# Delving Deep into One-Shot Skeleton-based Action Recognition with Diverse Occlusions

Kunyu Peng<sup>1</sup>, Alina Roitberg<sup>1</sup>, Kailun Yang<sup>1</sup>, Jiaming Zhang<sup>1</sup>, and Rainer Stiefelhagen<sup>1</sup>

*Abstract*—Occlusions are universal disruptions constantly present in the real world. Especially for sparse representations, such as human skeletons, a few occluded points might destroy the geometrical and temporal continuity critically affecting the results. Yet, the research of data-scarce recognition from skeleton sequences, such as one-shot action recognition, does not explicitly consider occlusions despite their everyday pervasiveness.

In this work, we explicitly tackle body occlusions for Skeletonbased One-shot Action Recognition (SOAR). We mainly consider two occlusion variants: 1) random occlusions and 2) more realistic occlusions caused by diverse everyday objects, which we generate by projecting the existing IKEA 3D furniture models into the camera coordinate system of the 3D skeletons with different geometric parameters, (e.g., rotation and displacement). We leverage the proposed pipeline to blend out portions of skeleton sequences of the three popular action recognition datasets (NTU-120, NTU-60 and Toyota Smart Home) and formalize the first benchmark for SOAR from partially occluded body poses. This is the first benchmark which considers occlusions for datascarce action recognition. Another key property of our benchmark are the more realistic occlusions generated by everyday objects, as even in standard recognition from 3D skeletons, only randomly missing joints were considered. We re-evaluate existing state-of-the-art frameworks for SOAR in the light of this new task and further introduce Trans4SOAR - a new transformer-based model which leverages three data streams and mixed attention fusion mechanism to alleviate the adverse effects caused by occlusions. While our experiments demonstrate a clear decline in accuracy with missing skeleton portions, this effect is smaller with Trans4SOAR, which outperforms other architectures on all datasets. Although we specifically focus on occlusions, Trans4SOAR additionally yields state-of-the-art in the standard SOAR without occlusion, surpassing the best published approach by 2.85% on NTU-120.

#### I. INTRODUCTION

**B**ENEFITING from the rapid progress of deep learning, conventional architectures for skeleton-based action recognition achieved remarkable results on a variety of benchmarks for body-pose based classification, *e.g.*, NTU-120 [2] and Toyota Smart Home [3]. However, most of the previous works regard relatively clean datasets as the default starting point for training models [2], [4], [5], [6]. Such datasets often explicitly ensure good body visibility through a suitable camera placement, but this assumption is rather naive, as in real-life the input is often disrupted through *occlusions*.



Fig. 1. An overview of the two proposed and reformulated occlusion scenarios, *i.e.*, REalistic synthesized occlusion (RE) proposed by us as depicted in (1) and RAndom occlusion (RA) reformulated by us depicted in (2). In order to generate realistic occlusion, we randomly choose a 3D furniture model from PIX3D dataset [1] and insert it into the world coordinate of the 3D pose using random translation and rotation. The occluded region of the skeletons are determined through camera view point projection, where the whole procedure is demonstrated in (3). In this work we investigate the influence brought by these two dominant occlusions for SOAR, as (4).

Skeleton-based action recognition algorithms operate on the coordinates of the 3D body joints and have attracted a great amount of attention [4], [5], [6], [7], [8], [9], [10], [11], [12] due to the increasing precision of depth sensors and their privacy-preserving characteristics, but occlusions are especially damaging for such sparse representations, where a few missing joints have a substantial impact on the geometric and temporal continuity.

Learning new concepts with only few labelled examples is often posed in the form of one- or few-shot recognition problem [13], [14], [15], [16], [17] and still remains one of the key problems in human action recognition. Especially if only few training examples are available, occlusions are critical since no diverse data is available for new categories and the quality of the few provided samples dominate the final results.

In this paper, we are interested in categorizing sequences of unseen 3D body poses from only one reference sample, where portions of these sequences are missing due to occlusions.

The research leading to these results was supported by the SmartAge project sponsored by the Carl Zeiss Stiftung (P2019-01-003; 2021-2026). The authors would like to thank the consortium for the successful cooperation. (*Corresponding author: Kailun Yang.*)

<sup>&</sup>lt;sup>1</sup>Authors are with Institute for Anthropomatics and Robotics, Karlsruhe Institute of Technology, Germany. (E-mail: {kunyu.peng, alina.roitberg, kailun.yang, jiaming.zhang, rainer.stiefelhagen}@kit.edu).

Code will be made publicly available at Trans4SOAR.

Since no past work on one-shot action recognition from skeleton data explicitly considers occlusions, we first introduce a new benchmark by blending out skeleton parts in three established action recognition datasets. Our idea is to use a library [1] of 3D objects to generate REalistic occlusions (RE), which we project onto the original data with different geometric parameters, such as rotation and displacement. We believe, that by projecting everyday objects (as shown in Figure 1), the occlusions are far more realistic than random dropping of data points, which has been considered in standard [18], [19], (*i.e.*, not data scarce) recognition from 3D skeletons in the past while preserving the dimension of skeleton joints compared with [20]. Still, we also consider random blending of body joints while considering both spatial and temporal information as a second occlusion variant, *i.e.*, RAndom occlusion (RA), in our benchmark.

We also introduce Trans4SOAR - a new transformer-based model which comprises three data streams and mixed fusion to overcome the challenges caused by occlusions. Until now, Skeleton-based One-shot Action Recognition (SOAR) has been mostly addressed with Convolutional Neural Networks (CNNs) coupled with metric learning [21], [22], [23] or meta learning [24]. While few recent works considered transformers networks in conventional video-based human activity classification [25], their potential as signal encoders of body movement, their transfer capabilities to new data-scarce activities *classes* and their ability to deal with body pose occlusions have been overlooked. For example, Skeleton-DML and SL-DML, the state-of-the-art approaches for the SOAR task [21], [22], both leverage a CNN-based encoder for signal-level skeleton representation and the deep metric learning paradigm. Our experimental strategy is therefore to first re-evaluate Skeleton-DML [21] and SL-DML [22] as the current state-of-theart approaches in the light of our new occlusion-centered task. Then, we revisit image-like modelling of skeleton dynamics with the rapidly emerging visual transformers within our Trans4SOAR model. Apart from being the first visual transformer-based architecture for encoding skeleton signals as image-like representations targeting at the SOAR task, we alleviate the adverse effects caused by occlusions through two novel design choices. First, we leverage complementary types of information (velocities, bones and joints) and propose the Mixed Attention Fusion Mechanism (MAFM) which learns how to link the information from diverse streams at the patch embedding level while considering the spatial and temporal neighborhood information. Secondly, we leverage the Latent Space Consistency (LSC) loss encouraging the model to output similar results with an additional auxiliary branch, if the embedding in the middle layer of the auxiliary branch has been altered by category agnostic prototypes, which suits naturally to the use-case of disturbances through occlusions.

This paper explicitly explores occlusions for SOAR and makes the following contributions:

• We for the first time tackle occlusions for Skeletonbased One-Shot Action Recognition (SOAR) and build a benchmark for this task by augmenting three established datasets for action recognition through our occlusion pipeline. Our pipeline considers different occlusion settings, where RAndom occlusion (RA) and REalistically synthesised occlusion (RE) based on everyday objects are the most important ones. We view the latter case as a more practical scenario closer to real-life applications and achieve it by using a library of 3D objects obtained from the IKEA 3D furniture dataset [1], which are inserted into the world coordinate system of the body poses with different geometric parameters.

- We introduce TRANS4SOAR an new three-stream transformer-based model specifically aimed at overcoming data occlusions by 1) leveraging diverse types of input (velocities, bones and joints) and enabling information exchange at the patch embedding level via the Mixed Attention Fusion Mechanism (MAFM) and 2) extensively augmenting the intermediate transformer representations through iteratively estimated category-specific prototypes and the Latent Space Consistency (LSC) loss.
- We conduct in-depth experiments in the SOAR task, evaluating two existing state-of-the art frameworks [22], [21] as well as our TRANS4SOAR network and its individual building blocks under four different types of occlusions. Unsurprisingly, introducing occlusions adversely impacts the outcome, marking an important future work direction. Our TRANS4SOAR model yields state-of-the-art on all three datasets under the presence of occlusions.
- As a side-observation, we discover that TRANS4SOAR also outperforms state-of-the-art in the standard SOAR, (*i.e.*, the not occluded SOAR task), surpassing the best previously published model on the challenging NTU-120 SOAR benchmark by > 2.8%.

## II. RELATED WORK

# A. One-shot Action Recognition

One-shot recognition, belonged to data-scarce representation learning paradigm, aims at the recognition of unseen category with only one reference samples given as guidance. Compared with existing works in one-shot image classification, where meta learning-based approaches [24], [27], [28], [29], [30], [31] dominate the important positions by reinitializing a new task set every epoch following the paradigm regarding learning to learn, Deep Metric Learning (DML) based approaches [21], [22], [23], which aims at achieving highly discriminative representation and closer representation distance for inter- and intra-category samples in the latent space, are well utilized for Skeleton-based One-shot Action Recognition (SOAR) benchmarked by NTU-120 [2] with predefined reference frames. One-shot action recognition has been well studied for several down-stream tasks, e.g., semantic segmentation [32] and video classification [13], [15], [16], [33], however the research of SOAR are much more sparse and mostly benchmarked on the NTU-120 dataset [2], [21], [34], [35], [36]. State-of-the-art recognition results are currently reached by the approaches of Memmesheimer et al. [22], [21], which use a CNN-based encoding of 3D skeletons represented as images and optimizes the framework with deep metric learning using a mixture of cross entropy and triplet margin losses. In this work, we investigate the transformer architecture



Fig. 2. An overview of the proposed TRANS4SOAR architecture, which is a **Trans**former for Skeleton-based **One-Shot Action Recognition**. (a) indicates the transformer block leveraged in TRANS4SOAR. This basic transformer attention block is proposed by LeViT [26], which builds up the transformer block in the later stage of our TRANS4SOAR architecture through stacking. (b) is the overview of the TRANS4SOAR training pipeline. First, the skeleton signals are encoded in three kinds of format, *i.e.*, joints, bones, and velocities. Image-like representations are formulated through the concatenation along the temporal axis of the skeleton data, which are further divided into several patches and fed into its corresponding patch embedding net. Then, the Mixed Attention Fusion Mechanism (MAFM) fuses the embeddings from these three different streams by using Mixed Fusion (MF) to achieve cross-stream aggregation on Key, Query, and Value together with the proposed Softmax Concentrated Aggregation (SCA). The Latent Space Consistency (LSC) loss  $L_{LSC}$  integrates an prototype augmented auxiliary branch and adopts cosine similarity loss to encourage the embeddings from the main branch **E** and the embedding from the auxiliary branch net adopts cosine similarity loss ( $L_{TPL}$ ), cross entropy loss ( $L_{CLS}$ ), and LSC loss ( $L_{LSC}$ ), are leveraged for discriminative representation learning. EMB indicates embedding generation layers, which are built based on multi-layer perceptrons (MLP). Head indicates a fully-connected (fc) layer based classification head. PE indicates the patch embedding network. (c) shows the workflow of the Mixed Fusion (MF) and (d) shows the Mixed Attention Fusion Mechanism (MAFM), where **Proj** indicates the fc-based projection layer, AVG indicates the average operation and LN indicates layer normalization.

and propose a new model named as TRANS4SOAR for the SOAR task while considering different occlusion scenarios. We build our optimization paradigm based on SL-DML [21] while further proposing a novel patch embedding level fusion approach considering different skeleton encoding formats.

#### B. Visual Transformers

Transformer networks [37] are rapidly gaining popularity in computer vision since their operationalization on image patched within the ViT [38] and DeiT [39] architectures. Recently, transformer models, known for capturing essential long-range context [37], become increasingly appealing in vision tasks since ViT [38] and DeiT [39] directly utilize a pure transformer on image patches. A large amount of transformerbased models are thereby put forward regarding this new trend, while some of them target at pursuing better accuracy on image classification task [40], [41], [42], resource-efficiency [26], [43] and the others are designed for more specific tasks, *e.g.*, semantic segmentation [32]. In action recognition task, the transformer-based models could be arranged into two main groups, which are video-based transformer [25], [44], [45], [46] and skeleton-based transformer [7], [9], [10], [11], [47], [4] using standard sequential skeleton as input. We for the first time investigate visual transformer for skeleton data by encoding the skeleton as image-like representation which has the same encoding procedure with SL-DML [22] while using an additional auxiliary branch and loss to achieve latent space consistency. Furthermore, in order to achieve high robustness against different occlusion scenarios, a novel feature extraction architecture TRANS4SOAR is proposed by integrating a Mixed Attention Fusion Mechanism (MAFM).

