



Fusing Hand and Body Skeletons for Human Action Recognition in Assembly

Dustin Aganian^(✉) , Mona Köhler, Benedict Stephan, Markus Eisenbach,
and Horst-Michael Gross

Ilmenau University of Technology, Neuroinformatics and Cognitive Robotics Lab,
98693 Ilmenau, Germany
dustin.aganian@tu-ilmenau.de

Abstract. As collaborative robots (cobots) continue to gain popularity in industrial manufacturing, effective human-robot collaboration becomes crucial. Cobots should be able to recognize human actions to assist with assembly tasks and act autonomously. To achieve this, skeleton-based approaches are often used due to their ability to generalize across various people and environments. Although body skeleton approaches are widely used for action recognition, they may not be accurate enough for assembly actions where the worker's fingers and hands play a significant role. To address this limitation, we propose a method in which less detailed body skeletons are combined with highly detailed hand skeletons. We investigate CNNs and transformers, the latter of which are particularly adept at extracting and combining important information from both skeleton types using attention. This paper demonstrates the effectiveness of our proposed approach in enhancing action recognition in assembly scenarios.

Keywords: Action Recognition · Skeleton-based · Fusion · Body Skeletons · Hand Skeletons · 3D/2D Skeletons · Assembly · Deep Learning

1 Introduction

Collaborative robots are playing an increasingly important role in the course of Industry 4.0 [9]. In order for the robot to collaborate with a human worker and assist in assembly processes, it first needs to visually perceive its environment, the current assembly state, and human actions [6, 15, 18]. For human action recognition, often RGB-based approaches are utilized in the state of the art, as they achieve the best results. However, RGB-based approaches face major difficulties when the target scenario deviates from the training scenario. They tend

This work has received funding from the Carl-Zeiss-Stiftung as part of the project engineering for smart manufacturing (E4SM).

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2023
L. Iliadis et al. (Eds.): ICANN 2023, LNCS 14254, pp. 207–219, 2023.
https://doi.org/10.1007/978-3-031-44207-0_18

to overfit to the environment and the persons seen, especially when the training dataset lacks diversity [15]. This limitation frequently applies to assembly datasets [2,3], which are often small and recorded at only a few locations. In contrast, skeleton-based approaches do not face these limitations, as they only process skeletons and, thus, can generalize much better to different environments.




Fig. 1. We combine body skeletons with hand skeletons for human action recognition. Some actions can be recognized primarily by the movement of the hands. The encoding of the skeleton sequences to images is explained in Sect. 4.1. Example frames from [2].

However, as shown in Fig. 1, some actions are difficult to recognize by the body skeleton alone. For example, the action of attaching a small object is mainly characterized by the object movement, as utilized in [1], and how the worker’s hands interact with it. For this assembly step, it is therefore also useful to utilize finer hand skeletons. This is already done for other assembly datasets such as Meccano [12] or Assembly101 [14], which are recorded in first-person view. However, using hand skeletons alone might not be sufficient for actions such as turning, rotating or pushing of workpieces. During these actions, the fingers are mostly rigid, and most of the movement takes place in the upper body.

Therefore, in this paper, we want to investigate how highly detailed hand skeletons can be combined with less detailed body skeletons to enhance the recognition of assembly actions on the ATTACH [2] and IKEA ASM [3] datasets. By doing so, we aim to recognize both types of actions.

Our study examines both 2D and 3D body skeletons. While 3D skeletons offer a more comprehensive representation of the person’s actions, 2D skeletons are more widely available in practical applications. In this paper, we demonstrate how 2D and 3D hand skeletons can be integrated with various body skeletons. One of the key challenges is that the hands are often occluded, either partially or entirely, which can complicate the estimation of hand positions and the fusion with body skeletons. We also explore the challenges associated with differently detailed skeletons. Specifically, a body skeleton typically has 18–32 joints, while two hand skeletons have most often 42 joints. Although, there are typically

more hand joints than body joints, the latter contains significantly more crucial information for many assembly actions. Therefore, in this paper, we describe how to address this dimension imbalance. Our contributions are as follows:

1. We investigate the use of hands in conjunction with body skeletons in both 2D and 3D to improve action recognition for assembly tasks.
2. We predict hand skeletons on the ATTACH and the IKEA ASM datasets and employ a selection process to identify the appropriate hands.
3. To the best of our knowledge, we are the first to employ the SwinV2 transformer [10] for skeleton-based action recognition.

