robustness and cross-dataset generalization are particularly relevant for healthcare applications, where models trained in one environment must reliably operate under unseen conditions without extensive retraining.




Fig. 9. Robustness Analysis per Augmentation Category

Fig. 9 shows the robustness improvement (cross-dataset accuracy on UCF101) for individual augmentation techniques and their combination. Values show the gain over the baseline HMDB51-trained model (85.1%). The combined augmentation produces the largest robustness improvement (+11.4%), indicating strong transferability to unseen data.

Conclusion

In this study, we introduced an Augmented Data Generator Framework (ADGF), a structured framework that applies four complementary augmentation strategies: Elastic Deformation, Random Erasing, Occlusion Simulation, and Viewpoint Angle Variation to enhance Human Action Recognition. These techniques employ validated parameter ranges, not as default conditions. AGDF is fine-tuned through a series of iterative experiments for diversity and perceptual analysis, validating variance, cosine distance, LPIPS and FVD metrics. Validation results and further details of these metrics were used only during the fine-tuning process; further details can be provided in the future in extension of this study. Following that, we used a 3D ResNet-18 as our backbone for the HMDB51 and its synthetic versions. We also performed validation for cross-dataset transfer on UCF101. Our method shows substantial gains in performance, and accuracy increased from 87.5% on the baseline to 97.7% with combined

augmentation. Final training loss decreased from 0.21 to 0.12, and convergence accelerated from 38 to 25 epochs. Cross-dataset accuracy rose from 85.1% to 96.5%, which indicated improved robustness and generalization.

Collectively, these results show that a systematically designed augmentation-driven framework can effectively improve accuracy, efficiency, and, more importantly, robustness when exposed to unseen data with variations, which is closer to real-world HAR implementations.

Therefore, ADGF delivered significant improvements in accuracy, convergence speed, and cross-dataset generalization. The current pipeline applies transformations independently at the frame level and does not explicitly model temporal consistency, which could introduce subtle artifacts for some actions. Future work will explore the integration of adversarial training and handcrafted feature fusion to further enhance robustness and extend its evaluation to multimodal datasets, including depth and skeletal data, and benchmark the framework on larger and more challenging datasets, such as Kinect, for scalability and transferability.

References

- [1] M. Karim, S. Khalid, A. Aleryani, J. Khan, I. Ullah, and Z. Ali, "Human Action Recognition Systems: A Review of the Trends and State-of-the-Art," *IEEE Access*, vol. 12, pp. 36372–36390, 2024, doi: 10.1109/ACCESS.2024.3373199.
- [2] G. Diraco, G. Rescio, A. Caroppo, A. Manni, and A. Leone, "Human Action Recognition in Smart Living Services and Applications: Context Awareness, Data Availability, Personalization, and Privacy," *Sensors*, vol. 23, no. 13, p. 6040, June 2023, doi: 10.3390/s23136040.
- [3] A. Sanchez-Comas, K. Synnes, and J. Hallberg, "Hardware for Recognition of Human Activities: A Review of Smart Home and AAL Related Technologies," *Sensors*, vol. 20, no. 15, p. 4227, July 2020, doi: 10.3390/s20154227.
- [4] L. E. Haddad, M. Hanoune, and A. Ettaoufik, "Computer Vision with Deep Learning for Human Activity Recognition: Features Representation," in *Engineering Applications of Artificial Intelligence*, A. Chakir, J. F. Andry, A. Ullah, R. Bansal, and M. Ghazouani, Eds., in *Synthesis Lectures on Engineering, Science, and Technology*, Cham: Springer Nature Switzerland, 2024, pp.

41–66. doi: 10.1007/978-3-031-50300-9_3.