# C. Skeleton-based Action Recognition with Occlusion

Most of the skeleton extraction approaches, *e.g.*, Alpha-Pose [48], [49], [50], tend to directly give zero output regarding the occluded human body joints. Thereby some researchers formulated the occluded action recognition task through randomly setting different body regions per frame as zeros to simulate spatial occlusion or setting the randomly selected frame as zeros to simulate temporal occlusion [18], [19], [51], [52], [53], [54] while self-occlusion caused by human body movement is considered in [55]. Notice that, all the aforementioned related works are for skeleton-based action recognition, which is not for SOAR, tackled by our work. In this work, we jointly consider both spatial occlusion and temporal occlusion at the same time to form random

occlusion which is a more reasonable randomly generated occlusion setting since temporal and spatial occlusion should be considered together if the whole skeleton sequence is seen as a sample for SOAR. Besides, for realistic synthesized occlusion, OSD dataset [20] for the first time projected 3D real world objects into image plane and then generated occluded skeleton in 2D image coordinate. The dimension of raw data is unfortunately not preserved which results in massive information loss by converting 3D data into 2D data. In order to tackle the occlusion issue in more realistic scenario while preserving dimension of the data, we thereby propose a dimension preserving realistic synthesized occlusion pipeline using an additional 3D model dataset PIX3D [1]. Moreover, we for the first time investigate SOAR under diverse occlusions while all the existing works target not-occluded one-shot action recognition.

#### III. BENCHMARK

To address the lack of related benchmarks, we collect and publicly release the first testbed for SOAR *under presence of occlusions*. Our benchmark augments three prominent datasets for SOAR with different occlusions, of which random spatiotemporal occlusions and the realistically synthesised occlusions derived from an existing 3D library of furniture objects being the most important ones. Next, we give a formal definition of the addressed task (Section III-A) and describe the data obstruction mechanism that we have developed to reach our design goal of realistic everyday occlusions (Section III-B) as well as the more conventional random occlusion pipeline (Section III-C).

# A. Problem Definition

The task we address is SOAR [21] where a priori knowledge acquired from data-rich action classes is transferred to categorize new data-scarce classes, while certain regions of the skeleton are not visible. Following the standard evaluation protocol for data-scarce action recognition, we build on the one-shot evaluation setting of NTU-120 [2], where new categories of the body pose sequences are categorized from a single reference sample. Formally, Cbase denotes the set of  $|C_{base}|$  data-rich categories available during training through large amount of labelled data  $D_{base} = \{(\mathbf{S}_i, l_i)\}_{i=1}^U, l_i \in C_{base}$ while U indicates the number of samples in  $D_{base}$ . Our goal is to distinguish the  $|C_{novel}|$  new activity classes  $C_{novel}$ , for which only  $\kappa = 1$  reference training example is available for each class. These data-scarce examples are referred to as support set  $D_{supp} = {\mathbf{S}_i}_{i=1}^O$  while *O* indicates the number of samples in  $D_{supp}$  and  $C_{base} \cap C_{novel} = \emptyset$ . The final task is then to assign a category  $l_n \in C_{novel}$  to each sample from the test set  $D_{test}$ containing examples from the data-scarce categories  $C_{novel}$ .

Since our idea is to study and address *occlusions* for skeleton-based recognition with little training data, we augment both, the support set  $D_{supp}$  and the test set  $D_{test}$  with our pipeline described in the upcoming sections. Note, that in our experiments, we consider both: 1) obstructing the reference examples from the support set and the test set examples and 2) considering occlusions in the test set only, while using complete sequences as our reference samples.



Fig. 3. An overview of Signal-to-Noise Ratio (SNR) distribution of the realistic synthesized occlusion dataset, where (a), (b) and (c) are for the NTU-120 [2], the NTU-60 [56] and the Toyota Smart Home [3] respectively. The legend indicates the corresponding SNR range.

# B. Realistic Synthesized Occlusion

Past work on skeleton-based activity recognition (without the data scarcity constraint) considered occlusions as random temporal or spatial obstructions only [18], achieved by randomly setting a fixed number of frames or a fixed number of joints to zero respectively. Such occlusions have highly controllable characteristics but are rather unusual in the real world, where objects are a common cause for obstructions and missing skeleton points exhibit specific geometric consistencies. To tackle this issue, we build new occluded versions of three public datasets, i.e., NTU-120 [2], NTU-60 [56] and Toyota Smart Home [3], by inserting the 3D IKEA furniture models obtained from the PIX3D dataset [1] into the world coordinates system of the human body. Note, that while NTU-120 [2] and NTU-60 [56] contain 3D skeletons, while Toyota Smart Home [3] covers 2D data, the process of augmenting the data with the realistic synthesized occlusions is different and will be explained in the following sections. The statistics of Signal-to-Noise Ratio (SNR) for each dataset with the proposed realistic synthesized occlusion is depicted in Figure 3.

3D realistic synthesized occlusion dataset generation. The NTU-120 [2] and NTU-60 [56] datasets cover different camera views. We therefore need to consider the cross-view consistency when obstructing the body poses with the furniture models. Unfortunately, the datasets do not provide the calibration data of the individual cameras, which would be the first essential piece of information when addressing this problem. Luckily, our goal is the skeleton data and each frame containing a skeleton provides sufficient world coordinate information about joints, which can be directly used to calibrate the relative position of the different cameras. A single skeleton sequence sample with T frames and J joints gives us a set of coordinates with number of  $T \cdot J$ . If N samples are provided in our dataset, the known number of the world coordinates thereby rises to N \* T \* J which is much higher than the rank of projection matrix between two different cameras. The calibration matrix  $\mathbf{F}_{ii}$  between two utilized cameras *i* and *j* can be estimated through the following equation:

$$\mathbf{F}_{ij} = (\mathbf{X}_i^T \mathbf{X})_i^{-1} \mathbf{X}_i^T \mathbf{X}_j'$$
(1)

where  $\mathbf{X}$  denotes the collection of human body joints with homogeneous coordinate format captured simultaneously by two different cameras. Thereby, the projection matrix between every two cameras can be obtained as  $\mathbf{F}_{ii} \in \mathbf{F}$ .

A detailed description of the 3D occlusion generation procedure is formalized in Alg. 1. First, we randomly select 3D object model from the existing IKEA furniture library [1] containing 395 different models from 9 categories and augment it via random rotation and translation regarding the vertical axis (bottom to up) and the horizontal plane respectively to simulate life-like occlusions while trying to ensure the bottom points of the skeleton and furniture are in the same vertical height level through vertical translation. To ensure the cross-view consistency, samples collected simultaneously by different cameras share the rotation and translation augmented furniture model by using the projection matrix from the calibration set F. Next, we need to determine which skeleton body joints are blended out by the current occluded object from the perspective of the camera. The skeleton body joints  $\mathbf{x} = [x_1, x_2, x_3]$  (where  $x_3$ indicates the axis which is parallel to the focus axis of the camera) and the points of the augmented furniture model  $\mathbf{z}$ are first horizontally projected along the focus axis of camera into  $\mathbf{x}^*$  and  $\mathbf{z}^*$ , as  $\mathbf{x}^* = [x_1^*/x_3^*, x_2^*/x_3^*]$  and  $\mathbf{z}^* = [z_1^*/z_3^*, z_2^*/z_3^*]$ .

Then, we build up a two-dimensional convex hull based on the projected points regarding  $\mathbf{z}^{\star}$  and use the following equation to determine if the human body joints fall into the occlusion convex hull from the camera point of view or not:

$$IsInHull = IsTrue(A \times (\mathbf{P}^{I}) \le Tile(-\mathbf{b}, (1, len(\mathbf{P}))), 0), (2)$$

where A denotes the construction equation of the 2D convex hull of  $\mathbf{z}^*$ , **b** denotes the last convex hull equation, **P** denotes the point that needs to be determined, Tile indicates whether all array elements along a given axis are able to be evaluated as True or not and  $len(\cdot)$  is the coordinate point dimension. Thereby the IsInHull is a binary indicator, while True marks the point being inside the convex hull and vice versa. Finally we will get a *Mask* for each queried skeleton data (see Alg. 1), and the occluded 3D points are set to zero.

2D realistic synthesized occlusion dataset generation. Since the third dataset we leverage, *i.e.*, Toyota Smart Home [3], only contains 2D skeletons in the image plane, the aforementioned pipeline for generating 3D realistic synthesized occlusion is modified to suit the 2D use-case. Instead of directly augmenting the 3D furniture model via rotation and translation, we use a randomly generated projection matrix, which transforms the points from the camera coordinates to the image coordinate system, to project the furniture model onto the image plane. Then, similar to the 3D realistic synthesized occlusion generation procedure, the convex hull of the occluded region is built up according to the projected points of the furniture model on the 2D image plane and an occlusion-aware mask is obtained through the ISInHull function. Finally, the corresponding 2D skeleton joints within the mask are changed to zeros.

# C. Random Occlusion

The second occlusion variant we considered is random data obstruction, which is a combination of random temporal and spatial occlusions used in past work on standard, (i.e., without the data scarcity constraint) skeleton-based action

#### Algorithm 1 3D realistic synthesized occlusion generation.

```
Input: F – the set of projection matrix for each camera pair; X –
   the set of skeleton data; Z - the collection of 3D furniture models
   from PIX3D dataset; R and T – random rotation and translation
   augmentations; X_{Occ} – a empty set for occluded skeleton data;
   [a, b] – predefined occluded SNR range for the acceptance,
   where a is the lower limitation and b is the upper limitation.
1: for all \mathbf{x} \in X: do
```

Set Accept = False. 2:

- 3: while Accept! = True do
- Set Found = False. 4:

```
5:
          Set N_{Occ} = 0.
```

{ $N_{Occ}$  is the occluded sample number for  $X_{d+1}$ .} 6:

```
7:
       while Found! = True do
```

- 8: Search  $\mathbf{X}_d$  collected simultaneously with  $\mathbf{x}$  from different views
- Extract the calibration set  $F_d$  for  $X_d$ . 9:
- 10: Randomly select  $\mathbf{z}$ , where  $\mathbf{z} \in \mathbf{Z}$ .
- Obtain augmented **z**, i.e.,  $\mathbf{z}'$ , by  $\mathbf{z}' = R(T(\mathbf{z}))$ . 11:
- Get  $Z_d$  by applying  $\mathbf{f}_d \in F_d$  on  $\mathbf{z}$ . 12:
- 13: Define  $Z_{d+1} = Z_d \cup \{z\}$  and  $X_{d+1} = X_d \cup \{x\}$
- if  $Z_{d+1}$  has no intersection with  $X_{d+1}$  for each corre-14: sponding element: then

15: Found = True

- end if 16:
- 17: end while
- 18: for  $(\mathbf{x}_d, \mathbf{z}_d) \in zip(X_{d+1}, Z_{d+1})$  do
- 19: Horizontally project  $\mathbf{z}_d$  and  $\mathbf{x}_d$  along focus axis of camera d into 2D plane as  $\mathbf{z}_d^*$  and  $\mathbf{x}_d^*$ . 20:

```
Build up 2D convex hull \Phi of \mathbf{z}_d^{\star}.
```

```
Mask_d = IsInHull(\Phi, \mathbf{x}_d^{\star}).
21:
               Calculate SNR_d = Sum(Mask_d)/len(Mask_d) for \mathbf{x}_d.
22:
               Occlude \mathbf{x}_d by \mathbf{x}_d[Mask_d] = zeros\_like(\mathbf{x}_d[Mask_d])
23:
              Append \mathbf{x}_d into X_d^{Occ}
24:
              if SNR_d in [a,b] then
25:
26:
                  N_{Occ} + = 1.
27:
               end if
           end for
28:
29:
           if N_{Occ} < T_{Occ} or N_{Rep} < T_{Rep} then
              Set Accept = False and N_{Rep} + = 1.
30:
31:
           else
```

```
Set Accept = True.
```

Del  $X_d$  from X and append the  $X_d^{Occ}$  into  $X_{Occ}$ .

end if

35: end while

36: end for

32:

33:

34:

recognition [18], [19]. For random temporal occlusions, a fixed number of frames are blended out randomly for each skeleton sequence to simulate full occlusion for certain points in time. For random spatial occlusions, a fixed number of joints are randomly set to zero in every frame of the skeleton data stream. This is a very specialized type of occlusions, since the exact number of joints are not visible in all frames. However, mixing both, random temporal occlusions and random spatial occlusions, is a more reasonable scenario with less predefined controllable conditions and a higher chance to happen in real-life. With  $\mathbf{x} \in \mathbb{R}^{T \times J \times B}$  being the skeleton data, we first flatten it into  $\mathbb{R}^{(T \times J) \times B}$ , after which a set of data points  $\gamma \cdot (T \times J)$  are randomly chosen with a predefined SNR ratio  $\gamma$  and blended out. Although we view the mixed spatial and temporal occlusion as a more reasonable option, we also conduct experiments with isolated random spatial and temporal occlusions for consistency. Overall, our experiments

described in the later sections will indicate, that the realistic synthesized occlusions are the most challenging ones.

