

# DEAR: Depth-Enhanced Action Recognition

Sadegh Rahmaniboldaji<sup>1</sup>, Filip Rybansky<sup>2</sup>, Quoc Vuong<sup>2</sup>, Frank Guerin<sup>1</sup>,  
and Andrew Gilbert<sup>1</sup>

<sup>1</sup> University of Surrey, Guildford, UK

<sup>2</sup> University of Newcastle, Newcastle, UK  
s.rahmani@surrey.ac.uk

**Abstract.** Detecting actions in videos, particularly within cluttered scenes, poses significant challenges due to the limitations of 2D frame analysis from a camera perspective. Unlike human vision, which benefits from 3D understanding, recognizing actions in such environments can be difficult. This research introduces a novel approach integrating 3D features and depth maps alongside RGB features to enhance action recognition accuracy. Our method involves processing estimated depth maps through a separate branch from the RGB feature encoder and fusing the features to understand the scene and actions comprehensively. Using the Side4Video framework and VideoMamba, which employ CLIP and VisionMamba for spatial feature extraction, our approach outperformed our implementation of the Side4Video network on the Something-Something V2 dataset. *Our code is available at:* <https://github.com/SadeghRahmaniB/DEAR>

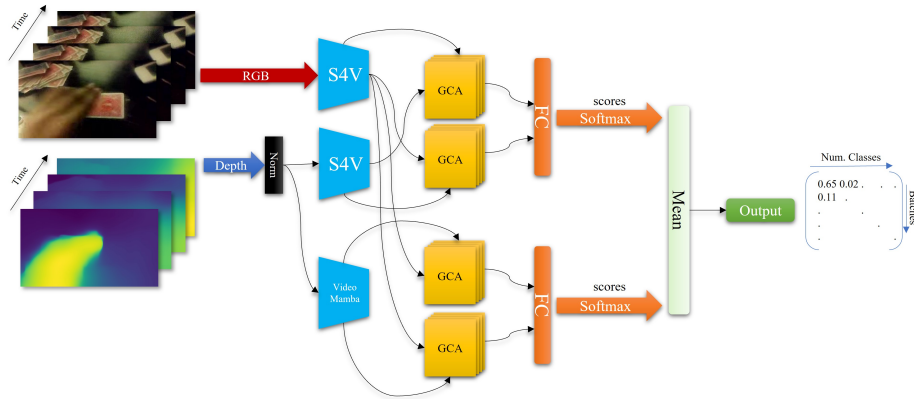
**Keywords:** supervised video understanding · Multi-modal representation learning · depth map · Action recognition

## 1 Introduction

Substantial attention in computer vision has focused on Human Action Recognition (HAR) from videos because of its wide range of applications, including surveillance, human-computer interaction and content-based video retrieval. Despite significant advancements, recognising actions in cluttered scenes remains a challenging problem. This difficulty arises from the inherent limitations in analysing 2D frames captured from a single camera perspective, which often fails to represent real-world environments’ three-dimensional (3D) nature.

Human vision can utilise 3D depth information to efficiently and robustly interpret people’s actions [29, 31]. Such information can be recovered from a combination of depth cues, such as optic flow and pictorial cues estimated from 2D frames [4, 32]. Significantly, depth information can enhance perception of critical 3D spatial arrangements of people and objects [8, 11], segment actions and objects from background clutter [23], and complement other visual cues to differentiate actions under challenging viewing conditions [5, 13, 15].

Using 3D features, such as depth maps, in HAR systems may thus provide critical information about the spatial arrangement of objects within a scene,



**Fig. 1:** Overview of our method. We perform supervised action recognition using a fusion of RGB and depth map frames. We used S4V [33] network for processing RGB and the same network plus VideoMamba [18] to extract depth features. Gated cross-attention (GCA) has been used for fusing modalities, while mean operation has been selected for fusing models’ scores.

offering a more holistic understanding that strictly 2D approaches may miss [17, 19, 21]. Previous studies used different methods to incorporate depth maps as additional input to enhance HAR. These methods include combining depth maps with RGB data in CNN-based networks [10, 14, 30], with inertial sensors such as accelerometers and gyroscopes [1], and with human posture data in CNN-based networks [16].

