# How Object Information Improves Skeleton-based Human Action Recognition in Assembly Tasks

Dustin Aganian, Mona Köhler, Sebastian Baake, Markus Eisenbach, and Horst-Michael Groß

Ilmenau University of Technology, Neuroinformatics and Cognitive Robotics Lab

98684 Ilmenau, Germany

dustin.aganian@tu-ilmenau.de, ORCID: 0009-0006-3925-6718

Abstract—As the use of collaborative robots (cobots) in industrial manufacturing continues to grow, human action recognition for effective human-robot collaboration becomes increasingly important. This ability is crucial for cobots to act autonomously and assist in assembly tasks. Recently, skeleton-based approaches are often used as they tend to generalize better to different people and environments. However, when processing skeletons alone, information about the objects a human interacts with is lost. Therefore, we present a novel approach of integrating object information into skeleton-based action recognition. We enhance two state-of-the-art methods by treating object centers as further skeleton joints. Our experiments on the assembly dataset IKEA ASM show that our approach improves the performance of these state-of-the-art methods to a large extent when combining skeleton joints with objects predicted by a state-of-the-art instance segmentation model. Our research sheds light on the benefits of combining skeleton joints with object information for human action recognition in assembly tasks. We analyze the effect of the object detector on the combination for action classification and discuss the important factors that must be taken into account.

# I. INTRODUCTION

In the course of industry 4.0, cooperation between humans and situation-aware collaborative robots (cobots) is becoming increasingly important for manufacturing processes [1]–[3]. In order to effectively assist in assembly processes, cobots must first be able to visually perceive the worker and recognize the current assembly state [4]–[9]. One crucial aspect of attaining this goal is human action recognition [7], [9].

For human action recognition, often RGB-based approaches are used. However, these approaches face several challenges in scenarios where the training dataset and the application environment do not have a significant overlap. Due to the difficulties in creating action recognition datasets for assembly processes [12], these datasets tend to be smaller compared to other areas such as object detection [13]. As a result, RGBbased approaches trained on these datasets tend to overfit and exhibit poor generalization capabilities. This can lead to difficulties when the person performing the action changes or when the background or working environment during deployment are drastically different to training. To overcome these limitations, skeleton-based action recognition models are the better alternative, as they tend to generalize much

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



Fig. 1. Overview of our approach. We combine skeleton joints and object information for two human action recognition methods, namely a) VA-CNN [10] and b) PoseConv3D [11]. Note that, for PoseConv3D, we actually have one heatmap per joint and stack frames in an additional 4<sup>th</sup> dimension. The different colors for object centers are used for visualization purposes only. Example frames taken from IKEA ASM dataset [12].

better under such circumstances [9]. Nevertheless, skeletonbased models have the drawback of not being able to process objects involved in actions, such as when a worker picks up a hammer or a screwdriver. Similarly, actions in which workers pick up similar-sized assembly parts, such as a cabinet side panel or a cabinet back panel, cannot be distinguished based on skeletons.

In contrast to action recognition, for object detection there are highly diverse, public datasets [13] available, which enable the training of very accurate and well generalizing object detectors [14]. Therefore, our goal in this paper is to develop a novel approach to improve human action recognition by incorporating additional object information into skeleton-based action recognition for trimmed sequences. We will demonstrate that a relatively straightforward enhancement of the input encoding (as shown in Fig. 1) can greatly improve the performance of two different methods. Thereby we use the state-of-the-art dataset IKEA ASM [12], a comprehensive assembly dataset for action recognition with human-object interactions. In order to analyze the performance gain of combining object information with skeleton information, we first experiment with ground truth object masks, which are independent of the accuracy of an object detector. Afterwards, we examine the effect of using predicted objects masks and present the impact on the performance caused by noisier predictions.

Our contributions in this paper can be summarized as follows:

- 1) We propose a novel approach to integrate object information into skeleton-based action recognition.
- Our experiments with two different methods for skeleton-based action recognition demonstrate the usefulness of combining skeletons and object positions.
- 3) We show the influence of ground truth versus predicted objects on the performance of action recognition.
- Overall, our approach massively improves the performance of action recognition on the IKEA ASM dataset [12].

# II. RELATED WORK

In this section, we will first present an overview of the current state of the art for skeleton-based human action recognition, and then specify the methods that we will modify for incorporating object information as well. Afterwards, we will outline other work in the field of human action recognition that specifically incorporate object information.

#### A. Skeleton-based Action Recognition

The general term human action recognition encompasses a wide range of subfields. In this paper, however, we want to focus on the action classification task of pre-trimmed video clips of human skeleton sequences. This task serves as a foundation, since other methods for tasks like action segmentation or action detection, such as PDAN [15] or MS-TCT [16], typically use models trained on action recognition of pre-trimmed clips as the backbone of their method.

Regarding skeleton-based action recognition, RNNs have been used for a long time [17]. Currently 2D CNNs such as VA-CNN [10], 3D CNNs like PoseConv3D [11] and graph convolution networks (e.g. 2S-AGCN [18], Si-GCN [19]) are typically used.

In this paper, we enhance the VA-CNN and PoseConv3D methods that both achieve state-of-the-art results on common skeleton-based action recognition benchmark datasets such as NTU RGB+D [20]. Furthermore, both methods allow for our modifications to incorporate object information sequences in addition to skeleton sequences. Here, VA-CNN suited itself because it is able to train on very long or short clips easily, since it can process any length of clips due to its input coding. Our preliminary studies have also shown that VA-CNN

achieved better results than other methods we tested, such as 2S-AGCN. Furthermore, we chose PoseConv3D because, unlike most other skeleton-based action recognition methods, it is able to easily handle multiple people in a frame and thus multiple instances of the same joint class. Why this is advantageous and how the two selected methods VA-CNN and PoseConv3D as well as our modifications work in detail will be described in Sec. IV.