# IV. METHODS: TRANS4SOAR

We introduce *Trans4SOAR* – a three-stream transformerbased model designed to overcome adverse effects of occlusions (an overview is provided in Fig. 2). The key ingredients of our model are 1) the Mixed Attention Fusion Mechanism (MAFM) which learns to aggregate three different types of skeleton information, (*i.e.*, joints, velocities and bones) at the patch embedding level and 2) an extensive data augmentation technique at the feature-level, where an auxiliary branch is augmented by category agnostic prototypes. The motivation of the latter component is to encourage better robustness against imperfect data brought by occlusions through an additional consistency cost computed between the obtained body pose embedding and its prototype-augmented version.

Next, we describe the basic components regarding the input encoding, the patch embedding procedure and the basic transformer blocks of TRANS4SOAR (Sec. IV-A) and MAFM – the central building block model responsible for the three-stream fusion at the patch embedding level (Sec. IV-B). Finally we introduce our auxiliary Latent Space Consistency (LSC) loss for encouraging invariance to transformations through consistency constraints and augmentations with previously learned action category prototypes (Sec. IV-C).

#### A. Illustration of the Base Components

Input encoding. We follow the skeleton encoding for body joints proposed by SL-DML [22] to cast the sequential skeleton data as image-like representations. Assuming that  $\mathbf{s} \in \mathbb{R}^{T \times J \times B}$  denotes a sequential skeleton sample, where T indicates the temporal length, J indicates the total number of joints, and B indicates the dimension of the coordinates of the skeleton joints. The input of our approach is achieved by interpolation, which transfers s from  $T \times J \times B$  to  $H \times W \times B$  to match the image-wise input. Moreover, TRANS4SOAR is a threestream architecture, which does not only consider joints (as in [22]) but also bones **b** and velocities **v** defined as  $\mathbf{v}_t = \mathbf{s}_t - \mathbf{s}_{t-1}$ , denoting the velocity for each joint during human body motion at time stamp t, and  $\mathbf{b}_{i,i} = \mathbf{s}_i - \mathbf{s}_i$ , for  $(i, j) \in \Omega_{bones}$ , denoting the vector of each bone of the human skeleton, respectively. These vectors are subsequently mapped to image-like arrays using the described above procedure.

Finally, after the skeleton format encoding and image-wise transformation, we obtain three individually image-wise inputs including joints, velocities, and bones which have the same shape as  $H \times W \times B$ . So at the end of input encoding we have three streams of input, *i.e.*, joints, velocities and bones.

**Patch embedding and transformer blocks.** Modern CNNs are excellent at preserving details, while transformers are known for capturing long-range dependencies, making the combination of CNN- and transformer blocks a potential double win. LeViT [26] opened the door for this combination by using a CNN with four layers for patch embedding, before the stack of transformer blocks. Standing on the shoulders of giants, we adopt the basic transformer blocks and patch

embedding layers proposed by LeViT [26] in TRANS4SOAR. The basic transformer attention block of LeViT [26] is depicted on the left hand side of Figure 2. After the acquisition of the Query Q, Key K and Value V through the projection layers  $\mathbf{Proj}_{O}$ ,  $\mathbf{Proj}_{K}$ , and  $\mathbf{Proj}_{V}$  respectively, the final attended output can be calculated as  $\operatorname{Proj}_{Top}(HardWish((Softmax(\mathbf{Q} \times \mathbf{K})) +$  $Bias_{att}$  × V), where each projection layer is composed of an 1x1 conv and a batch normalization layer, and Bias<sub>att</sub> denotes the attention bias. First, the leveraged three streams of inputs namely joints, velocities, and bones are separately divided into  $N_{patch} = (H/P) \times (W/P)$  patches with a predefined path size P and thereby three input sequences are produced, which are denoted as  $S_i$ ,  $S_v$ , and  $S_b$  for joints, velocities, and bones, respectively. Then, we build up patch embedding layers through a stack of convolutional neural networks, denoted as  $M_{\theta^j}$ ,  $M_{\theta^v}$ , and  $M_{\theta^b}$  to extract high-dimensional embeddings for the patch sequence of each stream, denoted separately as  $\mathbf{E}_{patch,j}$ ,  $\mathbf{E}_{patch,v}$ , and  $\mathbf{E}_{patch,b}$ . We follow the attention bias setting (depicted in Figure 2 (a)) instead of using position embeddings as introduced in [26]. The corresponding relationship is depicted in Eq. (3):

$$\mathbf{E}_{patch,j}, \, \mathbf{E}_{patch,v}, \, \mathbf{E}_{patch,b} = M_{\theta^j}(\mathbf{S}_j), \, M_{\theta^v}(\mathbf{S}_v), \, M_{\theta^b}(\mathbf{S}_b).$$
 (3)

The resulted three embedding streams are then fed into the key component of the proposed TRANS4SOAR architecture, *i.e.* Mixed Attention Fusion Mechanism (MAFM), for multi-stream fusion, which we now introduce in detail.

#### B. Multimodal Fusion at the Patch Embedding Level

Mixed Fusion (MF). The proposed Mixed Attention Fusion Mechanism (MAMF) builds upon the the Mixed Fusion (MF) strategy. The main purpose of MF is to transfer the important cues from the two auxiliary streams, *i.e.*, velocities and bones, to the main stream, *i.e.*, joints, to achieve multi-stream fusion of skeleton data on the patch embedding level. We propose to use a mixture of Value and Key for multi-stream fusion. While such a concept regarding the mixture of Key and Value is studied in MixFormer [57] for template matching, the design of our proposed MF mechanism is non-trivial. Unlike the mixture in template matching which aims to push the model to focus on similarity cues, our MF scheme is designed to harvest complementary cross-modality dependencies and enable a multi-stream agreement for discriminative embedding learning. Apart from using a concept of mixture of Key and Value, we design a unique three-stream patchembedding fusion architecture regarding MF and MAFM to suit the discriminative embedding learning the for SOAR. In the following, we introduce the proposed MF for multi-stream patch embedding fusion in detail. First, we encode the patch embeddings of the joints  $E_{patch,j}$  through two different linear projection layers, *i.e.*,  $\mathbf{Proj}_{O}^{jv}$  and  $\mathbf{Proj}_{O}^{bj}$ , as depicted in Eq. (4):

$$\mathbf{Q}_{jv}, \ \mathbf{Q}_{bj} = \mathbf{Proj}_Q^{jv}(\mathbf{E}_{patch,j}), \ \mathbf{Proj}_Q^{bj}(\mathbf{E}_{patch,j}).$$
(4)

Then, for Keys and Values of the jv branch and the bj branch, the input embeddings are aggregated together through concatenation, which is indicated as *Cat*. After that, for each single term, a projection layer is used for encoding. For

example,  $\mathbf{Proj}_{V}^{jv}$  is the projection layer for Value of the jvbranch. As a result,  $\mathbf{V}_{iv}$ ,  $\mathbf{K}_{jv}$ ,  $\mathbf{V}_{bj}$ , and  $\mathbf{K}_{bj}$  are yielded after the encoding:

$$\mathbf{V}_{iv} = \mathbf{Proj}_{V}^{jv}(Cat(\mathbf{E}_{patch,j}, \mathbf{E}_{patch,v})),$$
(5)

$$\mathbf{K}_{jv} = \mathbf{Proj}_{K}^{jv} (Cat(\mathbf{E}_{patch,j}, \mathbf{E}_{patch,v})),$$
(6)

$$\mathbf{V}_{bj} = \mathbf{Proj}_{V}^{bj}(Cat(\mathbf{E}_{patch,j}, \mathbf{E}_{patch,b})),$$
(7)

$$\mathbf{K}_{bj} = \mathbf{Proj}_{K}^{bj}(Cat(\mathbf{E}_{patch,j}, \mathbf{E}_{patch,b})).$$
(8)

After the aforementioned procedures, we have obtained Query, Key, and Value for the two branches, separately. Then the question lies in how to further aggregate these two branches. We introduce Softmax Concentrated Aggregation (SCA), which is realized through the following equations to achieve aggregation between  $\mathbf{V}_{jv}$  and  $\mathbf{V}_{bj}$ ,  $\mathbf{K}_{jv}$  and  $\mathbf{K}_{bj}$ , and  $\mathbf{Q}_{jv}$  and  $\mathbf{Q}_{bj}$ :

$$\mathbf{V} = (Softmax(\mathbf{V}_{j\nu})^T \mathbf{V}_{bj} + Softmax(\mathbf{V}_{bj})^T \mathbf{V}_{j\nu})/2, \qquad (9)$$

$$\mathbf{K} = (Softmax(\mathbf{K}_{jv})^T \mathbf{K}_{bj} + Softmax(\mathbf{K}_{bj})^T \mathbf{K}_{jv})/2, \quad (10)$$

$$\mathbf{Q} = (Softmax(\mathbf{Q}_{jv})^T \mathbf{Q}_{bj} + Softmax(\mathbf{Q}_{bj})^T \mathbf{Q}_{jv})/2.$$
(11)

After the SCA operation, we obtain the aggregated Query  $\mathbf{Q}$ , Key K, and Value V, which is merged together to formulate the desired attention by  $Att = Softmax(\mathbf{QK}^T/\sqrt{s_k})\mathbf{V}$ , where the scale factor is used to avoid the negative influence brought by the dot product on the variance and Att denotes the calculated attention value.

Mixed Attention Fusion Mechanism (MAFM). The MAFM is depicted on the upper right corner of Figure 2, which is designed for a further step of aggregation, while considering layer normalization (LN), averaged skip connection, and path drop out. First, the attended embedding  $\mathbf{E}_{att}$  is obtained through Eq. (12):

$$\mathbf{E}_{att} = MF(LN(\mathbf{E}_{patch,j}), LN(\mathbf{E}_{patch,v}), LN(\mathbf{E}_{patch,j})).$$
(12)

As shown in Figure 2, the original patch embeddings  $\mathbf{E}_{patch,j}$ ,  $\mathbf{E}_{patch,v}$ , and  $\mathbf{E}_{patch,b}$  are firstly averaged and then added with the path-dropped attended embedding  $\mathbf{E}_{att}$  to have  $\mathbf{E}_{asn}$ , an embedding after averaging (AVG) and applying an skip connection, as depicted in Eq. (13):

$$\mathbf{E}_{asn} = AVG(\mathbf{E}_{patch,j}, \mathbf{E}_{patch,v}, \mathbf{E}_{patch,b}) + DP(\mathbf{E}_{att}), \quad (13)$$

where DP indicates the drop path operation. Then, the final mixed resulted embedding  $\mathbf{E}_{mixed}$  is obtained via Eq. (14):

$$\mathbf{E}_{mixed} = DP(MLP(LN(\mathbf{E}_{asn}))) + \mathbf{E}_{asn}.$$
 (14)

Finally, the resulted mixed embedding is further fed into the stack of transformer blocks.

Algorithm 2 An overview of the training pipeline with latent space consistency (LSC) loss.