2 Related Work

In the following, we first present the state of the art of action recognition with skeleton sequences, before going into more detail about differences between hand and body skeleton action recognition and possibilities of fusing skeletons.

2.1 Methods for Skeleton-Based Action Recognition

Human action recognition encompasses various subfields, but in this paper, we focus on the action classification task of pre-trimmed video clips of human skeleton sequences, as this task serves as a foundation for other related problems, such as action segmentation or action detection. For skeleton-based action recognition, recently, 2D convolutional neural networks (CNNs) such as VA-CNN [20], 3D CNNs like PoseConv3D [5], graph convolution networks (GCNs) , and transformers like AcT [11] have been used.

In our paper, we adopt the skeleton encoding of [4] and use it like VA-CNN, which employs a ResNet50 backbone, as it has demonstrated superior or comparable results on the ATTACH Dataset [2] compared to GCN methods. In this approach, the skeleton sequence is encoded as an image so that typical image based classifiers can be used. The image encoding also provides the ability to weight the different skeletons based on their occupied image space which will be explained in Sect. 4.

Moreover, we are able to replace the CNN backbone with the SwinV2-T transformer [10], which has demonstrated excellent results in image-based pattern recognition.

2.2 Hand and Body Skeleton-Based Action Recognition

The idea of fusing less detailed body skeletons with highly detailed hand skeletons for action recognition has only been briefly addressed in the literature, and is still a new area of research. For instance, in NTU-X [17] body skeletons from NTU-RGBD 60/120 were extended to include highly detailed hand skeletons and facial features. In [16] a model was trained for every skeleton type to build an ensemble for classifying actions. It was demonstrated in [16,17] that additional hand skeletons from the NTU dataset for everyday and domestic actions (such as eating or blowing one’s nose) are helpful to the classification task.

In contrast, during assembly, the hand is often occluded by the object being worked on, and the quality of the estimated hand skeletons varies significantly. Typically, the state of the art for action recognition with hands focuses on gesture recognition, where the hands are usually unoccluded. For action recognition during assembly, hand skeletons have only been used in fine motor assembly (e.g., Meccano [12], Assembly101 [14]), where cameras are mounted either on the worker’s head or above the table and focus on the worker’s arms and hands. For instance, in the application scenario of fine-motor toy assembly, which is similar to ours, [14] demonstrated that estimated hand skeletons can be utilized for action recognition. This indicates that our approach of fusing body skeletons with hand skeletons shows promise for action recognition in general assembly tasks. Such tasks involve a combination of coarse actions, where the movement of the body is relevant, and fine motor actions (as in Fig. 1), where hand skeletons are primarily important. Therefore, in this paper, our goal is to explore how these differently detailed body and hand skeletons can be combined optimally.

3 Hand and Body Skeleton Dataset Preparation

Below, we first present the datasets we used. Afterwards, we explain how we estimated the hand skeletons and what to consider when processing them.

3.1 Datasets

To show our approach, we utilize two datasets that contain both small-grained assembly actions that can be mainly recognized by the hands movement as well as coarse assembly actions that involve the whole body, namely the ATTACH [2] and the IKEA ASM [3] datasets. Both datasets are captured from multiple views (three) and consist of assembly actions, where IKEA furniture are assembled. The action names for the action recognition task are composed of verb-object pairs. Below, we shortly discuss each dataset characteristics in detail.

ATTACH. The ATTACH dataset [2] provides different training splits, we focus on the person split in this paper as it is the most commonly used split for action recognition. Skeleton data are available in 3D from the Azure Kinect framework. Since the state of the art typically deals with 2D skeletons, we have also transformed the 3D skeletons into the 2D frame of the RGB camera. In our experiments, we consider both 3D and 2D body skeletons for combination with hand skeletons.

It is worth noting that actions are labeled for each hand independently. Moreover, some actions involve the use of tools such as wrenches, hammers or screwdrivers, where most of the movement occurs in the hand and fingers. Intuitively, this suggests that incorporating additional hand skeletons could potentially enhance the performance of skeleton-based action recognition methods.

IKEA ASM. We use the official splits provided in [3]. The dataset provides 2D skeletons for all views estimated by Keypoint R-CNN [8]. Unlike the Kinect

skeleton, Keypoint R-CNN only predict one single wrist joint per hand. Therefore, incorporating additional hand skeletons might also be useful for action recognition on the IKEA ASM dataset.