#### B. Action Recognition Incorporating Object Information

Object information can play a crucial role in improving action recognition. Research in this area typically focuses on advancing human action recognition. Some other studies are done in the field of human-object interaction recognition, with the goal of recognizing actions, where the objects used are included in the action labels.

An example of such work is [21] in which skeletons are linked with object position information. However, this combination does not take the object class into account, nor can more than one object be processed. In contrast, the presented method in [22] is not limited to the number of objects used. Here, using a rather simple graph network on a very small dataset, the authors demonstrated that the combination of object information with skeleton information generally works to train an action classifier. In the recently published work [23], object information is included in a graph convolutional network to perform action segmentation. In their experiments, ground truth object boxes were used exclusively, which benefited the developed method. However, in our experiments, we show, that the performance using ground truth objects does not resemble the real performance when using detected objects. Furthermore, in [23] neither the contribution of using the object boxes on the overall performance was shown, nor the use of detected objects was investigated.

In addition to these studies, there are works which investigate the use of object information through simulation. In [24], a method is introduced for simulating human actions and object movements to predict actions and activities. For the prediction, RNNs and Transformers were employed, but the performance of the developed methods has not been assessed on real datasets.

In contrast to related work, our approach can handle multiple objects, while simultaneously encoding class information along with the object positions. Furthermore, we show the general performance gain of combining object information with skeleton information compared to using solely skeleton information, an approach which has only been sparsely studied in the current state of the art. Moreover, we demonstrate how our methods perform on a large dataset with real estimated object information instead of simply using ground truth object masks.

#### III. USED DATASET

As previously described, our focus in this paper is to improve action recognition for cobots in the context of manufacturing processes. Action recognition datasets that are closest to this scenario are typically assembly datasets. Among them are datasets mostly recorded in first-person perspective or a very close surveillance perspective, such as Assembly101 [25] or Meccano [26]. Yet there are also datasets like IKEA ASM [12] and ATTACH [27], which were recorded from a surveillance perspective, which is very similar to the perspective of a cobot that observes a worker in order to help out.

For our experiments, we have decided to use the furniture assembly dataset IKEA ASM, which is a very extensive dataset, captured from multiple perspectives. It consists of scenes where different furniture types are assembled in various environments, making the dataset quite challenging. Moreover, in contrast to ATTACH, it also contains labeled object masks.

IKEA ASM consists of 371 distinct assembly processes, where 48 participants were instructed to assemble one of four different furniture types. Each assembly process is captured by three different Kinect V2 cameras, resulting in 1113 videos and 35 hours of footage at 24fps. For human action recognition, there are 17K labeled action instances, distributed over 33 atomic classes. We use the official splits provided in [12], where about  $\frac{2}{3}$  is in training and  $\frac{1}{3}$  in testing. The pieces of furniture consist of seven different object categories. For the object instances 1% of the data were manually labeled. The remaining annotations were obtained by overfitting several models on the annotated data, resulting in pseudo ground truth data. However, the object masks are only available for one view as shown in Fig. 2 at the top. For skeleton data IKEA ASM provides 2D skeleton predictions obtained from OpenPose [28] and Keypoint R-CNN [14], of which we use the better performing Keypoint R-CNN in our experiments.

The dataset provides official training and test splits, which we also use in our experiments.

#### IV. APPROACH

For integrating additional object information into skeletonbased action recognition, we build our approach upon two different state-of-the-art methods for skeleton-based action recognition, namely VA-CNN [10] and PoseConv3D [11].

#### A. VA-CNN Modification

VA-CNN [10] is a rather fast method while still achieving competitive results. Different to many other methods, the whole trimmed action sequence is utilized instead of a smaller clip. The main parts of VA-CNN are a specific input encoding, a view-adaptation module, and a classification network, which we will all describe in the following.

a) Input Encoding: VA-CNN uses the input encoding from [29]. The skeletons for the whole action sequence are transformed into a single RGB image. One column represents one frame of the action and in each column the skeleton joints are stacked in a fixed order. The XYZ coordinates of a joint are transformed to RGB by normalizing and scaling:

$$u_{t,j} = 255 \cdot \frac{v_{t,j} - c_{\min}}{c_{\max} - c_{\min}} ,$$
 (1)

where  $v_{t,j}$  denotes the 3D coordinates of the  $j^{th}$  joint of the  $t^{th}$  frame in a skeleton sequence,  $c_{\min}$  and  $c_{\max}$  are the minimum

and maximum of all joint coordinates in the training data, and  $u_{t,j}$  refers to the normalized coordinates. As we work with the 2D skeletons from IKEA ASM [12], we only have an image with red and green channels and an empty blue channel. The resulting output is shown in Fig. 1. For training and application, each resulting image can then be transformed to the appropriate input size for the classification network using image operations.

*b)* View-Adaptation Module: The view-adaption (VA) module aims at normalizing the skeletons to one resembling view. However, we have found that, when using 2D skeletons, this module does not work as intended and has no effect on the final performance. Therefore, we omit the VA part and refer to this modified VA-CNN as 2D-CNN in the remainder of this paper.

c) Classification Network: Here, any classification network can be utilized. Therefore, we use the popular and reliable ResNet-50 [30], which was also used in VA-CNN.