**Input:** S – a batch in  $D_{train}$ ; S<sub>p</sub> and S<sub>n</sub> – positive and negative anchor;  $f_{\delta}^{1 \to N-1}$  and  $f_{\theta}^{1 \to N-1}$  – first N-1 transformer layers of main and auxiliary branches;  $f_{\delta}^N$  and  $f_{\theta}^N$  – the N-th (last) transformer layer for main and auxiliary branches; *EMB* – Embedding layer;  $N_e$  – maximum training epochs;  $N_t$  – epoch threshold for the stage changing; E and  $E^*$  – embedding for main and auxiliary branches; PMB - Prototypes memory bank; WarmUpAug and PrototypeAug - Warm-up stage and prototypebased feature augmentation stage

1: for all  $epoch \in Range(N_e)$  do

2: for all  $S \in D_{train}$  do 3:

if  $epoch > N_t$  then

for all l in label<sub>S</sub> do Append(PMB[l])  $\rightarrow$  List<sub>p</sub> end for

 $E_P^* = Concat(List_p)$ 

endif

4: 5:

6:

7:

8:

9.

if BaseModel is not None then S = BaseModel(S)end if

10:  $\mathbf{E}_{patch} = PatchEmbeddingAndEncoding(\mathbf{S})$ 
$$\begin{split} \mathbf{E}_{parcn} &= \mathbf{I} \text{ acressing charge for a local mg (S)} \\ \mathbf{E}_{N-1} &= f_{\delta}^{1 \to N-1}(\mathbf{E}_{parch}), \ \mathbf{E}_{N-1}^* = f_{\theta}^{1 \to N-1}(\mathbf{E}_{parch}) \\ \text{if } epoch &< N_t \text{ then } \mathbf{E}_{aug}^* = WarmUpAug(\mathbf{E}_{N-1}^*) \\ \text{else } \mathbf{E}_{aug}^* = PrototypeAug(\mathbf{E}_{N-1}^*, \mathbf{E}_{P}^*) \\ \text{ond if } \end{split}$$
11: 12: 13: end if 14: 15:  $\mathbf{E} = EMB(f^N_{\delta}(\mathbf{E}_{N-1})), \ \mathbf{E}^* = EMB(f^N_{\theta}(\mathbf{E}^*_{aug}))$  $L_{tpl} = TripletMarginLoss(\mathbf{E}, \mathbf{E}_n, \mathbf{E}_p)$ 16:  $L_{LSC} = ConsistencyLoss(\mathbf{E}, \mathbf{E}^*) \rightarrow LSC \ loss$ 17:  $L_{class} = ClassificationLoss(Head(\mathbf{E}), label_{S})$ 18: 19:  $BackPropagation(WeightedSum(L_{tpl}, L_{class}, L_{LSC}))$ 20: end for if  $epoch < N_t - 1$  then 21: 22:  $CalculatePrototypes(D_{train}) \rightarrow Set(\mathbf{E}_{N-1}) \rightarrow PMB$ 23: end if 24: end for

# C. Prototype-based Latent Space Consistency Loss

To learn data-efficient one-shot action recognition, we put forward a Latent Space Consistency (LSC) loss, which encourages the consistency of the embeddings learned through the main branch and an additional prototype-based feature augmentation branch by cosine similarity loss, as illustrated in Alg. 2. The intention behind the design of LSC loss is to increase the robustness of the model by forcing the model to learn consistent embeddings even under the disruption of the feature-level augmentation, whose capability against occlusion is verified through our experiments. We build on top of a recent feature augmentation approach from semisupervised learning [58], but additionally propose a warmup self-augmentation phase and certain architecture changes, which have proven to be effective in improving both the accuracy and the robustness of the model.

**Estimating action category prototypes.** For the auxiliary branch augmentations at feature-level, we draw inspiration from FeatMatch [58], a recent method for semi-supervised image classification, where a learned weighted combined category-specific prototypes is used to enhance the intermediate features when referring to feature-level augmentations. Specifically, for each data-rich action category  $l_i \in C_{base}$ , we iteratively estimate its prototype in the latent space as the center of all training set embeddings of the specific action (we use the embeddings after the N-1 block if N is our number of transformer blocks). Note, that unlike FeatMatch, we use



Fig. 4. Overview of the self-augmentation at feature-level leveraged in the auxiliary branch of the proposed method. During the warm-up phase (top), the feature itself is used as the basis to compute attention masks used to self-augment the feature. At the later stage, we use action-specific prototypes Softmax-normalized along the channel dimension in order to augment the embedding (bottom).

the centers of the data-rich base categories available during training (while clustering is used in semi-supervised learning due to absence of labels). Every epoch, these action category prototypes are iteratively updated and stored into a fixed-sized vector by category-wise mean average, which we refer to as the *Prototype Memory Bank* (PMB). These action category prototypes are then used for feature augmentations in order to estimate the consistency cost.

**Prototype-based** feature enhancement with selfaugmentation warm-up. Leveraging prototype-based augmentation in context of one-shot learning requires further conceptual changes. First, since the prototypes indeed correspond to actual action categories from  $C_{base}$ (*i.e.* only one of the current training categories is correct), we first apply Softmax normalization across the channel dimension for prototypes vector  $\mathbf{E}_{P}^{*}$  and then refine it with the feature  $\mathbf{E}_{N-1}^*$  and project it into an embedding space as  $\mathbf{E}_{r,N-1}^* = g_{\mu}^2(Softmax(\mathbf{E}_P^*) \cdot \mathbf{E}_{N-1}^*)$ , where N indicates the total number of the transformer stage blocks.  $\mathbf{E}_{N-1}^*$  is obtained through  $\mathbf{E}_{N-1}^* = f_{\theta}^{1 \to N-1}(\mathbf{E}_{patch})$ , where  $f_{\theta}^i$ indicates the *i*-th transformer stage block for the auxiliary branch and Epatch is the mixed fused patch embedding  $E_{mixed}$ . At the same time, the feature  $\mathbf{E}_{N-1}^*$  is also projected as  $\mathbf{E}_{l,N-1}^* = g_{\mu}^1(\mathbf{E}_{N-1}^*)$ . Then, the attention weight **W** is calculated as  $\mathbf{W} = Softmax(\mathbf{E}_{l,N-1}^* \mathbf{E}_{r,N-1}^*)$ . After aggregating the information coming from prototypes vector  $\mathbf{E}_P^*$  to the original feature  $\mathbf{E}_{N-1}^*$  as depicted in Eq. (15):

$$\mathbf{E}_{agg,N-1}^{*} = g_{\mu}^{3}([\mathbf{W}\mathbf{E}_{r,N-1}^{*},\mathbf{E}_{l,N-1}^{*}]), \qquad (15)$$

the final augmented feature  $\mathbf{E}_{aug}^*$  is then obtained through a residual connection with the original feature  $\mathbf{E}_{N-1}^*$  by  $\mathbf{E}_{aug}^* = ReLU(\mathbf{E}_{N-1}^* + \mathbf{E}_{agg,N-1}^*)$ , where  $g_{\mu}^1$  and  $g_{\mu}^2$  indicate two fully-connected (fc) layers (no weight sharing), and  $g_{\mu}^3$  indicates a stack of two fc layers with ReLU in between. [·] denotes concatenation.

As in our case the prototypes are linked to true action categories from  $C_{base}$  (in contrast to unsupervised clustering necessary in self-supervised tasks), using centers of the assigned categories in the early training epochs would be unreliable. To alleviate this issue, we introduce an additional *warm-up phase*. The key idea is to leverage self-augmentation

instead of prototype-based augmentation until certain level of convergence is reached. At earlier stages, we therefore replace the attended prototype representation with the embedding  $\mathbf{E}_{N-1}^*$ . Figure 4 illustrates the difference between the selfaugmentation warm-up phase (top) and the prototype-based augmentation (bottom). Then, we switch to the phase at the bottom of Figure 4, while leveraging zero prototypes to achieve decenterization for a fixed number of epochs before using the class-agnostic prototype to do the feature-level prototype-based augmentation. The final augmented embeddings  $E^*$  can be obtained through  $\mathbf{E}^* = EMB(f_{\theta}^N(\mathbf{E}_{aug}^*)),$ where EMB indicates the multi-layer perceptron-based embedding generation layers. The final embedding from the main branch is obtained through  $\mathbf{E} = EMB(f_{\delta}^{N}(\mathbf{E}_{N-1}))$ . The embedding  $\mathbf{E}_{N-1}$  from the main branch is obtained through  $\mathbf{E}_{N-1} = f_{\delta}^{1 \to N-1}(\mathbf{E}_{patch})$ , where  $f_{\delta}^{i}$  indicates the *i*-th transformer stage block of the main branch. After the acquisition of the embeddings from the main branch E and the augmented

#### D. Deep Metric Learning Loss and Classification Loss

computed as  $L_{LSC} = 1 - cos(\mathbf{E}, \mathbf{E}^{\star})$ .

embeddings from the auxiliary branch  $E^*$ , the LSC loss is

**Triplet margin loss**. To harvest more discriminative embeddings, a triplet margin loss is leveraged in our training pipeline. Assuming  $\mathbf{a}_i$ ,  $\mathbf{p}_i$ , and  $\mathbf{n}_i$  denote the *i*-th selected anchor, the corresponding positive anchor, and the corresponding negative anchor in the latent space, where the positive anchor shares the same class with the selected anchor and the negative anchor has a different class compared to the selected anchor. The triplet margin loss is achieved through decreasing the distance between the selected anchor and the negative anchor while increasing the distance between the selected anchor and the negative anchor as depicted in Eq. (16):

$$L_{TPL} = \sum_{i=1}^{B} max \{ D(\mathbf{a}_i, \mathbf{n}_i) - D(\mathbf{a}_i, \mathbf{p}_i) + \sigma, 0 \} / B, \qquad (16)$$

where  $\sigma$  denotes the predefined margin and  $D(\cdot)$  denotes the pairwise distance. Assuming the pairwise distance between **a** and **n** is desired, the pairwise distance can be calculated as  $D(\mathbf{a}, \mathbf{n}) = ||\mathbf{a} - \mathbf{n} + \varepsilon||^2$ , where  $\varepsilon$  keeps as  $1e^{-6}$  and *B* denotes the batch size.

**Classification loss.** A cross-entropy loss is employed for the supervision of the training to ensure the classifiable capability of the learned embeddings in the latent space. Assuming  $\mathbf{p}_i$  denotes the prediction of the classifier of the model and  $\mathbf{y}_i$  denotes the label for the *i*-th sample inside a batch, the cross-entropy loss can be obtained through Eq. (17):

$$L_{CLS} = \sum_{i=1}^{B} -\mathbf{y}_i log(\mathbf{p}_i) + (1 - \mathbf{y}_i) log(1 - \mathbf{p}_i))/B$$
(17)

# V. EXPERIMENTS

#### A. Dataset Introduction

We perform comprehensive studies for the SOAR on three challenging datasets: NTU-60 [56], NTU-120 [2] and Toyota Smart Home [3]. We follow the SOAR protocol in NTU-120

 TABLE I

 EXPERIMENTS FOR SOAR WITHOUT OCCLUSION ON NTU-120 [2].

Encoder	Accuracy	F1	Precision	Recall					
Previously Published Appro	aches								
AN <sup>†</sup> [36]	41.0	-	-	-					
FC <sup>†</sup> [36]	42.1	-	-	-					
AP <sup>†</sup> [36]	42.9	-	-	-					
APSR [36]	45.3	-	-	-					
TCN-OneShot [34]	46.3	-	-	-					
SL-DML [22]	50.9	-	-	-					
Skeleton-DML [21]	54.2	-	-	-					
CNN-based Encoder Optimi	ized by DMI								
SL-DML (AlexNet [59])	40.33	39.14	42.42	40.35					
SL-DML (SqueezeNet [60])	42.55	40.52	41.88	42.51					
SL-DML (ResNet18 [61])	49.19	47.54	49.80	49.23					
Transformer-based Encoder	Optimized v	with DM	L (Ours)						
SL-DML (CaiT [42])	47.86	47.53	50.06	47.94					
SL-DML (ViT [38])	48.45	47.40	48.59	48.52					
SL-DML (Twins [40])	49.00	48.04	49.30	49.06					
SL-DML (ResT [43])	52.58	51.86	53.99	52.61					
SL-DML (Swin [41])	53.13	52.09	53.48	53.16					
SL-DML (LeViT [26])	53.19	52.22	53.85	53.29					
Our Proposed and Extended Approaches (Ours)									
SL-DML (LeViT) + LSC	55.94	54.29	55.80	56.04					
Trans4SOAR (Small)	56.27	56.43	58.59	56.32					
Trans4SOAR (Base)	57.05	55.90	57.26	57.12					

and formulate the evaluation protocols of Toyota Smart Home and NTU-60 to suit our data-scarce representation learning task. Additionally, we propose the occluded SOAR benchmarks for the first time building on top of these three datasets (see Sec. III). The NTU-120/NTU-60/Toyota Smart Home benchmarks feature 100/48/24 data-rich training categories and 20/12/7 data-scarce test categories respectively for one reference per unseen category. The protocols and occluded datasets will be publicly available in our benchmark.