This study introduces a novel framework that integrates depth maps with RGB frames to enhance action recognition performance. Uniquely, we estimate depth maps from 2D frames instead of using depth cameras, making the approach more affordable and comparable to 2D video benchmarks. Our method uses the Side4Video (S4V) [33] and VideoMamba [18] frameworks, which employ CLIP [25] and VisionMamba [34] for extracting spatial features from RGB frames. We extended this by adding a separate branch for processing the estimated depth maps, followed by a late fusion of the extracted features. This approach captures actions’ spatial and temporal dynamics more effectively by combining information from RGB and depth modalities.

## 2 Enhanced Action Recognition with Depth and RGB

We introduce a framework for taking depth maps and RGB and analyzing their spatial and temporal features to recognize actions. Based on Figure 1, our network’s architecture consists of two major branches, one for processing RGB frames and another for depth maps. In this network, we use a late fusion manner to integrate extracted features using gated cross-attentions and a mean opera-

tion on streams’ scores. In the next section, we describe the network architecture in more detail.

**RGB Encoder** The network in this branch aims to extract both the temporal and spatial features of RGB frames. To achieve this, we employed the Side4Video (S4V) model [33], which uses the vision component of OpenCLIP [25] as the frame feature extractor. Side4Video uses a side network, where the trainable layers are added in parallel with the larger frozen model layers instead of in series. Specifically, in our model inspired by S4V [33], instead of adding new layers in series with the frozen CLIP [25] layers, we introduce side branches for the new layers. This approach ensures that backpropagation occurs only through the new layers, bypassing the frozen layers of CLIP. Consequently, this innovation allows for faster network training with more affordable resources.

**Depth Encoder** Although there are some depth maps datasets like [27], we aim to use standard RGB video benchmarks. Therefore, we generate the depth directly from the 2D frames, using Midas [26] due to its balance between estimated depth maps’ quality and generation speed. Since depth maps are in the standard RGB format to extract the depth map features, we employ an additional Side4Video (S4V) branch to encode the depth data. In addition, we propose a parallel encoder based on the mamba architecture. The incorporation of two sub-branches for depth processing in our network is motivated by the comparative strengths and unique perspectives of the VideoMamba [18] and Side4Video (S4V) [33] models. Their respective benchmark results (see Table 1) show that VideoMamba is not as robust as S4V regarding overall performance. However, it offers a distinct understanding of the data. Since depth map features tend to be more sparse than RGB features, a less powerful but differently perceptive module, such as VideoMamba, is well-suited for this task. It is an architecture focused on sequence modelling to address sequence limitation and hardware consumption issues of transformers. We incorporate VideoMamba [18] as an additional encoder processing the depth, resulting in our network comprising two depth encoder sub-branches. It uses VisionMamba [34], to extract features of input images using Mamba, to process video frames to handle long-term high-quality videos that existing CNN and transformer-based networks might struggle with.

**Modality Fusion** Following encoding the three modalities, we implement a late fusion policy to combine the data to contextualize features bidirectionally. This approach involves processing RGB and depth features through two sets of gated cross-attention (GCA) layers to contextualize each input based on another. It incorporates a gating mechanism that selectively filters and prioritizes the information passed between modalities. This gate acts like a finely tuned sieve, ensuring that only the most relevant data influences the model’s decision-making process. For the two layers, one uses the RGB features as the key and value, the depth as the query, and the key and value for the depth input and the RGB as the query. This process is repeated twice: once for fusing S4V depth maps and RGB features and once for fusing Mamba depth maps and S4V RGB features. A mean operation is applied on two model streams to combine their scores.

**Table 1:** Out approach compared to other methods on SSV2. S4V\* is our implementation

Method	Modality	Backbone	pre-train	Top-1
SlowFast [9]	RGB	ResNet101	K400	63.1
TSM [20]	RGB	ResNet50	K400	63.4
TimeSformer [3]	RGB	ViT-L	IN-21K	62.4
Mformer [24]	RGB	ViT-L	IN-21K+K400	68.1
ViViT FE [2]	RGB	ViT-L	IN-21k+K400	65.9
VIMPAC [28]	RGB	ViT-L	HowTo100M	68.1
VideoMamba-S [18]	RGB	-	IN-1K	66.6
S4V* [33]	RGB	ViT-B	Clip-400M	70.2
Proposed	RGB+Depth(VideoMamba)	ViT-B	Clip-400M	70.0
Proposed	RGB+Depth	ViT-B	Clip-400M	70.3
<b>Proposed</b>	RGB+Depth+ Depth(VideoMamba)	ViT-B	Clip-400M	<b>70.8</b>

The overall model is trained with a cross-entropy loss function:

$$\mathcal{L} = - \sum_{i=1}^{Samples} \sum_{c=1}^{Classes} y_{i,c} \log(\hat{y}_{i,c}) \quad (1)$$

where  $y_{i,c}$  and  $\hat{y}_{i,c}$  are true label and predicted probability for the  $i$ -th sample and  $c$ -th class

### 3 Experiments

To assess the performance of our approach, we used Something-Something V2 [12] (SSV2) to evaluate the action recognition downstream task.