*d)* Incorporating Object Information: Our general idea for incorporating object information is to consider objects similar to further skeleton joints. Thus, we combine object data with skeleton data as early as in the input encoding. We modify the input encoding by appending additional rows below the skeleton joints. Specifically, we append one additional line for each object class that occurs in the dataset.

Since skeleton joints are represented as points, we also need to convert object masks to this format. Therefore, we calculate the center of mass for each object mask in order to represent each object as single point as well. These object coordinates can then be inserted in the appended lines below the skeleton joints. As not all object classes appear in all frames, we set the coordinates for the missing classes to zero, as can be seen by the black parts in Fig. 1.

Unfortunately, this encoding for 2D-CNN has a drawback. Since a human typically has only one of each type of skeleton joints, the encoding was not designed to include multiple instances of the same joint class. Thus, we are presented with a challenge when dealing with varying numbers of objects in each sequence or sometimes even each frame. For example, during assembling a table, there are multiple table legs, but only one of these legs can be encoded. Even if each object class occurs only once, this issue can still occur when working with a real object detector, which also produces false positives that cannot be filtered completely by tuning the confidence threshold.

Therefore, we decided to only use one object coordinate of each object class. This leaves us with the question which object to use if there are multiple instances of the same class. To overcome this challenge, we use a simple heuristic: The objects a human interacts with are more relevant to action recognition than others. For interacting with objects, humans typically use their hands. Therefore, we calculate the euclidian distance of all object coordinates to the predicted hand joints, and choose the object with the lowest distance to both hand joints for each class. Alternatively, we also experimented with the lowest distance to only one hand joint, which led to slightly



Fig. 2. Instance segmentation of our Mask R-CNN with Swin-Tiny backbone trained on IKEA ASM [12]. The colors of the seven object classes are shown on the right. Ground truth labels are only available for the top view. We also visualize the center of the object masks with a white dot. All predictions with a confidence score s > 0.1 are shown, as this liberal threshold performed best in our action recognition experiments.

worse results. As we can see in Fig. 2, our heuristic chooses the table leg that is most important for the current assembly step.

In our experiments, we show that our enhancement for incorporating object information results in an exceptionally high boost in action recognition performance.

## B. PoseConv3D Modification

In contrast to VA-CNN, PoseConv3D [11] is a more complex method, using 3D convolutions, while processing 2D skeletons. Moreover, PoseConv3D is able to process multiple instances of the same skeleton joint class, since it can handle multiple detected human skeletons, thereby processing the joints of the group of persons.

For PoseConv3D the input encoding is also the most interesting part: The skeleton joints of a single frame are embedded in heatmaps – one heatmap for each joint class – which are stacked along the channel axis, resulting in a heatmap tensor of size  $J \times H \times W$ , where J is the number of distinct joint classes and H and W refer to the height and width of the cropped frame. The heatmaps are obtained as following: The  $j^{th}$  joint coordinate  $(c_x^{(j)}, c_y^{(j)}, s^{(j)})$  with the location  $(c_x^{(j)}, c_y^{(j)})$  and the confidence score  $s^{(j)}$  is embedded in a heatmap K with a Gaussian kernel centered at the joint:

$$K_{j,x,y} = e^{-\frac{\left(x - c_x^{(j)}\right)^2 + \left(y - c_y^{(j)}\right)^2}{2 \cdot \sigma^2}} \cdot s^{(j)}, \tag{2}$$

where  $\sigma$  denotes the variance of the Gaussian and is set according to the original PoseConv3D implementation [11]. For each joint, the multiple heatmaps over time are then stacked in an additional time axis T, resulting in a 4-dimensional input  $x \in \mathbb{R}^{J \times T \times H \times W}$ , which is illustrated in Fig. 1. This input is then fed into the 3D-CNN SlowOnly [31], which is directly inflated from ResNet-50.

Similar to our 2D-CNN method, we can append object information to the skeletons, here by using additional heatmaps. For this, we again consider the object centers as further skeleton joints and append one additional heatmap for every object class. In contrast to 2D-CNN, PoseConv3D offers the major advantage that it can handle multiple instances of the same joint class. This means that we are not restricted to one single object per class, but can simply draw a Gaussian for each object of the same class in one single heatmap. Likewise, for detected object masks, we have the advantage of also modeling the estimated confidence directly with the Gaussian.

In order to prove the usefulness of object information for human action recognition, we experiment with both ground truth as well as detected object masks from an instance segmentation model.

#### V. SETUP

For our implementation, we use the publicly available code from VA-CNN<sup>1</sup> for our 2D-CNN method and MMAction2 [32] for PoseConv3D. Overall, we followed the original pipelines and hyperparameters and only adapted the following: For training, we use the Adam optimizer [33] and a OneCycle learning rate scheduler with a 10% warmup and different maximum learning rates. In general, we always perform a linear learning rate search close to the original learning rates and repeat each configuration two to three times to obtain consistent results.

We evaluate the performance of our models using mean class accuracy (mAcc) as well as top-1 accuracy (top1) as these are commonly used for pre-trimmed sequences in many action recognition papers [25], [26], [34], [35]. For the top1 metric, all action clips contribute equally. In contrast, mAcc averages accuracy across all action classes, which compensates for class imbalances.