#### **B.** Implementation Details

For TRANS4SOAR training we set the warm-up phase threshold  $N_t = 20$  while using another 10 epochs for decenterization. We train our model optimized by AdamW [62] with Cosine Annealing Scheduler for 50 epochs and batch size of 32 using Nvidia A100 GPU with PyTorch 1.8.0 to reproduce the best performance. We use an initial learning rate of  $3.5e^{-5}$ with the weights of the three losses, *i.e.*, Triplet Margin Loss  $(\sigma = 0.2)$ , Cross Entropy Loss and LSC loss as 1.0, 0.4 and 0.1. Our TRANS4SOAR (Small) has  $D_{Key}$ : 1,  $N_{head}$ : [2,2,2],  $H_{dep}$ : [2,4,4] and  $C_{dim}$ : [384,512,512] with 23M parameters while TRANS4SOAR has  $D_{Kev}$ : 32,  $N_{head}$ :[6,9,12],  $H_{dep}$ : [4,4,4] and  $C_{dim}$ : [384,512,768] with 43M parameters, where  $D_{Kev}$ ,  $N_{head}$ ,  $H_{dep}$  and  $C_{dim}$  denote dimension of Key, number of the attention head, number of the basic transformer attention block inside each transformer block and the unified feature dimension inside each transformer block respectively. Both of our approaches have three main transformer blocks. To ensure that there is no information leakage caused by the data augmentation to the occlusion region, the occlusion is generated before the data augmentation for both the realistic occlusion scenario and the random occlusion scenario. The protocols and the occlusion benchmarks will be released.

# C. Analyses for SOAR Without Occlusion

Performance analyses regarding different components. As in Table I, we firstly empirically evaluate the benefits brought by the LSC loss achieved through prototype-based feature augmentation and an additional auxiliary branch. The baseline we chose is SL-DML [22], which has the same data preprocessing technique with our approach. Specifically, we use the SL-DML pipeline and equip the selected transformer-based architecture, *i.e.*, LeViT [26], with an additional auxiliary branch for attention-based augmentations via feature-level prototypes. The LSC loss is obtained through the calculation of cosine similarity loss between the embedding from the main branch and the embedding from the auxiliary branch. The aforementioned approach with LSC loss is denoted as SL-DML (LeViT) + LSC compared with SL-DML (LeViT), which replaces CNN by the LeViT in the SL-DML pipeline. Although LSC loss does not have any influence on the architecture at test-time, it performs surprisingly well for the SOAR task without occlusion. For instance, SL-DML (LeViT) + LSC Loss leads to accuracy gains by 2.75% (NTU-120, Table I), 3.22% (NTU-60, Table II (a)) and 2.24% (Toyota Smart Home, Table III (a)), compared with SL-DML (LeViT), which has overall better performance compared with SL-DML [22] and Skeleton-DML [21] for the SOAR without occlusion. We observe the benefits of our LSC loss on NTU-120 [2], surpassing the previous two approaches, *i.e.*, SL-DML [22] by >5% and Skeleton-DML [21] by 1.74% (Table I). Our ablation experiments regarding the main components of LSC loss are shown in Table IV regarding the last three experiments, showing the importance of the warm-up stage and de-centerization stage which bring a performance improvement by 1.5% compared with LSC loss without both the aforementioned components. Then, the combination of the LSC loss and the MAFM, mixing three streams of input at patch embedding level, further contributes a remarkable performance gain regarding the SOAR without occlusion compared with the existing state-of-theart works [21], [21]. On the NTU-120 [2], TRANS4SOAR (Base) surpasses Skeleton-DML [21] and SL-DML [22] by 2.85% and 6.15% for accuracy while outperforming SL-DML (LeViT) + LSC by 1.11%, indicating an incremental performance enhancement considering the discriminative ability of the learned embedding by using MAFM and LSC loss. We also conduct experiments to showcase the individual performance gain brought by LSC and MAFM in Table IV regarding the first three experiments. Furthermore, consistent improvements are achieved by TRANS4SOAR in the other two datasets, e.g., NTU-60 in Table II (a) and Toyota Smart Home [3] in Table III (a) for the SOAR without occlusion. The NTU-60 [56] has less training categories than the NTU-120 [2], thus, it is used to evaluate the generalizability of the leveraged models, which means realizing the SOAR with less a prior knowledge. In Table II (a), our TRANS4SOAR (Base) surpasses SL-DML [22] and Skeleton-DML [21] by 19.37% and 18.65% for accuracy, indicating that, given less a prior knowledge, TRANS4SOAR has better capability to harvest more discriminative representation. Furthermore, the Toyota Smart Home [3] contains 2D skeleton data in image coordinate

#### TABLE II

# EXPERIMENTS ON THE NTU-60 [56] FOR SOAR CONSIDERING THE SCENARIOS (A) WITHOUT OCCLUSION, (B) WITH REALISTIC OCCLUSION (RE) AND (C) WITH RANDOM OCCLUSION (RA).

	(a	) Withou	t Occlus	ion	(b) With RE					(c) Wi	th RA	
Encoder	Acc.	F1	Prec.	Rec.	Acc.	F1	Prec.	Rec.	Acc.	<b>F1</b>	Prec.	Rec.
Previously Published App	oroaches											
SL-DML [22] Skeleton-DML [21]	54.82 55.54	54.31 50.88	56.72 53.13	54.65 51.24	36.90 42.66	35.86 40.90	36.59 41.50	37.05 42.82	45.28 60.43	43.13 59.66	45.00 61.37	45.42 60.54
Transformer-based Encod	Transformer-based Encoder Optimized by DML (Ours)											
SL-DML (Swin [41]) SL-DML (LeViT [26])	56.99 64.45	56.24 64.17	58.67 66.35	56.99 64.47	51.71 52.72	50.60 52.19	52.54 54.90	51.82 52.86	64.65 56.73	63.74 55.89	66.57 57.57	64.77 56.85
Our Extended and Evalu	ated App	oroached	(Ours)									
SL-DML (LeViT) + LSC Trans4SOAR (Small) Trans4SOAR (Base)	67.67 69.74 <b>74.19</b>	67.87 70.52 <b>74.34</b>	68.74 72.45 <b>75.91</b>	67.67 69.82 <b>74.20</b>	53.79 56.84 <b>59.28</b>	52.76 55.84 <b>58.96</b>	54.18 58.27 <b>59.91</b>	53.88 56.98 <b>59.40</b>	60.78 67.90 <b>72.59</b>	58.75 67.32 <b>71.82</b>	59.97 68.94 <b>73.89</b>	60.90 68.01 <b>72.66</b>

 TABLE III

 EXPERIMENTS ON THE TOYOTA SMART HOME [3] FOR SOAR (A) WITHOUT OCCLUSION, (B) WITH REALISTIC OCCLUSION (RE) AND (C) WITH RANDOM OCCLUSION (RA).

Fuendau	(a)	Withou	t Occlusi	on		(b) Wi	ith RE			(c) Wi	ith RA	
Elicouer	Acc.	F1	Prec.	Rec.	Acc.	F1	Prec.	Rec.	Acc.	F1.	Prec.	Rec.
Previously Published App	roaches											
SL-DML [22]	58.98	27.15	27.64	35.00	38.93	25.16	32.93	28.48	53.79	26.28	27.24	29.67
Skeleton-DML [21]	47.31	18.45	18.58	23.80	47.67	24.86	27.93	27.35	48.91	21.60	25.00	21.75
Transformer-based Encod	ler Optin	nized by	DML (C	Ours)								
SL-DML (Swin [22])	58.76	28.83	29.17	32.34	35.43	18.48	23.24	23.80	65.50	29.20	30.78	29.69
SL-DML (LeViT [26])	62.22	31.98	37.56	35.16	38.48	22.58	27.66	24.62	61.96	26.42	28.52	29.20
Our Extended and Evaluation	ated App	roached	(Ours)									
SL-DML (LeViT) + LSC	64.46	31.91	34.07	33.58	41.82	24.34	29.02	26.67	63.77	27.72	29.09	29.90
Trans4SOAR (Small)	66.87	28.08	31.47	34.63	55.12	26.90	29.41	30.69	68.47	28.86	29.56	32.25
Trans4SOAR (Base)	70.22	33.96	37.81	35.33	60.15	25.50	33.12	31.86	68.91	29.27	34.15	31.45

TABLE IV Ablation study of LSC and MAFM used in the Trans4SOAR on NTU-60 [56] without occlusion.

With LSC	Self-aug. wp	De-centerization	MAFM	Accuracy	F1	Precision	Recall
				64.45	64.17	66.35	64.47
√	√	√		67.67	67.87	68.74	67.67
			$\checkmark$	71.55	71.85	73.45	71.63
√			$\checkmark$	72.69	72.80	74.27	72.73
√		√	$\checkmark$	73.09	73.39	74.54	73.14
√	√	√	$\checkmark$	74.19	74.34	75.91	74.20

TABLE V A comparison to other encoder architectures.

Methods	Accuracy	F1	Recall	Precision
SL-DML (CTR-GCN[5])	43.92	41.38	45.21	43.89
SL-DML (STTR[10])	39.56	39.45	41.92	39.58
SL-DML (LeViT) + LSC	55.94	54.29	55.80	56.04
Trans4SOAR (Small)	56.27	56.43	58.59	56.32
Trans4SOAR (Base)	57.05	55.90	57.26	57.12

format, delivering a valuable data format to explore the SOAR task. In Table III (a), our TRANS4SOAR (Base) undoubtedly shows the best performance over all the previous approaches with large margin. Observing the other three metrics, *i.e.*, F1-score, precision and recall, since the first two datasets have balanced distributed samples for different categories, theses three terms do not have large difference compared with the accuracy. However, since the action categories on the Toyota Smart Home [3] is not equal distributed, these three terms are able to showcase whether the true prediction is balanced distributed in the test set or not. Our TRANS4SOAR surpasses all the approaches in terms of all metrics on the investigated datasets. In order to ablate the effect of different model scales, we construct TRANS4SOAR (Small) with only

 TABLE VI

 Experiments for SOAR without occlusion on NTU-120 [2]

 UNDER GAUSSIAN NOISE DISRUPTION.

Gaussian Noise	$\sigma = 0.1$	$1, \mu = 0$		$\sigma = 0.05, \mu = 0$				
Encoder	Acc.	F1	Prec.	Rec.	Acc.	F1	Prec.	Rec.
SL-DML [22]	21.42	11.83	8.50	21.71	21.76	12.23	8.70	21.86
SL-DML (LeVIT) SL-DML (LeVIT) + LSC Trans4SOAR	52.54 53.09	51.16 51.89	51.61 53.05	52.65 53.15	51.91 54.74	50.08 54.65	51.67 56.33	52.01 54.83

 TABLE VII

 EXPERIMENTS REGARDING REALISTIC SYNTHESIZED OCCLUSION (A)

 AND RANDOM OCCLUSION (B) FOR SOAR ON NTU-120 [2].