- [5] Y. Zhang, B. Li, H. Fang, and Q. Meng, “Current Advances on Deep Learning-based Human Action Recognition from Videos: a Survey,” in 2021 20th IEEE International Conference on Machine Learning and Applications (ICMLA), Pasadena, CA, USA: IEEE, Dec. 2021, pp. 304–311. doi: 10.1109/ICMLA52953.2021.00054.
- [6] I. Laptev, “On Space-Time Interest Points,” *Int. J. Comput. Vis.*, vol. 64, no. 2–3, pp. 107–123, Sept. 2005, doi: 10.1007/s11263-005-1838-7.
- [7] C. Shorten and T. M. Khoshgoftaar, “A survey on Image Data Augmentation for Deep Learning,” *J. Big Data*, vol. 6, no. 1, p. 60, July 2019, doi: 10.1186/s40537-019-0197-0.
- [8] H. Wang and C. Schmid, “Action Recognition with Improved Trajectories,” in 2013 IEEE International Conference on Computer Vision (ICCV), Sydney, Australia: IEEE, Dec. 2013, pp. 3551–3558, doi: 10.1109/ICCV.2013.441.
- [9] M. Ahsan Habib et al., “Exploring Progress in Text-to-Image Synthesis: An In-Depth Survey on the Evolution of Generative Adversarial Networks,” *IEEE Access*, vol. 12, pp. 178401–178440, 2024, doi: 10.1109/ACCESS.2024.3435541.
- [10] N. E. Khalifa, M. Loey, and S. Mirjalili, “A comprehensive survey of recent trends in deep learning for digital images augmentation,” *Artif. Intell. Rev.*, vol. 55, no. 3, pp. 2351–2377, Mar. 2022, doi: 10.1007/s10462-021-10066-4.
- [11] T. Kumar, R. Brennan, A. Mileo, and M. Bendeche, “Image Data Augmentation Approaches: A Comprehensive Survey and Future Directions,” *IEEE Access*, vol. 12, pp. 187536–187571, 2024, doi: 10.1109/ACCESS.2024.3470122.
- [12] J. Wang, Y. Chen, S. Hao, X. Peng, and L. Hu, “Deep Learning for Sensor-based Activity Recognition: A Survey,” *Pattern Recognit. Lett.*, vol. 119, pp. 3–11, Mar. 2019, doi: 10.1016/j.patrec.2018.02.010.
- [13] L. Bibbò and M. M. B. R. Vellasco, “Human Activity Recognition (HAR) in Healthcare,” *Appl. Sci.*, vol. 13, no. 24, p. 13009, Dec. 2023, doi: 10.3390/app132413009.
- [14] R. Ullah et al., “Vision-Based Activity Recognition for Unobtrusive Monitoring of the Elderly in Care Settings,” *Technologies*, vol. 13, no. 5, p. 184, May 2025, doi: 10.3390/technologies13050184.
- [15] O. D. Lara and M. A. Labrador, “A Survey on Human Activity Recognition using Wearable Sensors,” *IEEE Commun. Surv. Tutor.*, vol. 15, no. 3, pp. 1192–1209, 2013, doi: 10.1109/SURV.2012.110112.00192.
- [16] A. Herath, M. Harandi, and F. Porikli, “Going Deeper into Action Recognition: A Survey,” *Image Vis. Comput.*, vol. 60, pp. 4–21, Apr. 2017, doi: 10.1016/j.imavis.2017.01.010.
- [17] P. Kumar, S. Chauhan, and L. K. Awasthi, “Human Activity Recognition (HAR) Using Deep Learning: Review, Methodologies, Progress and Future Research Directions,” *Arch. Comput. Methods Eng.*, vol. 31, no. 1, pp. 179–219, Jan. 2024, doi: 10.1007/s11831-023-09986-x.
- [18] C. I. Patel, D. Labana, S. Pandya, K. Modi, H. Ghayvat, and M. Awais, “Histogram of Oriented Gradient-Based Fusion of Features for Human Action Recognition in Action Video Sequences,” *Sensors*, vol. 20, no. 24, p. 7299, Dec. 2020, doi: 10.3390/s20247299.
- [19] S. P. Sahoo, S. Ari, K. Mahapatra, and S. P. Mohanty, “HAR-Depth: A Novel Framework for Human Action Recognition Using Sequential Learning and Depth Estimated History Images,” *IEEE Trans. Emerg. Top. Comput. Intell.*, vol. 5, no. 5, pp. 813–825, Oct. 2021, doi: 10.1109/TETCI.2020.3014367.
- [20] I. D. Mienye and T. G. Swart, “A Comprehensive Review of Deep Learning: Architectures, Recent Advances, and Applications,” *Information*, vol. 15, no. 12, p. 755, Nov. 2024, doi: 10.3390/info15120755.
- [21] A. Deng, T. Yang, and C. Chen, “A Large-scale Study of Spatiotemporal Representation Learning with a New Benchmark on Action Recognition,” in 2023 IEEE/CVF International Conference on Computer Vision (ICCV), Paris, France: IEEE, Oct. 2023, pp. 20462–20474. doi: 10.1109/ICCV51070.2023.01876.
- [22] T. Singh and D. K. Vishwakarma, “Video benchmarks of human action datasets: a review,” *Artif. Intell. Rev.*, vol. 52, no. 2, pp. 1107–1154, Aug. 2019, doi: 10.1007/s10462-018-9651-1.
- [23] L. Yuan, Z. He, Q. Wang, and L. Xu,