<sup>&</sup>lt;sup>1</sup>https://github.com/microsoft/View-Adaptive-Neural-Networks-for-Skeleton-based-Human-Action-Recognition

Our experiments are sectioned into two parts: First, in Sec. VI, we analyze the impact of integrating additional object information on human action recognition using ground truth data. Afterwards, in Sec. VII, we train our own instance segmentation model and use the resulting object mask predictions in combination with skeletons for action recognition.

#### VI. EXPERIMENTS WITH GROUND TRUTH OBJECT MASKS

In the following, we analyze the extent to which combining object centers with skeleton joints can improve skeleton-based human action recognition. In order to demonstrate the general viability of our approach and to be independent from the performance of an object detector, we first use the ground truth data of IKEA ASM for objects (Sec. VI-B). We also present an ablation study on a modified IKEA ASM dataset to investigate the benefits of the combination, even if the object classes have a small influence on the classes to be predicted (Sec. VI-C). The subsequent section Sec. VII then presents our experiments with object masks, determined by an instance segmentation model.

# A. Preliminary Considerations for Using Ground Truth Objects

When training with ground truth data for object information, we are faced with the issue that the IKEA ASM dataset has three views, but objects are only annotated in the top view as shown in Fig. 2. To solve this issue, we considered the following solutions:

a) Only Use Top View Data: One option could be to only use the camera view on which the objects have been labeled and discard the other two views. However, this leads to a smaller and simpler data subset of only one third. Furthermore, it limits comparability to predicted object masks in Sec. VII.

b) Use All Views, but Add Objects Only to Top View: Another option could be to use skeleton data from all views and combine them with object data only in case of the top view. This allows us to compute results over the whole test set, leading to better comparability. However, as 2/3 of the dataset would contain no object information at all, the impact of additional object information might be limited and it might be hard to asses the actual potential.

c) Add Object Centers to Every Skeletal Perspective: In order to increase the impact of additional object information, we could also add the object centers from the top view to all views. This seems counterintuitive at first, as the object center coordinates do not match the skeleton joint coordinates for front and side view. As we are only dealing with 2D data and have neither depth information nor extrinsics for all three views, we are also unable to transform the data to a corresponding perspective. Yet, in preliminary experiments with 2D-CNN, we found that option c) leads to the best results, i.e., 47.8 mAcc, compared to only 37.5 mAcc for option b). This already shows how important object information seems to be, if the mAcc is up to 10 percentage points worse, if 2/3 of the training and test data consisted of skeletal data only. Therefore, we used option c) for all experiments in this section,

enabling us to evaluate the potential of combining object information with skeleton joints for human action recognition.<sup>2</sup>

# B. Assessing the Potential of Skeleton-Object-Combination

a)  $\ddagger$  – Skeleton Only Baseline: To establish a general baseline, we first trained both 2D-CNN and PoseConv3D using only skeletons. As we can see in the first row of Tab. I, the much more complex PoseConv3D is also significantly better on only skeleton sequences, so the results from the state of the art also apply to the IKEA ASM dataset.

 TABLE I

 Results using Ground Truth Object Masks on IKEA ASM [12]

			2D-CNN		PC3D		
£	Î	used objects	mAcc	top1	mAcc top1		
1			37.7	70.3	39.5 75.0		
	\ \	most relevant all	44.0	67.0 —	45.8 75.0 51.7 78.5		
\ \	\ \	most relevant all	47.8	79.6 —	56.5 84.3 <b>58.5 85.4</b>		

b)  $\hbar + \Im - Skeleton-Object-Combination:$  For analyzing the combination of objects and skeletons, it must be taken into account that there are multiple ways to combine object information with skeleton joints, depending on the used methods 2D-CNN and PoseConv3D (as described in Sec. IV). Therefore, we carried out various experiments in order to be able to compare the methods as fair as possible.

For our 2D-CNN method, we are restricted to only one object per class, which we will refer to as the *most relevant objects* in the following. In contrast, for PoseConv3D this restriction does not apply. Thus, first we trained PoseConv3D with *all objects* in every frame and second also only with the *most relevant objects* for better comparability to 2D-CNN. Tab. I presents the results of these experiments.

Compared to using only skeletons, the combination with object information improves the performance to a large extent for both methods. For 2D-CNN, we can gain about 10 percentage points mAcc. For PoseConv3D the improvement is even greater with 19 percentage points mAcc. As expected, PoseConv3D performs better when it has access to all object data. These results show how important object information is for skeleton-based action recognition when the acting person interacts with objects. And thus, these results also show that purely skeleton-based inputs are severely disadvantaged in such action recognition tasks and that fundamentally important information is typically lost when using only skeleton sequences.

c)  $\bigcirc$  – Only Ground Truth Object Information: Due to this large performance gain, we decided to investigate how well the action classifier performs when we train our models using only ground truth object information, and consequently

 $<sup>^{2}</sup>$ As a side note: The combination of skeleton data and object data according to option a), i.e., using only the top view for training and testing, the 2D-CNN method achieved a mAcc of 46.6% on the limited test dataset.

omit skeleton data. The results are shown in the middle rows of Tab. I. Surprisingly, the models trained only on ground truth object information show a very large performance gain compared to the models trained only on skeleton information. This result shows how much a perfect object detector can contribute to robust action recognition. Nevertheless, it must be taken into account that this comparison is not completely fair here, since the object models were trained and tested exclusively on the top view perspective. A more fair comparison in which we only use real predicted object information on all perspectives is presented in Sec. VII.

Furthermore, we have to keep in mind, that the actions of the IKEA ASM dataset such as *pick up leg* or *pick up shelf* can be much better distinguished by moving objects than by a moving skeleton. Therefore, we also carried out an ablation study on a dataset with combined action classes, which we will describe in the following section.