Encoder		(a) W	ith RE		(b) With RA				
Encouer	Acc.	F1	Prec.	Rec.	Acc.	F1	Prec.	Rec.	
SL-DML [22]	39.82	37.85	39.32	39.86	42.53	42.24	44.79	42.56	
Skeleton-DML [21]	49.21	46.82	48.10	49.18	35.15	32.59	34.29	35.22	
SL-DML (LeViT [26])	44.22	42.29	44.20	44.31	35.00	33.24	41.45	35.10	
SL-DML (Swin [41])	47.19	45.64	46.78	47.29	47.19	45.64	46.78	47.29	
SL-DML (LeViT) + LSC	48.28	46.03	47.58	48.31	38.04	35.93	37.87	38.11	
Trans4SOAR (Small)	51.64	50.47	52.36	51.70	53.27	51.33	53.80	53.35	
Trans4SOAR (Base)	52.35	48.79	52.87	52.43	53.17	52.89	54.50	53.21	

23M parameters which pursues both light model structure and high accuracy, and achieves second best performance, showcasing that the LSC loss and MAFM are helpful for learning discriminative features via different model variants. We also conduct experiments in Table V to compare with graph convolutional approach [5] and skeleton transformer approach [10], however the performance of these two encoder architectures for the SOAR task even without occlusion is not satisfied compared with TRANS4SOAR and SL-DML (LeViT).

Tolerance to noisy inputs. The quality of the skeleton data

TABLE VIII EXPERIMENTS WITH DIFFERENT REALISTIC SYNTHESIZED OCCLUSION RATIO ON THE NTU-60 [56] FOR THE SOAR.

Model	RE_Range	Accuracy	F1	Precision	Recall
SL-DML [22]		36.90	35.86	36.59	37.05
Skeleton-DML [21]		35.15	32.59	34.29	35.22
SL-DML (LeViT [26])	0.05-0.2	52.72	52.19	54.90	52.86
SL-DML (LeViT) + LSC		53.79	52.76	54.18	53.88
Trans4SOAR (Small)		56.84	55.84	58.27	56.98
Trans4SOAR (Base)		59.28	58.96	59.91	59.40
SL-DML [22]		39.26	38.71	39.59	39.43
Skeleton-DML [21]		38.52	38.74	39.23	38.64
SL-DML (LeViT [26])	0.05.0.25	53.17	52.52	54.16	53.34
SL-DML (LeViT) + LSC	0.05-0.55	53.58	52.75	54.07	53.77
Trans4SOAR (Small)		61.69	61.60	64.01	61.81
Trans4SOAR (Base)		58.27	56.63	58.81	58.40
SL-DML [22]		34.89	32.63	31.85	35.07
Skeleton-DML [21]		42.83	42.33	42.46	42.93
SL-DML (LeViT [26])	0.05-0.5	54.84	54.07	57.06	54.99
SL-DML (LeViT) + LSC		55.07	55.01	57.56	55.21
Trans4SOAR (Small)		59.59	59.21	59.49	59.70
Trans4SOAR (Base)		57.52	57.21	59.61	57.64

TABLE IX EXPERIMENTS FOR DIFFERENT FUSION TECHNIQUES ON NTU60 [56] UNDER DIFFERENT OCCLUSION SCENARIOS.

Fusion Method	OCC	Accuracy	F1	Precision	Recall
Single (Joints)	RE	53.79	52.76	54.18	53.88
Single (Bones)	RE	54.22	53.73	54.86	54.33
Single (Velocities)	RE	56.93	56.10	57.97	57.03
Addition	RE	56.37	54.48	55.68	56.51
Multiplication	RE	53.35	51.91	53.69	53.50
Concatenation	RE	58.61	57.21	57.63	58.73
Late Fusion	RE	56.93	56.10	57.97	57.03
MAFM	RE	59.28	58.96	59.91	59.40
Single (Joints)	RA	60.78	58.75	59.97	60.90
Single (Bones)	RA	55.15	53.56	56.63	54.16
Single (Velocities)	RA	33.15	30.54	29.67	33.82
Addition	RA	65.09	65.03	66.36	65.18
Multiplication	RA	67.54	67.51	68.65	67.63
Concatenation	RA	68.05	68.54	70.90	68.13
Late Fusion	RA	71.16	71.58	73.16	71.22
MAFM	RA	72.59	71.82	73.89	72.66
Single (Joints)	Ν	67.67	67.87	68.74	67.67
Single (Bones)	Ν	61.45	61.44	63.50	61.57
Single (Velocities)	Ν	49.74	50.08	51.31	49.89
Addition	Ν	67.05	66.88	68.09	67.12
Multiplication	Ν	64.63	65.05	66.34	64.75
Concatenation	Ν	67.75	67.79	69.56	67.86
Late Fusion	Ν	57.15	56.52	57.57	57.26
MAFM	Ν	74.19	74.34	75.91	74.20

is influenced by a variety of factors, such as sensor noise or occlusions. First, a larger gap between the TRANS4SOAR and standard DML trained on the Toyota Smart Home [3] (which is noisier than the more controlled NTU-datasets) hints towards its advantages specifically for imperfect input. To validate if this is the case, we evaluate the model for inputs corrupted by different magnitudes of Gaussian noise and discover a remarkable tolerance of TRANS4SOAR (in Table VI). While the prediction quality diminishes for basic DML-based models, the utilizing of LSC loss on the SL-DML (LeViT) is more robust when confronted with unreliable data, which showcases the superiority of the proposed LSC loss against Gaussian noise input.

In particular, the performance for the SL-DML (LeViT) with the LSC loss falls from 55.94% on clean data to 51.91% for Gaussian noise with  $\sigma$ =0.05, while this decline is much higher (53.19%  $\rightarrow$  21.97%) for the SL-DML (LeViT). We attribute this to the extensive learned augmentations at the featurelevel taking place in the auxiliary branch while formulating the LSC loss. The LSC loss encourages the model to output similar results if the embedding has been altered, which suits naturally to the use-case of noise disturbances. Furthermore, TRANS4SOAR (Base) surpasses all the other investigated

TABLE X EXPERIMENTS REGARDING DIFFERENT RANDOM OCCLUSION RATIO ON THE NTU-60 [56] FOR THE SOAR.

Model	RA_ratio	Accuracy	F1	Precision	Recall
SL-MDL [22]		45.28	43.13	45.00	45.42
Skeleton-DML [21]		60.43	59.66	61.37	60.54
SL-DML (LeViT [26])	0.1	56.73	55.89	57.75	56.85
SL-DML (LeViT) + LSC	0.1	60.78	58.75	59.97	60.90
Trans4SOAR (Small)		69.74	70.52	72.45	69.82
Trans4SOAR (Base)		72.59	71.82	73.89	72.66
SL-DML [22]	1	46.39	42.82	46.69	46.54
Skeleton-DML [21]		58.93	56.07	58.45	59.05
SL-DML (LeViT [26])	0.2	46.32	43.78	43.94	46.40
SL-DML (LeViT) + LSC	0.5	47.82	45.02	48.41	47.91
Trans4SOAR (Small)		66.57	66.26	67.94	66.65
Trans4SOAR (Base)		72.39	72.81	74.68	72.43
SL-DML [22]		43.44	38.46	41.30	43.57
Skeleton-DML [21]		44.69	41.89	45.74	44.79
SL-DML (LeViT [26])	0.5	35.77	32.56	36.22	35.94
SL-DML (LeViT) + LSC		40.53	37.38	38.33	40.59
Trans4SOAR (Small)		52.92	50.78	55.13	53.02
Trans4SOAR (Base)		54.82	55.01	58.01	54.93

TABLE XI EXPERIMENTS FOR REFERENCE W/ OR W/O OCCLUSIONS ON NTU-60 [56].

Model	000	OCCVal	Accuracy	F1	Procision	Decall
Model	occ	OCCVAI	Accuracy	11	Trecision	Recan
SL-MDL [22]			48.74	46.46	47.45	48.88
Skeleton-DML [21]			49.30	48.57	49.62	49.45
SL-DML (LeViT [26])	DA	т	53.47	52.35	54.94	53.63
SL-DML (LeViT) + LSC	KA	1	53.57	53.73	56.55	53.72
Trans4SOAR (Small)			72.16	72.42	73.67	72.23
Trans4SOAR (Base)			71.59	72.22	73.95	71.67
SL-DML [22]			45.28	43.13	45.00	45.42
Skeleton-DML [21]			60.43	59.66	61.37	60.54
SL-DML (LeViT [26])	DA	E	56.73	55.89	57.57	56.85
SL-DML (LeViT) + LSC	KA	г	60.78	58.75	59.97	60.90
Trans4SOAR (Small)			67.90	67.32	68.94	68.01
Trans4SOAR (Base)			72.59	71.82	73.89	72.66
SL-DML [22]			36.90	35.86	36.59	37.05
Skeleton-DML [21]			42.66	40.90	41.50	42.82
SL-DML (LeViT [26])	DE	- T	52.72	52.19	54.90	52.86
SL-DML (LeViT) + LSC	KE	1	53.79	52.76	54.18	53.88
Trans4SOAR (Small)			56.84	55.84	58.27	56.98
Trans4SOAR (Base)			59.28	58.96	59.91	59.40
SL-DML [22]			39.51	39.64	40.82	39.64
Skeleton-DML [21]	DE	F	44.29	43.10	44.26	44.46
SL-DML (LeViT [26])	RE .	l r	55.12	55.22	57.51	55.26
SL-DML (LeViT) + LSC			55.07	55.01	57.56	55.21
Trans4SOAR (Small)			54.37	52.97	55.08	54.38
Trans4SOAR (Base)			58.48	57.10	57.75	58.61

approaches with no doubt by 54.74% and 53.09% in terms of accuracy for the Gaussian noise conditioned by  $\sigma$ =0.05 and  $\sigma$ =0.1 respectively.

# D. Analyses for REalistic Synthesized Occlusion (RE)

We conduct experiments regarding RE for NTU-120 [2], NTU-60 [56] and Toyota Smart Home [3] in Table VII (a), Table II (b), and Table III (b) with SNR range  $0.05 \rightarrow 0.2$  and with occlusion on the reference set. First, the performance of all investigated approaches for the SOAR with RE benchmark is degraded compared to the SOAR without occlusion benchmark, indicating that the proposed RE is very challenging for discriminative representation learning. In Table VII (a), our TRANS4SOAR (Base) shows the best performance by 52.35%, 48.79%, 52.87% and 52.43% for accuracy, F1-score, precision and recall, indicating that the performance is equally distributed among the investigated classes on the NTU-120 [2]. The TRANS4SOAR (Small) achieves second best performance among all the metrics on NTU-120 with RE by 51.64% for accuracy. Note, that the SL-DML (LeViT) demonstrates worse performances on all the conducted datasets with RE for SOAR. On NTU-120 [2] with RE, the SL-DML (LeViT) approach only has an accuracy of 44.22% which is lower than the Skeleton-DML [21] with an accuracy of 49.21%. However,

TABLE XII EXPERIMENTS FOR RANDOM TEMPORAL AND SPATIAL OCCLUSION.