However, it should be noted that some actions are difficult to recognize even with hand skeletons. For example, actions such as *pick up back panel*, *pick up front panel*, and *pick up side panel* can only be distinguished by the object used [1], which is not present in the skeleton data.

3.2 Hand Skeleton Estimation

For estimating hand skeletons, the hands need to be clearly visible in the current frame. However, due to their small size in the IKEA ASM dataset, we first cropped a 300×300 patch of the RGB image around the wrist joint of the body skeleton. For the ATTACH dataset, we can skip this first step.




Fig. 2. Overview of our different fusion approaches. (a) As a baseline we train models with only the body skeleton. H is the height of the input image. (b) As a simple way of fusing both skeleton types we merge them into a single image while investigating different ratios between body and hand skeletons. N_h is the number of hand joints (42 in our case) and s is a scaling factor. (c) We treat both skeletons types as different modalities and apply them as distinct input images.

To detect hands and estimate hand skeletons, we used MediaPipe [19]. However, since the predictions can be rather noisy, we filtered the hands by discarding all hands where the distance between the wrist joints of the predicted hand skeleton and the body skeleton exceeded a certain threshold. We kept at most two hands per image. In cases where hand skeletons were missing, we simply took skeletons from past frames to attribute for the missing data.

MediaPipe predicts both 2D and 3D hand skeletons with 21 joints each. While 2D hand skeletons are represented in the image plane, the 3D hand skeletons are represented in a metric space, where the origin is located on the surface of each hand. Therefore, when working with 3D data, we transformed the 3D hands into the frame of the 3D body skeletons.

4 Approach

In the following, we describe our approach to action recognition of pre-trimmed skeleton sequences. In Sect. 4.1, we first present our baseline with body skeletons, before discussing different variations for incorporating hand skeletons in Sect. 4.2.

4.1 Baseline: Body Skeleton Approach

For our baseline, we only use the body skeleton, without incorporating additional hand skeletons¹. For this, we encode the skeletons from a trimmed action sequence into one single RGB image, similar to [4, 20]. One column of the image represents one frame, where the skeleton joints are stacked in a fixed order. To transform a joint to RGB, the XYZ coordinates are normalized and scaled.² For 2D skeleton data, we have just two channels. These images (see Fig. 1 for a visualization) can then be used as input to typical image-based classification architectures such as ResNet50 (ResNet).

Furthermore, while ResNet is typically used in the state of the art [2, 20], we additionally use a SwinV2-T transformer (Swin) [10] for the first time to classify skeleton sequences. Moreover, Swin offers another possibility for fusing hand and body skeleton data, which we will describe in the following.

4.2 Approaches for Fusing Hand and Body Skeletons

For incorporating additional hand skeleton data, we experiment with different methods, as illustrated in Fig. 2. Figure 2a serves as a schematic representation of our baseline approach. In the following, we describe two approaches for encoding the sequence of body skeletons with additional hand skeletons. The first approach involves encoding the hand and body skeletons in a single image, while the second approach creates multiple images that are then combined in the network, similar to multimodal networks that integrate color data with depth data [7, 13].

Single Image Fusion. Figure 2b illustrates our single image fusion approach. Naively, the hand skeleton joints could be appended below the body skeleton joints in the skeleton encoded image. For example, for a Kinect Azure skeleton and MediaPipe hands, the first 32 rows would contain the body skeleton, followed by the right hand and the left hand, each with 21 rows. In this way, however, the body skeleton would account for just under 43% of the input, while the hands would account for the remaining 57%. Such a division, in which the number of hand joints of both hands is predominant, is typical for the relevant skeletons used in the state of the art. This example is shown on the left side of Fig. 2b.

However, for recognizing assembly actions, the body skeleton provides more relevant information than the fine hand skeletons, which should only serve as support. With such a naive partitioning of the image, the classifier is given a bias by devoting a larger input space to the hand skeletons.

To address this issue, we investigate another option to fuse the skeletons into one image, which is shown on the right side of Fig. 2b. Here, we keep the original scaling resolution of the body skeleton as in the baseline (see Fig. 2a).

¹ Our preliminary experiments on the ATTACH dataset using hand skeletons solely showed far inferior results compared to body skeletons solely and are thus not investigated further.

² For ResNet, the image is resized to 224×224 with pixel values ranging from 0-255. For Swin, we use a resolution of 256×256 with pixel values from 0-1.