#### C. Verb-focused Action Recognition Study

The IKEA ASM dataset contains classes for action recognition, which can only be distinguished by the object used, such as: attach back panel, attach side panel, or attach shelf. For the subject area of human-object interaction recognition, this question is highly relevant, but for some applications, such as general human action recognition, it is sufficient to know only the action in general, i.e., the verb describing the action. In order to investigate the usefulness of combining object information with skeleton joints for such application domains, we modify the IKEA ASM dataset such that the action classes to be predicted are no longer so heavily dependent on the objects used. Thus, we summarize the 33 action classes based on the verbs, resulting in an only verbs version of the IKEA ASM dataset. This leads to 12 remaining action classes (align, attach, flip, insert, lay down, pick up, position, push, rotate, slide, spin, tighten).

Subsequently, we repeated experiments conducted above on the *only verbs* dataset and present results in Tab. II. Note that these results cannot directly be compared to the results in other tables, as the *only verbs* dataset consists of combined action classes, leading to fewer confusions between different classes.

 TABLE II

 Results on our Only Verbs Version of IKEA ASM [12]

0			2D-CNN		PC3D	
Т	<b>S</b>	used objects	mAcc	top1	mAcc	top1
1			63.7	82.6	64.7	85.8
	\ \	most relevant all	63.8	78.3	74.2 77.4	83.6 88.9
√ √	\ \	most relevant all	69.4 —	86.7 —	<b>81.4</b> 81.1	91.5 <b>92.0</b>

First, generally speaking, the same trends can be seen as on the original IKEA ASM dataset. The more computationally intensive PoseConv3D is generally always better than the 2D-CNN. Likewise, the combination of skeleton joints and object information again shows to be enormously beneficial. The mAcc for 2D-CNN is improved by 6 percentage points and for PoseConv3D by 17 percentage points compared to training only on skeleton joints. While using only objects results in similar mAcc performance to using only skeletons for 2D-CNN, we are still observing a large performance gain of 13 percentage points for PoseConv3D.

Therefore, we could effectively demonstrate that the performance gain of incorporating object information cannot only be attributed to the way the action classes are defined in IKEA ASM dataset. On the contrary, even for the verbs-only action classes, object information still increases performance to a large extent.

# VII. EXPERIMENTS WITH PREDICTED OBJECT MASKS

In the previous section, we presented the potential of combining object information with skeleton joints by utilizing the ground truth object masks. For the following experiments, we first train an instance segmentation model (Sec. VII-A). Afterwards, we combine its predicted object masks with skeleton joints for human action recognition (Sec. VII-B). Thus, we show the impact of using non-perfect predicted object masks, which typically also contain false positive as well as false negative predictions.

As mentioned before, so far we could only use object coordinates from top view, since there are no labels for the other two views. In this section, we will counteract this issue by applying our trained instance segmentation model to predict object masks for all views. This allows us to combine object and skeleton information from corresponding perspectives. Moreover, we can investigate how non-matching object and skeletal perspectives might have negatively affected the learning process, leading to potentially worse performance.

#### A. Instance Segmentation

For training an instance segmentation model on the IKEA ASM dataset, the 1% manually labeled frames from top view are used. The authors of the IKEA ASM dataset [12] already provide pretrained weights for Mask R-CNN [14] with ResNet-50, ResNet-101 as well as ResNeXt-101 backbone. Moreover, we also train our own Mask R-CNN with a Swin Transformer Tiny [37] backbone, as the Swin Transformer is a state-of-the-art model outperforming ResNet in many applications. For training on the IKEA ASM dataset, we used a model pretrained on MS COCO [13] and finetuned it for

TABLE III Performance of Mask R-CNN instance segmentation models with different backbones on IKEA ASM [12].

backbone	weights	mask AP AP50		bbox AP AP50	
ResNet-50 [30] ResNet-101 [30] ResNeXt-101 [36] Swin-Tiny [37]	[12] [12] [12]	58.1 62.1 <b>65.9</b> 63.9	77.2 82.0 85.3 <b>89.4</b>	59.5 64.6 69.5 <b>72 5</b>	77.7 81.8 86.4 92.1

12 epochs using the AdamW optimizer [38] and a OneCycle learning rate scheduler with a maximum learning rate of  $5 \cdot 10^{-5}$ . While having similar complexity with ResNet-50, our Swin-Tiny Mask R-CNN reaches an even higher performance than the ResNeXt-101 model in most metrics, as shown in Tab. III. Qualitative results of our trained model on all three views are shown in the second row of Fig. 2. Overall, we observed very few false negative predictions, but – depending on the confidence threshold – several false positive predictions.

After training our model, we predicted object masks for all frames of all views. We saved all predictions with a confidence score s > 0.1, which allows us to later experiment with different thresholds for human action recognition.

#### B. Human Action Recognition

In the following, we first describe the main motivation of the series of experiments in this section and the resulting different training settings we need to consider. We also describe how we choose which of the dedicated object masks to use for training and testing, and the different resulting training configurations. Finally, we conclude with the experimental results and compare them to our baseline, as well as the results on ground truth object masks.

a) Preliminary Considerations: The purpose of the following experiments on human action recognition is twofold and thus we have to consider two different training settings.