Model	(a) Ra	ndom tei	nporal o	cclusion	(b) Ra	ndom sp	atial occl	lusion
	Acc.	F1.	Prec.	Rec.	Acc.	F1.	Prec.	Rec.
Experiments on NTU	-120 wit	th rando	m tempo	ral occlus	sion.			
SL-DML [22]	38.15	34.87	38.51	38.11	38.15	35.26	36.76	38.13
Skeleton-DML [21]	27.20	24.43	26.75	27.12	27.93	25.91	28.24	27.93
Trans4SOAR (Small)	51.60	50.73	52.65	50.99	46.99	46.24	49.71	47.07
Trans4SOAR (Base)	54.11	52.93	53.85	54.21	49.43	49.08	51.35	49.48
Experiments on NTU	-60 with	randon	ı tempor	al occlusi	on.			
SL-DML [22]	58.68	58.46	60.20	58.72	52.48	50.59	54.06	52.65
Skeleton-DML [21]	51.81	50.50	53.06	51.95	46.38	43.68	45.94	46.54
Trans4SOAR (Small)	71.45	71.32	72.94	71.51	68.94	69.61	71.84	69.01
Trans4SOAR (Base)	75.01	74.75	75.76	75.06	69.08	69.18	71.19	69.14
Experiments on Toyo	ta Smart	Home	with ran	dom tem	poral occ	lusion.		
SL-DML [22]	53.36	22.97	28.17	24.58	60.36	20.10	24.52	20.89
Skeleton-DML [21]	53.65	23.90	31.54	25.33	41.95	26.36	32.83	27.97
Trans4SOAR (Small)	63.66	29.90	31.76	34.06	66.76	31.76	33.14	35.66
Trans4SOAR (Base)	68.48	31.11	33.81	33.80	64.49	32.43	35.80	34.29

compared with SL-DML (LeViT), SL-DML (LeViT) + LSC loss still has a better accuracy of 48.28%, indicating that LeViT architecture is not good at dealing with RE, while LSC loss can alleviate the negative influence. After the using of the MAFM to form our TRANS4SOAR (Base), a superior performance of 52.35% in accuracy shows up, indicating that the disruption issue caused by RE can be well addressed through the triplet stream encoding and MAFM. These experimental results illustrate the importance of the proposed MAFM on dealing with the disruption brought by the RE through aggregating three different skeleton encoding formats, which contains potential de-occlusion cues, and also show the superiority of our reformulated TRANS4SOAR regarding the robustness against the occlusion disruption from the real life compared with the LeViT, on which we build up our TRANS4SOAR based. Similar comparison and analyses could be found on the other two datasets, i.e., NTU-60 [56] in Table II (b) and Toyota Smart Home [3] in Table III (b), where the TRANS4SOAR (Base) surpasses Skeleton-DML [21] and SL-DML [22] by 16.62% and 22.38% on NTU-60 [56], and 12.48% and 21.22% on Toyota Smart Home [3], while TRANS4SOAR (Small) also shows competitive performances. We also conduct experiments by using different Sigal-to-Noise Ratio (SNR) range for the SOAR with RE as depicted in Table VIII, TRANS4SOAR shows promising and stable performance > 56% in terms of accuracy considering both TRANS4SOAR (Base) and TRANS4SOAR (Small) for three SNR ranges, *i.e.*, 0.05 - 0.2, 0.05 - 0.35 and 0.05 - 0.5 on NTU-60 [56].

# E. Analyses Regarding Random Occlusion (RA)

Random occlusion, considered as a combination between random temporal and spatial occlusions, is leveraged as the second main occlusion in our work on NTU-120 [2], NTU-60 [56] and Toyota Smart Home [3], depicted in Table VII (b), Table II (c) and Table III (c), with SNR = 0.1 and without occlusion on reference set respectively. The proposed TRANS4SOAR (Base) keeps surpassing all the existing approaches by large margins, *e.g.*, SL-DML [22] and Skeleton-DML [21] by 10.64% and 18.02% on NTU-120 [2]. The performance of Skeleton-DML [21] under RA is worse than that of SL-DML [22], while the case is reversed on RE, which means most of the existing approaches can not be robust

against different occlusions. However, TRANS4SOAR overcomes this issue and demonstrates promising performances over different occlusions, especially for RE, which is an important ability for learning discriminative representation. The proposed MAFM is also proved to have strong capability while dealing with different occlusions, which is well addressed through taking the three stream of skeleton patch embedding as input. MAFM is further illustrated as the best fusion architecture among all the investigated fusion methods regarding the two main occlusions in Table IX which will be introduced later. Considering the three streams encoding, first, since bone and velocity use temporal and spatial difference respectively, more cues regarding the neighbourhood of the occluded region could be encoded in different perspectives. Furthermore, while tackling with RA on the NTU-60 [56] and Toyota Smart Home [3], TRANS4SOAR (Base) and TRANS4SOAR (Small) also demonstrate promising state-ofthe-art performance across all the leveraged metrics, which reflects the strengths of the proposed models in multiple point of views. Similar ablations regarding the SNR ratio, *i.e.*, 0.1, 0.2 and 0.3, of the RA, are done in Table X, where the performances of TRANS4DARC (Base) and TRANS4DARC (Small) surpass all the leveraged approaches. The experiments are done with occlusion on reference set. Especially for SNR =0.1 and SNR = 0.3, TRANS4DARC (Base) achieves 72.59% and 72.39% for accuracy while the Skeleton-DML [21] only achieves 60.43% and 58.93%. Due to the large disruption by using SNR = 0.5 with RA, the performance of TRANS4SOAR (Base) only achieves 54.82% accuracy while still outperforming the state-of-the-art approach by 10.13%.

#### F. Analyses for Occlusion Disruption on Reference Samples

Experiments are conducted in Table XI to investigate different occlusions on the NTU-60 [56] reference set with the occlusion state denotes by OCCVal, where T and F indicate with occlusion and without occlusion. We set SNR = 0.1 for RA and SNR range  $0.05 \rightarrow 0.2$  for RE, which is comparable regarding averaged SNR. SL-DML [22] and Skeleton-DML [21] have absolute performance fluctuation for accuracy by 3.46% and 11.13% for RA, and 2.61% and 1.63% for RE. What we desire is that the model should have small absolute fluctuation regarding different OCCVal setting. The experimental results of TRANS4SOAR suit this desire with absolute fluctuation 1.00% for RA and 0.80% for RE, illustrating the strong ability against the occlusion on the reference set. Considering RA and RE with OCCVal = F, both the performances of SL-DML [22] and Skeleton-DML [21] are worse with RE than that with RA, indicating that RE is more challenging.

#### G. Analyses for Ablation of Fusion Mechanisms

To demonstrate the efficiency of MAFM, we conduct comparison experiments among several fusion approaches in Table IX, where we set SNR = 0.1 for RA and SNR range  $0.05 \rightarrow 0.2$  for RE. The mostly leveraged fusion technique is late fusion which conducts fusion at the decision level. However the design of late fusion triplicates the model size as 113*M* while the others are at 40*M* level. Here, we consider



Fig. 5. An overview of the qualitative experimental results on NTU-60 [56] with RE for SOAR, where GT indicates the groundtruth and Trans4SOAR indicates the prediction of TRANS4SOAR-Base. The true prediction is marked as green, while the false prediction is marked as red.



Fig. 6. Failure case examples for SOAR with RE on NTU-60 [56].

to propose a efficient fusion mechanism at patch embedding level which takes both the model performance and size into consideration. The baselines for patch embedding level fusion includes the addition, multiplication and concatenation, which are directly executed after the acquisition of the patch embeddings for the three streams. Another method we compared with is late fusion, which conducts addition after obtaining the final embeddings of the three streams and has a three times larger model size. The experimental results indicate that MAFM has great performance compared with all the leveraged patch-embedding level fusiom baselines and the competitive late fusion on the NTU-60 [56] under different occlusions, e.g., No occlusion (N), REalistic synthesized occlusion (RE) and RAndom occlusion (RA). Specifically, TRANS4SOAR with MAFM surpasses the late fusion by 2.35%, 1.43%, and 17.04% on the RE, RA, and N respectively, while having a smaller model size for both inference and training. Simultaneously. TRANS4SOAR with MAFM surpasses the investigated approach with the best performance among the basic patch embedding level fusion approach by 0.67%, 4.54% and 6.44% for RE, RA and N.



Fig. 7. TSNE visualizations for (a) Skeleton-DML under RA, (b) TRANS4SOAR-Base under RA, (c) Skeleton-DML under RE and (d) TRANS4SOAR-Base under RE on NTU-60 [56].

# H. Analyses for Random Temporal and Spatial Occlusions

In order to show the performance of all the leveraged models with the two existing occlusions in the related work, *i.e.*, random *temporal* and *spatial* occlusion, which might also be interesting to the community regarding the specific occlusion considering temporal and spatial components individually, we conducted experiments on three datasets while using the most effective approaches investigated in our work, as described in Table XII (a) and Table XII (b), where we choose the occluded frame number as 10 and the occluded joints number as 5 respectively. Compared with RE, these two leveraged occlusions which is specific controlled through predefined

TABLE XIII THE COMPARISON IN TERMS OF ACCURACY, THE NUMBER OF PARAMETERS (#PARAMS), AND GFLOPS ON NTU-120 WITHOUT OCCLUSION [2]

Encoder	Accuracy	#Params	GFLOPS
Previously Published Approaches			
AN <sup>†</sup> [36]	41.0	-	-
FC <sup>†</sup> [36]	42.1	-	-
AP <sup>†</sup> [36]	42.9	-	-
APSR [36]	45.3	-	-
TCN-OneShot [34]	46.3	3.5M	8.5
SL-DML [22]	50.9	11.2M	23.8
Skeleton-DML [21]	54.2	11.2M	23.8
CNN-based Encoder Optimized by DML			
SL-DML (AlexNet [59])	40.33	57.1M	9.2
SL-DML (SqueezeNet [60])	42.55	0.7M	9.7
SL-DML (ResNet18 [61])	49.19	11.2M	23.8
GCN-based Encoder Optimized with DML (Ours)			
SL-DML (CTR-GCN [5])	43.92	1.6M	9.2
SL-DML (STTR [10])	39.56	7.0M	37.4
Transformer-based Encoder Optimized with DML (Ours)			
SL-DML (CaiT [42])	47.86	120.8M	53.9
SL-DML (ViT [38])	48.45	53.6M	27.1
SL-DML (Twins [40])	49.00	25.2M	75.1
SL-DML (ResT [43])	52.58	57.8M	61.3
SL-DML (Swin [41])	53.13	87.3M	29.3
SL-DML (LeViT [26])	53.19	38.9M	30.4
Our Proposed and Extended Approaches (Ours)			
SL-DML (LeViT) + LSC	55.94	38.9M	30.4
Trans4SOAR (Small)	56.27	23.1M	34.1
Trans4SOAR (Base)	57.05	43.8M	47.9

occluded frame and joint numbers are easier to be addressed as their randomness is not satisfied. But the important thing is that our proposed TRANS4SOAR (Base) and TRANS4SOAR (Small) still surpass all existing works by large margins on all datasets with these two occlusions which further illustrates the efficiency of our model against different occlusions, even the occlusion is specifically controlled by predefined concepts, *e.g.*, the occluded frame number.

# I. Analysis for Qualitative and TSNE Experimental Results

Qualitative analysis. The qualitative results are given in Figure 5 for SOAR with RE on the NTU-60 [56], where the occluded body joints are marked as red dots. TRANS4SOAR has overall great performance while comparing with Skeleton-DML [21] with true prediction 3:2 among the selected 4 samples. The occlusion of the joints which is dominant to the action has a large influence on the model, e.g., arm and hand joints for TakeOffGlasses action, where Skeleton-DML [21] gives a false prediction while TRANS4SOAR pursues a true prediction. However, due to the high similarity between several actions, e.g., WipeFace and DrinkingWater, TRANS4SOAR still has false prediction but the predicted DrinkingWater action is more similar with WipeFace compared with TearUpPaper predicted by Skeleton-DML [21], showing that there is still research space for the future research. We further present failure cases in Figure 6 to investigate the cause of the false prediction of our model for SOAR under RE. Since SOAR is only able to harvest informative classification cues from the information of the given fixed number of the human body joints, occlusion, which is assigned to the most dominant joint region during a specific movement type, causes large information decrease during the feature extraction

procedure resulting in false prediction for SOAR. Considering the first sample in Figure 6, the hand region is occluded when the person is picking something up, however the hand and arm region is the dominant region for the action throw. The information decrease on the dominant region makes our model predict pointing at something, which is a false prediction. Considering the second sample, most of the leg region are occluded by the projected object while the leg region is the dominant region during falling. The missing information also causes a negative effect on our model for the SOAR task.

**TSNE analysis.** In Figure 7, a TSNE [63] comparison among (a) Skeleton-DML under RA, (b) TRANS4SOAR-Base under RA, (c) Skeleton-DML under RE, and (d) TRANS4SOAR-Base under RE is shown to deliver a better understanding regarding the learned features on the latent space. First, compared with (a) and (c), (b) and (d) harvest clearer boundaries for more classes in the SOAR task, which showcases that TRANS4SOAR has the capability to obtain embeddings with more discriminative cues. Then, if we look at the same approach under different occlusions, smaller changes are shown for TRANS4SOAR-Base, as demonstrated in (b) and (d) while the the shape and structure of the latent space embeddings extracted from Skeleton-DML has larger changes as shown in (a) and (c). Overall, our TRANS4SOAR approach shows better robustness against the different occlusions from the perspective of the change of the learned features on the latent space.