- “Advancing Human Motion Recognition with SkeletonCLIP++: Weighted Video Feature Integration and Enhanced Contrastive Sample Discrimination,” *Sensors*, vol. 24, no. 4, p. 1189, Feb. 2024, doi: 10.3390/s24041189.
- [24] E. Goceri, “Medical image data augmentation: techniques, comparisons and interpretations,” *Artif. Intell. Rev.*, vol. 56, no. 11, pp. 12561–12605, Nov. 2023, doi: 10.1007/s10462-023-10453-z.
- [25] N. Akhtar, M. Saddique, K. Asghar, U. I. Bajwa, M. Hussain, and Z. Habib, “Digital Video Tampering Detection and Localization: Review, Representations, Challenges and Algorithm,” *Mathematics*, vol. 10, no. 2, p. 168, Jan. 2022, doi: 10.3390/math10020168.
- [26] T. Salimans et al., “Improved Techniques for Training GANs,” in *Adv. Neural Inf. Process. Syst. (NeurIPS)*, 2016, pp. 2234–2242.
- [27] C. Xu, B. Chai, J. He, X. Zhang, and E. W. M. Ma, “Deep Learning-Based Human Action Recognition: A Survey,” *Pattern Recognit.*, vol. 139, p. 109524, July 2023, doi: 10.1016/j.patcog.2023.109524.
- [28] M. A. Hassan, M. U. G. Khan, R. Iqbal, O. Riaz, A. K. Bashir, and U. Tariq, “Predicting humans future motion trajectories in video streams using generative adversarial network,” *Multimed. Tools Appl.*, vol. 83, no. 5, pp. 15289–15311, Sept. 2021, doi: 10.1007/s11042-021-11457-z.
- [29] M.-Y. Liu, X. Huang, J. Yu, T.-C. Wang, and A. Mallya, “Generative Adversarial Networks for Image and Video Synthesis: Algorithms and Applications,” *Proc. IEEE*, vol. 109, no. 5, pp. 839–862, May 2021, doi: 10.1109/JPROC.2021.3049196.
- [30] L. Perez and J. Wang, “The Effectiveness of Data Augmentation in Image Classification using Deep Learning,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Honolulu, HI, USA: IEEE, July 2017, pp. 1–8, doi: 10.1109/CVPRW.2017.19.
- [31] A. Schäfer, G. Reis, and D. Stricker, “The Gesture Authoring Space: Authoring Customised Hand Gestures for Grasping Virtual Objects in Immersive Virtual Environments,” in *Mensch und Computer 2022*, Darmstadt Germany: ACM, Sept. 2022, pp. 85–95. doi: 10.1145/3543758.3543766.
- [32] M. Ehatisham-Ul-Haq et al., “Robust Human Activity Recognition Using Multimodal Feature-Level Fusion,” *IEEE Access*, vol. 7, pp. 60736–60751, 2019, doi: 10.1109/ACCESS.2019.2913393.
- [33] B. X. B. Yu, Y. Liu, X. Zhang, S. Zhong, and K. C. C. Chan, “MMNet: A Model-Based Multimodal Network for Human Action Recognition in RGB-D Videos,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, no. 3, pp. 3522–3538, Mar. 2023, doi: 10.1109/TPAMI.2022.3177813.
- [34] I. M. De La Jara, C. Rodriguez-Opazo, E. Marrese-Taylor, and F. Bravo-Marquez, “An empirical study of the effect of video encoders on Temporal Video Grounding,” in *2023 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, Paris, France: IEEE, Oct. 2023, pp. 2842–2847. doi: 10.1109/ICCVW60793.2023.00306.
- [35] M. Almushyti and F. W. B. Li, “Recognising Human-Object Interactions Using Attention-based LSTMs,” *Comput. Graph. Vis. Comput. CGVC*, p. 5 pages, 2019, doi: 10.2312/CGVC.20191269.
- [36] Y. Huang, H. Chen, K. Huang, M. Chen, J. Wang, and B. Liu, “Optimization of Space-Time image velocimetry based on deep residual learning,” *Measurement*, vol. 232, p. 114688, June 2024, doi: 10.1016/j.measurement.2024.114688.
- [37] D. Tran, H. Wang, L. Torresani, J. Ray, Y. LeCun, and M. Paluri, “A Closer Look at Spatiotemporal Convolutions for Action Recognition,” in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Salt Lake