- Top View: On the one hand, we want to evaluate the potential performance drop of using predicted instead of ground truth object masks. In order to compare these results with our experiments in Sec. VI, we also use predicted object masks only for the top view, but combine them with skeletons from the other two views as well.
- 2) All Views: On the other hand, we want to investigate the benefit of combining real predicted object masks and skeleton joints for action recognition closer to a real-world application. Thus, the perspective of the predicted objects mask should always match the perspective of the skeleton joints. This should also make it easier for the models to combine the data appropriately and facilitate the learning process.

b) Selecting the Detected Object: For 2D-CNN, we are restricted to one object per class, which we determine by calculating the distance to the hand joints. However, different confidence thresholds for the predicted objects lead to different chosen objects. Therefore, we performed an equidistant line search to determine the best confidence threshold for all experiments. Overall, a threshold of  $\tau_s = 0.1$  led to the best results, and we present all results using this threshold. We assume that the distance to the hand joints is a good heuristic for obtaining the most relevant object, even when the predictions contain several false positives as shown in Fig. 2. We therefore hypothesize that higher thresholds with more false negatives (i.e., more missing objects) have a greater adverse effect than smaller thresholds with more false positives. In the following



Fig. 3. Comparison of action recognition performance on IKEA ASM [12] with  $\hat{\lambda}$ : only skeletons,  $\mathfrak{D}$ : only objects and  $\hat{\lambda} + \mathfrak{D}$ : the combination of both.

experiments, models trained with this kind of object preselection are again marked as *most relevant* objects used.

For PoseConv3D, this restriction does not apply. We are able to use all objects, and scale the maximum values in the heatmaps by the predicted confidence scores. Implicitly, this should give the model the intuition that objects with a higher confidence value might be more relevant. Experimental results from this training setting are consequently marked as *all* objects used. However, for comparison reasons we also experiment on PoseConv3D with *most relevant* objects used.

We use the predicted object masks from Mask R-CNN with Swin-Tiny backbone in all of our experiments in this section. As the results on instance segmentation in Tab. III have already indicated, using predictions from other models led mostly to worse action recognition results in preliminary experiments and for the more computationally expensive model with ResNext-101 backbone to no clear advantage.

c) Results with Detected Object Masks: We present the results for the settings described above for our experiments of combining detected object masks and skeleton joints in Tab. IV.

Likewise, in Fig. 3 we present these results in comparison to our skeleton only baseline and the experiments with the combi-

 TABLE IV

 Results using Detected Object Masks on IKEA ASM [12]

				2D-CNN	PC3D
	犬	Î	used objects	mAcc top1	mAcc top1
top	\	\	most relevant	43.9 75.6	44.4 75.1
view	\	\	all		41.6 72.1
all	\	\	most relevant	47.2 77.7	49.0 <b>80.2</b>
views	\	\	all		<b>49.9</b> 79.7

nation of ground truth object masks. For a better overview, this figure shows only the best results of PoseConv3D regarding the setting which objects were used.

Examining these results, the obvious first observation is again that PoseConv3D performs better than 2D-CNN on the combination of skeleton joints and objects. However, the difference is nowhere near as large as in the previous comparisons. When comparing the skeleton-only baseline and the combination of skeleton joints with detected object masks, a large performance gain can be seen in Fig. 3. On *all views*, for example, both the 2D-CNN and the PoseConv3D are about 10 percentage points better on the mAcc. This shows that our modified methods can handle imperfectly predicted object masks well, and thus we highly encourage to also combine skeletons with predicted object masks for human action recognition in real-world applications.

When comparing the results with ground truth object masks and the results with predicted object masks, we need to look at the experiments on *top view* respectively. Here we can see that the models produce worse results, with the PoseConv3D showing a much greater performance drop than the 2D-CNN. This shows that our instance segmentation model, which achieves a very high segmentation AP50 of 89.4, still has significant weaknesses in predicting object masks, and that these false detections make action recognition significantly more difficult compared to the usage of ground truth object masks.

However, our instance segmentation model offers a major benefit in that it can predict object masks for all perspectives. When comparing the results of combining predicted object masks and skeleton joints for *top view* and *all views*, we clearly see a considerable improvement. The 2D-CNN effectively utilized predicted object masks on *all views* and skeleton joints to achieve results that are on par with those obtained with ground truth object masks. This indicates that the influence of noisy predictions can be better compensated by the models when the perspective of the predictions fits the perspective of the skeleton joints.

The impact of noisy predictions is also shown by the PC3D results when comparing the difference between *most relevant* and *all* objects over *top view* and *all views*. When training with less data (*top view*), it has shown to be necessary to filter the *most relevant* objects to handle the noisy predictions. When training on more data (*all views*), the model learns to compensate for the noisy predictions by itself.

Finally, we examine the results utilizing only object masks without skeletons. As expected, the models consistently achieve worse results than with ground truth object masks. However, we can also see that the results for predicted object masks are significantly better on *all views* compared to only using *top view*. Again this shows that the models trained on *all views* benefit from more diverse training data from different perspectives to compensate for the noisy predictions.

d) Training on Ground Truth and Testing on Detected Object Masks: For human action recognition, one could also consider training with ground truth objects and then apply the trained action recognition model and test with predicted objects. However, we want to emphasize that this setting leads to considerably worse results. Our 2D-CNN trained with ground truth object masks and tested with predicted object masks (both from *top view*) achieves a mAcc of only 40.2, which is 4 percentage points worse than if we had trained directly with predicted object masks from *top view*. These results show that the performance drops considerably, if the object detector used in a real-world application does not match the object detector used for training. This is consistent with literature, where this phenomenon is described as sudden concept drift between training and deployment [39], [40].