The body skeleton is scaled up to the original input resolution of the classifier, and subsequently, another image with upscaled hand skeletons is stacked below. We investigate scaling factors $s \in [1, 8]$, where we scale the height of the encoded hand skeleton images (i.e., $N_h = 42$ for MediaPipe skeletons), where $s=8$ resembles the scaling of the body skeleton image.

Multiple Image Input. As an alternative to the previous approach, the skeleton data can be split into different images, and the resulting features can be fused in the network. Recent work on the EMSAFormer [7] has shown that the SwinV2 transformer is particularly suitable for multimodal processing. In their study, the Swin transformer was extended in such a way that RGB and depth images of a scene are fed into the same Swin network as two different images.

We propose a similar approach for processing the encoded body skeleton images and the encoded hand skeleton images. Figure 2c (left) shows how we create two images, one for the body skeleton and one for both hands. The first image is encoded on the first 64 channels of the feature map in the patch embedding, and the second image is encoded on the last 32 channels. After the first attention block, the network combines the information and passes it on to the subsequent blocks, whereby the Swin architecture was not changed.

Alternatively, we can split the skeletons into three images, as shown in Fig. 2c (right). In this case, three images are created, and each is embedded on 32 channels and given to the respective attention head. With this approach, the network itself can decide how to further use the combined information.

5 Experiments

Below we present the results of our experiments on fusing body skeletons with hand skeletons. In Sect. 5.2, we show experiments with 3D body skeletons, before moving on to 2D body skeletons in Sect. 5.3. First, we describe our training setup.

5.1 Setup

Our networks were trained for 100 epochs using the Adam optimizer and a one cycle learning rate scheduler with 10% of epochs as warmup and several maximum learning rates ranging from $5 \cdot 10^{-3}$ to $5 \cdot 10^{-5}$. We validated after each epoch and chose the best epoch for testing. The performances of our trained networks are evaluated using mean class accuracy (mAcc) and top-1 accuracy (top1), two widely used metrics in action recognition literature [2, 12, 14, 17].

Our networks are initialized with ImageNet weights, which improves performance, although the encoded images generated from skeleton data differ a lot from real images. However, performance is still fluctuating, which is why we trained with at least five well-functioning learning rates and repeated training three times for each setup. We present our result using box plots, where each box plots summarizes at least 15 trainings.

5.2 Experiments with 3D Body Skeletons

In the following, we present results solely on the ATTACH dataset [2]. While the IKEA ASM dataset [3] also includes 3D skeletons, they are only available for one camera perspective and captured at a very low frame rate, which makes them rather unsuitable for skeleton-based action recognition.

Baseline – 3D Body Skeletons This subsection serves as a benchmark for our experiments with fused inputs, as we optimize hyperparameters to create a strong baseline using body skeletons solely. On the left side of Fig. 3, we present the results of the baseline experiments with 3D body skeletons on the ATTACH dataset. We compare the performance of two models with similar complexity, namely the SwinV2-T transformer (Swin) and the ResNet50 CNN (ResNet). Our results demonstrate that Swin outperforms ResNet, with a median improvement of more than six percentage points and a maximum improvement of more than four percentage points. Even the worst performing Swin model performs better than the best ResNet model, indicating that Swin is a suitable model for processing skeleton sequences encoded as images.

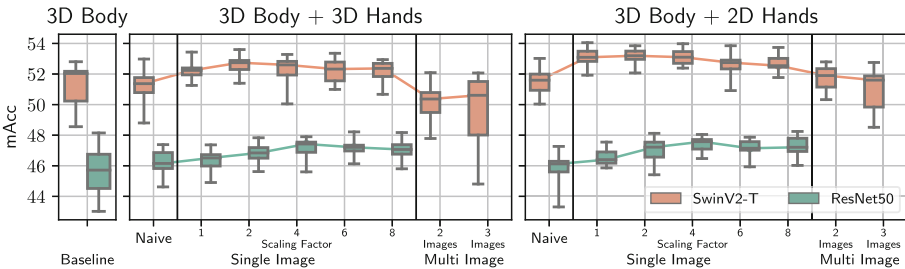


Fig. 3. Results using 3D body skeletons on the ATTACH dataset for our baseline models as in Fig. 2a and our different fusion methods: Naive concatenation as in Fig. 2b left, image concatenation with scaling of encoded hand skeleton image as in Fig. 2b right, multi image input as in Fig. 2c. Best results are listed in Table 1.