**Implementation Details.** To assess the performance of our approach, we used an input video clip with 8 frames each of size  $224 \times 224$ , and frames were chosen using static stride. The encoder that we are using is the ViT-B/16 architecture [7], for which we used AdamW optimizer. [22]

**Results** Table 1 presents the overall results and comparison with other methods. Our framework outperforms recent works in top-1 accuracy. Our full model **RGB+Depth+Depth(VideoMamba)** achieves 70.8% accuracy; this is 4.2% and 0.6% greater than the individual data models VideoMamba-S and S4V that our approach is inspired by. The table also shows the fusion of depth and RGB, **RGB+Depth**, is 0.5% lower than when we fuse the mamba logits too.

We also examined the impact of depth maps on various classes and found significant improvements in specific activities. For example, in class 145, *Stuffing something into something*, the depth map helped it go up from 63% to 73%. Also, classes 12 and 98, *Dropping something onto something* and *Pushing something so that it almost falls off but doesn't*, achieved a 5% improvement each by introducing depth maps. However, some classes, like 122, *Rolling something on*

a *flat surface*, experienced a reduction of accuracy by 8%, likely due to the lack of depth needed by this class for classification.

## 4 Conclusion

In this work, we introduced a novel approach and architecture for enhancing action recognition by integrating depth maps. Our method estimates the depth maps of RGB frames, processes them through a separate branch from the RGB feature extractor, and subsequently fuses these features to understand the scene and actions comprehensively. We employed the Side4Video framework, leveraging the vision component of CLIP [25] to extract spatial features. Our approach demonstrated superior performance compared to our implementation of the Side4Video network. We benchmarked our method on the Something-Something V2 dataset, which provided a rigorous evaluation environment. For future work, we plan to apply this network to the EPIC-Kitchens dataset [6] to further validate its effectiveness across diverse and complex scenarios. This study underscores the potential of using depth maps alongside RGB features to improve action recognition, highlighting the importance of multi-modal data fusion in developing robust models for understanding intricate human-object interactions.

## Acknowledgements

Leverhulme Trust Research Project Grant RPG-2023-079 funded this work.

## References

1. Ahmad, Z., Khan, N.: Cnn based multistage gated average fusion (mgaf) for human action recognition using depth and inertial sensors (2020)
2. Arnab, A., Dehghani, M., Heigold, G., Sun, C., Lučić, M., Schmid, C.: Vivit: A video vision transformer. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 6836–6846 (2021)
3. Bertasius, G., Wang, H., Torresani, L.: Is space-time attention all you need for video understanding? In: ICML. vol. 2, p. 4 (2021)
4. Brenner, E., Smeets, J.B.J.: Depth Perception, pp. 1–30. John Wiley & Sons, Ltd
5. Chainay, H., Humphreys, G.W.: The real-object advantage in agnosia: Evidence for a role of surface and depth information in object recognition. *Cognitive Neuropsychology* **18**(2), 175–191 (2001), PMID: 20945210
6. Damen, D., Doughty, H., Farinella, G.M., Furnari, A., Kazakos, E., Ma, J., Moltisanti, D., Munro, J., Perrett, T., Price, W., et al.: Rescaling egocentric vision: Collection, pipeline and challenges for epic-kitchens-100. *International Journal of Computer Vision* pp. 1–23 (2022)
7. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)