#### VIII. CONCLUSION

In this paper, we extended two state-of-the-art methods for skeleton-based action recognition to be able to process object information in addition to skeleton joints. Using these modified methods, we then performed investigations on the challenging assembly dataset IKEA ASM [12]. We evaluated the benefit of object information for action classification of human-object interaction classes, as well as for verb-based action classification only. We also showed the importance of matching the perspective of objects to skeletons and examined the influence of noisy predictions of an object detector on the performance of action classification.

Overall, incorporating additional object information into skeleton-based action recognition improves the performance to a large extent. Our proposed approach works extremely well for the two considered state-of-the-art skeleton-based action recognition methods 2D-CNN and PoseConv3D. Moreover, we assume that our approach of treating objects similar to further skeleton joints can also be transferred to other endto-end deep-learning-based methods for skeleton-based action recognition.

However, the combination of object and skeleton information is a relatively new area of study with ample opportunities for future exploration. For example, one could also think of using complete object masks, instead of the center of mass, e.g., the input encoding from PoseConv3D would be very suitable for this. Moreover, instead of combining the inputs, one could also use a dual-branch network and fuse skeleton and object information in the network. Furthermore, this approach could also be applied to other applications besides assembly, where object information also plays a great role, such as cooking applications (e.g. on the dataset EPIC-KITCHENS-100 [41]). Perhaps it could even improve action recognition for more general applications, for which the Kinetics400 dataset [34] could be utilized in combination with a general object detector trained on MS COCO [13].

With our findings, we hope to have paved the way for more researchers to experiment with combining object information and skeleton data for human action recognition. We expect that more research in this area enables cobots to recognize human actions more robustly and accurately. This would improve human-robot collaboration, which is one of the major goals of industry 4.0.

#### REFERENCES

- E. Matheson, R. Minto, E. G. Zampieri, M. Faccio, and G. Rosati, "Human–robot collaboration in manufacturing applications: A review," *Robotics*, vol. 8, no. 4, p. 100, 2019.
- [2] F. Sherwani, M. M. Asad, and B. S. K. K. Ibrahim, "Collaborative robots and industrial revolution 4.0 (IR 4.0)," in *Int. Conf. on Emerging Trends* in Smart Technologies (ICETST). IEEE, 2020, pp. 1–5.
- [3] A. K. Inkulu, M. R. Bahubalendruni, A. Dara, and K. SankaranarayanaSamy, "Challenges and opportunities in human robot collaboration context of industry 4.0 – a state of the art review," *Industrial Robot: Int. Journal of Robotics Research and Application*, vol. 49, no. 2, pp. 226–239, 2021.
- [4] L. Wang, R. Gao, J. Váncza, J. Krüger, X. V. Wang, S. Makris, and G. Chryssolouris, "Symbiotic human-robot collaborative assembly," *CIRP annals*, vol. 68, no. 2, pp. 701–726, 2019.
- [5] M. Al-Amin, W. Tao, D. Doell, R. Lingard, Z. Yin, M. C. Leu, and R. Qin, "Action recognition in manufacturing assembly using multimodal sensor fusion," *Procedia Manufacturing*, vol. 39, pp. 158–167, 2019.
- [6] S. Li, P. Zheng, J. Fan, and L. Wang, "Toward proactive human-robot collaborative assembly: A multimodal transfer-learning-enabled action prediction approach," *IEEE Transactions on Industrial Electronics*, vol. 69, no. 8, pp. 8579–8588, 2021.
- [7] M. Eisenbach, D. Aganian, M. Köhler, B. Stephan, C. Schröter, and H.-M. Gross, "Visual scene understanding for enabling situation-aware cobots," in *IEEE Int. Conf. on Automation Science and Engineering* (CASE), 2021.
- [8] J. Fan, P. Zheng, and S. Li, "Vision-based holistic scene understanding towards proactive human–robot collaboration," *Robotics and Computer-Integrated Manufacturing*, vol. 75, p. 102304, 2022.
- [9] M. Terreran, M. Lazzaretto, and S. Ghidoni, "Skeleton-based action and gesture recognition for human-robot collaboration," in *Int. Conf.* on Intelligent Autonomous Systems (IAS). Springer, 2022, pp. 29–45.
- [10] P. Zhang, C. Lan, J. Xing, W. Zeng, J. Xue, and N. Zheng, "View adaptive neural networks for high performance skeleton-based human action recognition," *IEEE transactions on pattern analysis and machine intelligence (TPAMI)*, vol. 41, no. 8, pp. 1963–1978, 2019.
- [11] H. Duan, Y. Zhao, K. Chen, D. Lin, and B. Dai, "Revisiting skeletonbased action recognition," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 2969–2978.
- [12] Y. Ben-Shabat, X. Yu, F. Saleh, D. Campbell, C. Rodriguez-Opazo, H. Li, and S. Gould, "The ikea asm dataset: Understanding people assembling furniture through actions, objects and pose," in *IEEE Winter Conf. on Applications of Computer Vision (WACV)*, 2021, pp. 847–859.
- [13] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft COCO: Common Objects in Context," in *European Conf. on Computer Vision (ECCV)*, 2014, pp. 740–755.
- [14] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask r-cnn," in *IEEE Int. Conf. on Computer Vision (ICCV)*, 2017, pp. 2961–2969.
- [15] R. Dai, S. Das, L. Minciullo, L. Garattoni, G. Francesca, and F. Bremond, "Pdan: Pyramid dilated attention network for action detection," in *IEEE Winter Conf. on Applications of Computer Vision (WACV)*, 2021, pp. 2970–2979.
- [16] R. Dai, S. Das, K. Kahatapitiya, M. S. Ryoo, and F. Brémond, "Mstct: Multi-scale temporal convtransformer for action detection," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 20041–20051.
- [17] L. L. Presti and M. La Cascia, "3d skeleton-based human action classification: A survey," *Pattern Recognition*, vol. 53, pp. 130–147, 2016.
- [18] L. Shi, Y. Zhang, J. Cheng, and H. Lu, "Two-stream adaptive graph convolutional networks for skeleton-based action recognition," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 12 026–12 035.
- [19] R. Liu, C. Xu, T. Zhang, W. Zhao, Z. Cui, and J. Yang, "Si-gcn: Structure-induced graph convolution network for skeleton-based action recognition," in *Int. Joint Conf. on Neural Networks (IJCNN)*. IEEE, 2019, pp. 1–8.
- [20] A. Shahroudy, J. Liu, T.-T. Ng, and G. Wang, "Ntu rgb+d: A large scale dataset for 3d human activity analysis," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016.