However, we want to emphasize that training with Swin is significantly more challenging than with ResNet, which is usually very robust regarding hyperparameters. With Swin, it is crucial to select an appropriate learning rate schedule, as training can fail with even slightly too high learning rates. Conversely, slightly too low learning rates do not produce significant improvements over ResNet. We found that the best results were achieved with learning rates only marginally smaller than the ones that caused training to fail.

Fusion of Hand Skeletons with 3D Body Skeletons Figure 3 illustrates the results of our fusion experiments using our transformed 3D hand skeletons from MediaPipe in the middle, and 2D hand skeletons on the right.

3D Hands: Overall, an improvement of the median and variance can be observed when using the single image fusion approach with the correct scaling factor. While no improvement of the maximum for ResNet is observable, for Swin the incorporation of 3D hands increased performance by about one percentage point. This shows that there is relevant information in the hand skeletons that helps making the training more consistent or even improves the general quality of the models. Moreover, it shows that Swin is significantly better at combining the relevant information from the estimated hand skeletons with the full body skeletons. However, since the 3D hand skeletons in MediaPipe are estimated on 2D color images, a poor estimation of the hand joints may have led to only slight improvements. Therefore, we explore to combine the 3D body skeleton with 2D hand skeletons in the following.

2D Hands: The results of fusing 2D hand skeletons with 3D body skeletons are shown in the right half of Fig. 3. First and foremost, this fusion can be challenging due to the different frames of reference. The 3D skeletons exist in a metric space while the 2D skeletons are given in image coordinates. This means that the different parts of the input image for the single image fusion approach need to be normalized independently.

For ResNet, using the 2D hands results in similar performance compared to 3D hand skeletons. On the other hand, Swin demonstrates that this fusion works very well, and in some cases, it performs even better than the fusion with 3D hands. In fact, the maximum improvement over the baseline is more than one percentage point. This highlights Swin’s ability to handle the challenges of using disparate input spaces.

These results also confirm our assumption that the estimated 3D hand skeletons from MediaPipe are less accurate than the 2D hand skeletons.

Fusion Variations: When comparing the different fusion approaches that we examined, both ResNet and Swin yielded similar results. The naive approach, which involves stacking the hand and body skeleton joints and then scale the encoded skeleton image (Fig. 2b left), produced inferior results compared to stacking the encoded images for hand and body skeleton joints (Fig. 2b right). This highlights the importance of scaling up the body skeleton image with a higher upscaling factor, similar to the body skeleton baseline (Fig. 2a).

However, we observed different results when comparing how much the hand skeleton joints need to be upscaled. Swin performed better with a smaller scale factor, while ResNet achieved better results with a larger scale factor. This could possibly be attributed to the different convolutions in the first layer of the respective networks - ResNet uses a 7×7 convolution with stride 2, while Swin’s patch embedder is a 4×4 convolution with stride 4.

We also compared single image fusion approaches to multiple image approaches in the Swin transformer. Unfortunately, the multiple image approaches were inferior to all other approaches. The median and maximum results were significantly worse, and the variance was much larger. This suggests

that this approach for multimodal input to a Swin network cannot be easily applied.

The lower performance of the multiple image approaches in Swin could potentially be attributed to the patch embedding process. This involves splitting the convolutions to different images to obtain the feature maps with the needed channel sizes. Furthermore, we experimented with larger patch embeddings as in [7], where the body skeleton image is processed into 96 channels of the feature map and the hands into 32 or both into 64. Although this improved the models and made them perform similarly to the single image approaches, it significantly increased the needed computational power and training time. In [7], it was shown that appropriate pre-training can be crucial. However skeleton-based pre-training is not typical in literature and also not the focus of this paper.

5.3 Experiments with 2D Body Skeletons

Most datasets and state-of-the-art approaches utilize 2D skeletons. Therefore, we also experiment with 2D skeletons and show results on the ATTACH [2] and the IKEA ASM [3] dataset. First, we present results of our body only baseline and afterwards the fusion with hand skeletons.