8. Fairchild, G.T., Holler, D.E., Fabbri, S., Gomez, M.A., Walsh-Snow, J.C.: Naturalistic object representations depend on distance and size cues. *bioRxiv* (2024)
9. Feichtenhofer, C., Fan, H., Malik, J., He, K.: Slowfast networks for video recognition. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 6202–6211 (2019)
10. Garcia, N.C., Morerio, P., Murino, V.: Modality distillation with multiple stream networks for action recognition. In: *Proceedings of the European Conference on Computer Vision (ECCV)* (September 2018)
11. Gomez, M.A., Skiba, R.M., Snow, J.C.: Graspable objects grab attention more than images do. *Psychological Science* **29**(2), 206–218 (2018), PMID: 29215960
12. Goyal, R., Ebrahimi Kahou, S., Michalski, V., Materzynska, J., Westphal, S., Kim, H., Haenel, V., Fruend, I., Yianilos, P., Mueller-Freitag, M., et al.: The "something something" video database for learning and evaluating visual common sense. In: *Proceedings of the IEEE international conference on computer vision*. pp. 5842–5850 (2017)
13. Humphrey, G.K., Jolicoeur, P.: An examination of the effects of axis foreshortening, monocular depth cues, and visual field on object identification. *The Quarterly Journal of Experimental Psychology Section A* **46**(1), 137–159 (1993)
14. Imran, J., Kumar, P.: Human action recognition using rgb-d sensor and deep convolutional neural networks. In: *2016 International Conference on Advances in Computing, Communications and Informatics (ICACCI)*. pp. 144–148 (2016)
15. Jiang, Y., Boehler, C.N., Nönnig, N., Düzel, E., Hopf, J.M., Heinze, H.J., Schoenfeld, M.A.: Binding 3-D Object Perception in the Human Visual Cortex. *Journal of Cognitive Neuroscience* **20**(4), 553–562 (04 2008)
16. Kamel, A., Sheng, B., Yang, P., Li, P., Shen, R., Feng, D.D.: Deep convolutional neural networks for human action recognition using depth maps and postures. *IEEE Transactions on Systems, Man, and Cybernetics: Systems* **49**(9), 1806–1819 (2019)
17. Kondratyuk, D., Yuan, L., Li, Y., Zhang, L., Tan, M., Brown, M., Gong, B.: Movinets: Mobile video networks for efficient video recognition (2021)
18. Li, K., Li, X., Wang, Y., He, Y., Wang, Y., Wang, L., Qiao, Y.: Videomamba: State space model for efficient video understanding (2024)
19. Li, Y., Wu, C.Y., Fan, H., Mangalam, K., Xiong, B., Malik, J., Feichtenhofer, C.: Mvitv2: Improved multiscale vision transformers for classification and detection (2022)
20. Lin, J., Gan, C., Han, S.: Tsm: Temporal shift module for efficient video understanding. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 7083–7093 (2019)
21. Liu, M., Chen, X., Zhang, Y., Li, Y., Rehg, J.M.: Attention distillation for learning video representations (2020)
22. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101* (2017)
23. Nakayama, K., Shimojo, S., Silverman, G.H.: Stereoscopic depth: Its relation to image segmentation, grouping, and the recognition of occluded objects. *Perception* **18**(1), 55–68 (1989), PMID: 2771595
24. Patrick, M., Campbell, D., Asano, Y., Misra, I., Metze, F., Feichtenhofer, C., Vedaldi, A., Henriques, J.F.: Keeping your eye on the ball: Trajectory attention in video transformers. *Advances in neural information processing systems* **34**, 12493–12506 (2021)

25. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision (2021)
26. Ranftl, R., Lasinger, K., Hafner, D., Schindler, K., Koltun, V.: Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer (2020)
27. Shahroudy, A., Liu, J., Ng, T.T., Wang, G.: Ntu rgb+d: A large scale dataset for 3d human activity analysis (2016)
28. Tan, H., Lei, J., Wolf, T., Bansal, M.: Vimpac: Video pre-training via masked token prediction and contrastive learning. arXiv preprint arXiv:2106.11250 (2021)
29. Vanrie, J., Verfaillie, K.: Perceiving depth in point-light actions. *Perception & Psychophysics* **68**(4), 601–612 (May 2006)
30. Wang, H., Song, Z., Li, W., Wang, P.: A hybrid network for large-scale action recognition from rgb and depth modalities. *Sensors* **20**(11) (2020)
31. Wang, Y., Jiang, Y.: Integration of 3d structure from disparity into biological motion perception independent of depth awareness. *PLOS ONE* **9**(2), e89238 (2014)
32. Welchman, A.E.: The human brain in depth: How we see in 3d. *Annual Review of Vision Science* **2**, 345–376 (2016)
33. Yao, H., Wu, W., Li, Z.: Side4video: Spatial-temporal side network for memory-efficient image-to-video transfer learning (2023)
34. Zhu, L., Liao, B., Zhang, Q., Wang, X., Liu, W., Wang, X.: Vision mamba: Efficient visual representation learning with bidirectional state space model (2024)