- [21] L. Xu, C. Lan, W. Zeng, and C. Lu, "Skeleton-based mutually assisted interacted object localization and human action recognition," *IEEE Transactions on Multimedia*, 2022.
- [22] C. R. Dreher, M. Wächter, and T. Asfour, "Learning object-action relations from bimanual human demonstration using graph networks," *IEEE Robotics and Automation Letters (RA-L)*, vol. 5, no. 1, pp. 187– 194, 2019.
- [23] H. Xing and D. Burschka, "Understanding spatio-temporal relations in human-object interaction using pyramid graph convolutional network," in *IEEE Int. Conf. on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 5195–5201.
- [24] M. Zakour, A. Mellouli, and R. Chaudhari, "Hoisim: Synthesizing realistic 3d human-object interaction data for human activity recognition," in *IEEE Int. Conf. on Robot & Human Interactive Communication (RO-MAN)*. IEEE, 2021, pp. 1124–1131.
- [25] F. Sener, D. Chatterjee, D. Shelepov, K. He, D. Singhania, R. Wang, and A. Yao, "Assembly101: A large-scale multi-view video dataset for understanding procedural activities," in *IEEE Conf. on Computer Vision* and Pattern Recognition (CVPR), 2022, pp. 21 096–21 106.
- [26] F. Ragusa, A. Furnari, S. Livatino, and G. M. Farinella, "The meccano dataset: Understanding human-object interactions from egocentric videos in an industrial-like domain," in *IEEE Winter Conf. on Applications of Computer Vision (WACV)*, 2021, pp. 1569–1578.
- [27] D. Aganian, B. Stephan, M. Eisenbach, C. Stretz, and H.-M. Gross, "ATTACH dataset: Annotated two-handed assembly actions for human action understanding," in *Int. Conf. on Robotics and Automation (ICRA)*. IEEE, 2023.
- [28] Z. Cao, T. Simon, S.-E. Wei, and Y. Sheikh, "Realtime multi-person 2d pose estimation using part affinity fields," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 7291–7299.
- [29] Y. Du, Y. Fu, and L. Wang, "Skeleton based action recognition with convolutional neural network," in *IEEE IAPR Asian Conf. on pattern recognition (ACPR)*, 2015, pp. 579–583.
- [30] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE Conf. on Computer Vision and Pattern Recognition* (CVPR), 2016, pp. 770–778.
- [31] C. Feichtenhofer, H. Fan, J. Malik, and K. He, "Slowfast networks for video recognition," in *IEEE Int. Conf. on Computer Vision (ICCV)*, 2019, pp. 6202–6211.
- [32] MMAction2 Contributors, "Openmmlab's next generation video understanding toolbox and benchmark," https://github.com/open-mmlab/ mmaction2, 2020.
- [33] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in Int. Conf. on Learning Representations (ICLR), 2015.
- [34] J. Carreira and A. Zisserman, "Quo vadis, action recognition? a new model and the kinetics dataset," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [35] J. Liu, A. Shahroudy, M. Perez, G. Wang, L.-Y. Duan, and A. C. Kot, "Ntu rgb+ d 120: A large-scale benchmark for 3d human activity understanding," *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 42, no. 10, pp. 2684–2701, 2019.
- [36] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He, "Aggregated residual transformations for deep neural networks," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 1492–1500.
- [37] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, "Swin transformer: Hierarchical vision transformer using shifted windows," in *IEEE Int. Conf. on Computer Vision (ICCV)*, 2021, pp. 10012–10022.
- [38] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in Int. Conf. on Learning Representations (ICLR), 2019.
- [39] J. Lu, A. Liu, F. Dong, F. Gu, J. Gama, and G. Zhang, "Learning under concept drift: A review," *IEEE Transactions on Knowledge and Data Engineering*, vol. 31, no. 12, pp. 2346–2363, 2019.
- [40] F. Bayram, B. S. Ahmed, and A. Kassler, "From concept drift to model degradation: An overview on performance-aware drift detectors," *Knowledge-Based Systems*, p. 108632, 2022.
- [41] D. Damen, H. Doughty, G. M. Farinella, A. Furnari, E. Kazakos, J. Ma, D. Moltisanti, J. Munro, T. Perrett, W. Price *et al.*, "Rescaling egocentric vision: collection, pipeline and challenges for epic-kitchens-100," *Int. Journal of Computer Vision (IJCV)*, vol. 130, no. 1, pp. 33–55, 2022.