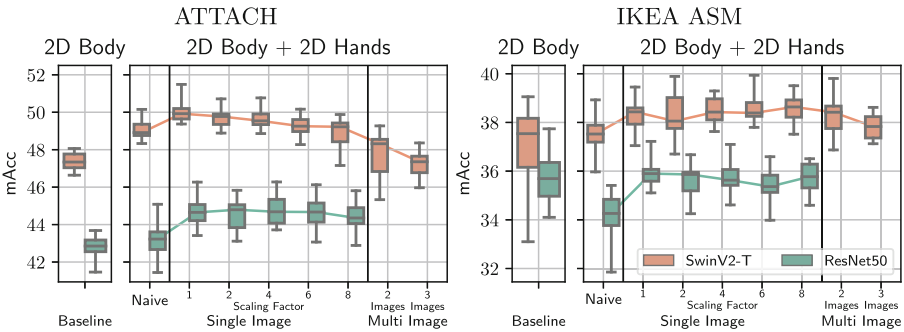


Fig. 4. Results using 2D body skeletons on the ATTACH and IKEA ASM datasets for our baseline models as in Fig. 2a and our different fusion methods: Naive concatenation as in Fig. 2b left, image concatenation with scaling of encoded hand skeleton image as in Fig. 2b right, multi image input as in Fig. 2c. Best results are listed in Table 1.

Baseline – 2D Body Skeletons In Fig. 4, we present the results of the baseline experiments with 2D body skeletons for each dataset. Firstly, it is important to note that the 2D body skeleton baseline results are worse compared to the 3D skeleton baseline results. This can be attributed to the loss of depth information when using 2D skeletons.

As observed in the previous section on using 3D skeletons, Swin outperforms ResNet on both datasets. However, as explained in Sect. 3.1 the skeleton-based

action recognition problem is very challenging on IKEA ASM due to a differentiation of actions by objects, which are not encoded in skeleton data. This could explain the smaller improvement in accuracy on IKEA ASM than on ATTACH.

Fusion of Hand Skeletons with 2D Body Skeletons Right to the respective baseline results in Fig. 4, we present the results of the fusion experiments with 2D hand and body skeletons. The comparison between the 2D body skeleton baseline and the fusion approaches reveals a notable improvement in classification performance for both the ATTACH and IKEA ASM datasets. Thus, the inclusion of hand skeletons in addition to body skeletons emerges as a highly effective strategy to elevate the accuracy of action recognition in assembly applications. Below, we go into more detail on the results for each dataset individually.

ATTACH: A closer look on the results on the ATTACH dataset and the comparison with 3D body skeletons reveals that hand skeletons are crucial for achieving improved performance with 2D body skeletons, as indicated by the greater improvement over the corresponding baseline. This holds true for both Swin and ResNet models, highlighting the significance of hand skeletons in mitigating the loss of depth information when only 2D body skeletons are available.

IKEA ASM: The results on the IKEA ASM dataset are less conclusive. Although the addition of hand skeletons generally leads to better medians and smaller variances, the improvement is not as clear as on the ATTACH dataset. Specifically, while the Swin and EMSAFormer models show clear improvement with the addition of hand skeletons, the ResNet only shows improvement in median. One possible explanation for this difference is that predicting hand skeletons on the IKEA ASM dataset is more challenging due to the small size of the hands, which often results in missing hand skeleton estimations. The attention mechanisms in the Swin transformer may be better suited to handle this issue of jumps in the temporal sequence, while the ResNet struggles with it and therefore processes the information contained in the hand skeletons less effectively.

Table 1. Best results of our experiments. We report the mean class accuracy mAcc and in parentheses the top-1 accuracy for the ATTACH [2] and IKEA ASM [3] datasets.

	ATTACH				IKEA ASM			
	3D Body (Baseline)	3D Body 3D Hand	3D Body 2D Hand	2D Body(Baseline)	2D Body 2D Hand	2D Body (Baseline)	2D Body 2D Hand	
ResNet50	48.2 (56.5)	48.2 (55.7)	48.3 (55.9)	43.7 (52.3)	46.3 (54.6)	37.7 (70.3)	37.2 (72.6)	
SwinV2-T	52.8 (60.3)	53.6 (61.0)	54.1 (61.7)	48.1 (56.0)	51.5 (58.9)	39.1 (72.6)	39.9 (73.9)	

6 Conclusion

Our work demonstrates a successful fusion of hand and body skeletons, which improves assembly action recognition notably. While hand skeletons contain important information, they are often prone to noise and misinformation due

to difficulties in estimation, such as occlusion or object/tool manipulation. To avoid this issue, our approach specifically handles the importance of the body skeletons to prevent the hand skeletons from dominating the input representation.

Furthermore, our approach demonstrates improved action recognition for two state-of-the-art assembly datasets, not only with 3D body skeletons but also with more commonly available 2D body skeletons. We have demonstrated a successful approach for preparing hand skeletons for action recognition and provided guidance on the key considerations for successful training with the Swin transformer. Overall, our work makes an important contribution to the field of action recognition in mobile robotics and collaborative robots.

References

1. Aganian, D., Köhler, M., Baake, S., Eisenbach, M., Gross, H.M.: How object information improves skeleton-based human action recognition in assembly tasks. In: IEEE International Joint Conference on Neural Networks (IJCNN) (2023)
2. Aganian, D., Stephan, B., Eisenbach, M., Stretz, C., Gross, H.M.: ATTACH dataset: annotated two-handed assembly actions for human action understanding. In: IEEE International Conference on Robotics and Automation (ICRA) (2023)
3. Ben-Shabat, Y., et al.: The IKEA ASM dataset: understanding people assembling furniture through actions, objects and pose. In: IEEE Winter Conference on Applications of Computer Vision (WACV) (2021)
4. Du, Y., Fu, Y., Wang, L.: Skeleton based action recognition with convolutional neural network. In: IEEE IAPR Asian Conference on Pattern Recognition (ACPR) (2015)
5. Duan, H., Zhao, Y., Chen, K., Lin, D., Dai, B.: Revisiting skeleton-based action recognition. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
6. Eisenbach, M., Aganian, D., Köhler, M., Stephan, B., Schröter, C., Gross, H.M.: Visual scene understanding for enabling situation-aware cobots. In: IEEE International Conference on Automation Science and Engineering (CASE) (2021)
7. Fishedick, S., Seichter, D., Schmidt, R., Rabes, L., Gross, H.M.: Efficient multi-task scene analysis with RGB-D transformers. In: IEEE International Joint Conference on Neural Networks (IJCNN) (2023)
8. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask R-CNN. In: IEEE International Conference on Computer Vision (ICCV) (2017)
9. Inkulu, A.K., Bahubalendruni, M.R., Dara, A., SankaranarayanaSamy, K.: Challenges and opportunities in human robot collaboration context of industry 4.0 - a state of the art review. *Ind. Robot: Int. J. Robot. Res. Appl.* **49**(2) (2021)
10. Liu, Z., et al.: Swin transformer v2: scaling up capacity and resolution. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
11. Mazzia, V., Angarano, S., Salvetti, F., Angelini, F., Chiaberge, M.: Action transformer: a self-attention model for short-time pose-based human action recognition. *Pattern Recogn.*, 124 (2022)
12. Ragusa, F., Furnari, A., Livatino, S., Farinella, G.M.: The MECCANO dataset: understanding human-object interactions from egocentric videos in an industrial-like domain. In: IEEE Winter Conference on Applications of Computer Vision (WACV) (2021)

13. Seichter, D., Köhler, M., Lewandowski, B., Wengefeld, T., Gross, H.M.: Efficient RGB-D semantic segmentation for indoor scene analysis. In: International Conference on Robotics and Automation (ICRA) (2021)
14. Sener, F., et al.: Assembly101: a large-scale multi-view video dataset for understanding procedural activities. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
15. Terreran, M., Lazzaretto, M., Ghidoni, S.: Skeleton-based action and gesture recognition for human-robot collaboration. In: International Conference on Intelligent Autonomous Systems (IAS). Springer (2022). https://doi.org/10.1007/978-3-031-22216-0_3
16. Trivedi, N., Sarvadevabhatla, R.K.: PSUMNet: unified modality part streams are all you need for efficient pose-based action recognition. In: ECCV Workshop and Challenge on People Analysis (WCPA). Springer (2022). https://doi.org/10.1007/978-3-031-25072-9_14
17. Trivedi, N., Thatipelli, A., Sarvadevabhatla, R.K.: NTU-X: an enhanced large-scale dataset for improving pose-based recognition of subtle human actions. In: Indian Conference on Computer Vision, Graphics and Image Processing (ICVGIP). ACM (2021)
18. Wang, L., et al.: Symbiotic human-robot collaborative assembly. *CIRP annals* **68**(2) (2019)
19. Zhang, F., et al.: MediaPipe hands: on-device real-time hand tracking. In: Workshop on Computer Vision for AR/VR (CV4ARVR) (2020)
20. Zhang, P., Lan, C., Xing, J., Zeng, W., Xue, J., Zheng, N.: View adaptive neural networks for high performance skeleton-based human action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* (2019)