# J. Analyses for the Model Efficiency.

To have a detailed look at the efficiency of different approaches, the accuracy for SOAR on no-occluded NTU-120, the number of the parameters, and the GFLOPS during inference are listed in Table XIII. The number of parameters and GFLOPS for the first four approaches under Previously Published Approaches are not available. Compared with the visual transformer-based approaches, CNN-based approaches and GCN-based approaches preserve a smaller number of parameters and the GFLOPS while mostly delivering an unsatisfied performance for the SAOR task. The high performance of the visual transformer-based approaches are not exactly due to using larger models, since SL-DML (CaiT) has the largest number of the parameter and SL-DML (ResT) has the largest GFLOPS, but they do not achieve better performances compared with SL-DML (LeViT), which has 53.19% in accuracy, 38.9M parameters, and 30.4 GFLOPS. TRANS4SOAR has a competitive amount of parameters and GFLOPS compared with other visual transformer approaches while delivering the best performance for the SOAR task. Especially, TRANS4SOAR (Small) shows 56.27% in accuracy with only 23.1M parameters and 34.1 GFLOPS. Since our model is a multi-modality model, a reasonable increment in terms of the number of parameters and the GFLOPS is expected. TRANS4SOAR (Base) achieves a > 70M parameter decrease compared with the late fusion approach while harvesting a better performance for SOAR, which demonstrates the superiority of TRANS4SOAR from the perspective of multi-modality fusion.

# VI. CONCLUSION

In this work, we look into the problem of data-scarce recognition of daily activities through the lens of one-shot recognition, while considering diverse occlusions. First, we propose realistic synthesized and random occlusion to better address the occlusion problem. Then, a novel architecture TRANS4SOAR is put forward to provide discriminative representations for skeleton input and enhance the robustness against different scenarios. We design a Mixed Attention Fusion Mechanism (MAFM), featuring a three-stream of skeleton encoding inputs to realize efficient fusion on the patchembedding level. Inspired by recent success of augmentationbased methods in semi-supervised learning, we further introduce the latent space consistency loss, which leverages an additional auxiliary branch encouraging the embedder to produce similar results despite extensive augmentations at the feature level. TRANS4SOAR sets the new state of the art on both normal and occluded SOAR benchmarks established on three datasets. In the future, occluded one-shot recognition based on video data is still attractive to be researched.

#### REFERENCES

- [1] X. Sun *et al.*, "Pix3D: Dataset and methods for single-image 3D shape modeling," in *CVPR*, 2018.
- [2] 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*, 2020.
- [3] S. Das et al., "Toyota smarthome: Real-world activities of daily living," in *ICCV*, 2019.
- [4] Y. Zhang, B. Wu, W. Li, L. Duan, and C. Gan, "STST: Spatial-temporal specialized transformer for skeleton-based action recognition," in *MM*, 2021.
- [5] Y. Chen, Z. Zhang, C. Yuan, B. Li, Y. Deng, and W. Hu, "Channelwise topology refinement graph convolution for skeleton-based action recognition," in *ICCV*, 2021.
- [6] S. Yan, Y. Xiong, and D. Lin, "Spatial temporal graph convolutional networks for skeleton-based action recognition," in AAAI, 2018.
- [7] R. Bai, "Hierarchical graph convolutional skeleton transformer for action recognition," arXiv preprint arXiv:2109.02860, 2021.
- [8] Y.-B. Cheng, X. Chen, D. Zhang, and L. Lin, "Motion-transformer: Self-supervised pre-training for skeleton-based action recognition," in *MMAsia*, 2021.
- [9] V. Mazzia, S. Angarano, F. Salvetti, F. Angelini, and M. Chiaberge, "Action transformer: A self-attention model for short-time human action recognition," *Pattern Recognition*, 2022.
- [10] C. Plizzari, M. Cannici, and M. Matteucci, "Skeleton-based action recognition via spatial and temporal transformer networks," *Computer Vision and Image Understanding*, 2021.
- [11] ——, "Spatial temporal transformer network for skeleton-based action recognition," in *ICPRW*, 2021.
- [12] Y.-F. Song, Z. Zhang, C. Shan, and L. Wang, "Constructing stronger and faster baselines for skeleton-based action recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [13] C. Cao, Y. Li, Q. Lv, P. Wang, and Y. Zhang, "Few-shot action recognition with implicit temporal alignment and pair similarity optimization," *Computer Vision and Image Understanding*, 2021.
- [14] Z. Gao, L. Guo, W. Guan, A.-A. Liu, T. Ren, and S. Chen, "A pairwise attentive adversarial spatiotemporal network for cross-domain few-shot action recognition-R2," *IEEE Transactions on Image Processing*, 2021.
- [15] J. Hong, M. Fisher, M. Gharbi, and K. Fatahalian, "Video pose distillation for few-shot, fine-grained sports action recognition," in *ICCV*, 2021.
- [16] J. Patravali, G. Mittal, Y. Yu, F. Li, and M. Chen, "Unsupervised fewshot action recognition via action-appearance aligned meta-adaptation," in *ICCV*, 2021.
- [17] T. Perrett, A. Masullo, T. Burghardt, M. Mirmehdi, and D. Damen, "Temporal-relational CrossTransformers for few-shot action recognition," in CVPR, 2021.

- [18] Y.-F. Song, Z. Zhang, and L. Wang, "Richly activated graph convolutional network for action recognition with incomplete skeletons," in *ICIP*, 2019.
- [19] Y.-F. Song, Z. Zhang, C. Shan, and L. Wang, "Richly activated graph convolutional network for robust skeleton-based action recognition," *IEEE Transactions on Circuits and Systems for Video Technology*, 2021.
- [20] Y. Wu, H. Qiu, J. Wen, and R. Feng, "OSD: An occlusion skeleton dataset for action recognition," in *Big Data*, 2020.
- [21] R. Memmesheimer, S. Häring, N. Theisen, and D. Paulus, "Skeleton-DML: Deep metric learning for skeleton-based one-shot action recognition," WACV, 2022.
- [22] R. Memmesheimer, N. Theisen, and D. Paulus, "SL-DML: Signal level deep metric learning for multimodal one-shot action recognition," in *ICPR*, 2021.
- [23] Y. Zou, Y. Shi, Y. Wang, Y. Shu, Q. Yuan, and Y. Tian, "Hierarchical temporal memory enhanced one-shot distance learning for action recognition," in *ICME*, 2018.
- [24] Y. Zou, Y. Shi, D. Shi, Y. Wang, Y. Liang, and Y. Tian, "Adaptationoriented feature projection for one-shot action recognition," *IEEE Transactions on Multimedia*, 2020.
- [25] A. Arnab, M. Dehghani, G. Heigold, C. Sun, M. Lučić, and C. Schmid, "ViVit: A video vision transformer," in *ICCV*, 2021.
- [26] B. Graham et al., "LeViT: a vision transformer in ConvNet's clothing for faster inference," in ICCV, 2021.
- [27] W. Xue and W. Wang, "One-shot image classification by learning to restore prototypes," in AAAI, 2020.
- [28] S. Tsutsui, Y. Fu, and D. Crandall, "Meta-reinforced synthetic data for one-shot fine-grained visual recognition," in *NeurIPS*, 2019.
- [29] S. X. Hu, D. Li, J. Stühmer, M. Kim, and T. M. Hospedales, "Pushing the limits of simple pipelines for few-shot learning: External data and fine-tuning make a difference," in *CVPR*, 2022.
- [30] Y. Hu, S. Pateux, and V. Gripon, "Squeezing backbone feature distributions to the max for efficient few-shot learning," *Algorithms*, 2022.
- [31] Y. Bendou *et al.*, "EASY: Ensemble augmented-shot y-shaped learning: State-of-the-art few-shot classification with simple ingredients," *arXiv* preprint arXiv:2201.09699, 2022.
- [32] J. Zhang, K. Yang, A. Constantinescu, K. Peng, K. Müller, and R. Stiefelhagen, "Trans4Trans: Efficient transformer for transparent object segmentation to help visually impaired people navigate in the real world," in *ICCVW*, 2021.
- [33] X. Wang *et al.*, "Semantic-guided relation propagation network for fewshot action recognition," in *MM*, 2021.
- [34] A. Sabater, L. Santos, J. Santos-Victor, A. Bernardino, L. Montesano, and A. C. Murillo, "One-shot action recognition in challenging therapy scenarios," in *CVPRW*, 2021.
- [35] J. Liu, A. Shahroudy, D. Xu, A. C. Kot, and G. Wang, "Skeletonbased action recognition using spatio-temporal LSTM network with trust gates," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.
- [36] J. Liu, G. Wang, P. Hu, L.-Y. Duan, and A. C. Kot, "Global contextaware attention LSTM networks for 3D action recognition," in *CVPR*, 2017.
- [37] A. Vaswani et al., "Attention is all you need," in NeurIPS, 2017.
- [38] A. Dosovitskiy *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," in *ICLR*, 2021.
- [39] H. Touvron, M. Cord, M. Douze, F. Massa, A. Sablayrolles, and H. Jégou, "Training data-efficient image transformers & distillation through attention," in *ICML*, 2021.
- [40] X. Chu *et al.*, "Twins: Revisiting the design of spatial attention in vision transformers," in *NeurIPS*, 2021.
  [41] Z. Liu *et al.*, "Swin transformer: Hierarchical vision transformer using
- [41] Z. Liu *et al.*, "Swin transformer: Hierarchical vision transformer using shifted windows," in *ICCV*, 2021.
- [42] H. Touvron, M. Cord, A. Sablayrolles, G. Synnaeve, and H. Jégou, "Going deeper with image transformers," in *ICCV*, 2021.
- [43] Q. Zhang and Y. Yang, "ResT: An efficient transformer for visual recognition," in *NeurIPS*, 2021.
- [44] S. Li et al., "GroupFormer: Group activity recognition with clustered spatial-temporal transformer," in ICCV, 2021.
- [45] Y. Zhang *et al.*, "VidTr: Video transformer without convolutions," in *ICCV*, 2021.
- [46] K. Peng, A. Roitberg, K. Yang, J. Zhang, and R. Stiefelhagen, "Trans-DARC: Transformer-based driver activity recognition with latent space feature calibration," in *IROS*, 2022.
- [47] F. Shi et al., "STAR: Sparse transformer-based action recognition," arXiv preprint arXiv:2107.07089, 2021.
- [48] H.-S. Fang, S. Xie, Y.-W. Tai, and C. Lu, "RMPE: Regional multi-person pose estimation," in *ICCV*, 2017.

- [49] J. Li, C. Wang, H. Zhu, Y. Mao, H.-S. Fang, and C. Lu, "CrowdPose: Efficient crowded scenes pose estimation and a new benchmark," in *CVPR*, 2019.
- [50] Y. Xiu, J. Li, H. Wang, Y. Fang, and C. Lu, "Pose Flow: Efficient online pose tracking," in *BMVC*, 2018.
- [51] D. Li and W. Shi, "Partially occluded skeleton action recognition based on multi-stream fusion graph convolutional networks," in CGI, 2021.
- [52] X. Ding, S. Zhu, W. Qu, and W. Chen, "Generalized graph convolutional networks for action recognition with occluded skeletons," in *ICCPR*, 2020.
- [53] M. Ghafoor and A. Mahmood, "Quantification of occlusion handling capability of 3D human pose estimation framework," *IEEE Transactions* on *Multimedia*, 2022.
- [54] F. Angelini, Z. Fu, Y. Long, L. Shao, and S. M. Naqvi, "2D posebased real-time human action recognition with occlusion-handling," *IEEE Transactions on Multimedia*, 2020.
- [55] M. Liu, H. Liu, and C. Chen, "Robust 3D action recognition through sampling local appearances and global distributions," *IEEE Transactions* on *Multimedia*, 2018.
- [56] A. Shahroudy, J. Liu, T.-T. Ng, and G. Wang, "NTU RGB+D: A large scale dataset for 3D human activity analysis," in *CVPR*, 2016.
- [57] Y. Cui, C. Jiang, L. Wang, and G. Wu, "MixFormer: End-to-end tracking with iterative mixed attention," in CVPR, 2022.
- [58] C.-W. Kuo, C.-Y. Ma, J.-B. Huang, and Z. Kira, "FeatMatch: Featurebased augmentation for semi-supervised learning," in *ECCV*, 2020.
- [59] A. Krizhevsky, "One weird trick for parallelizing convolutional neural networks," arXiv preprint arXiv:1404.5997, 2014.
- [60] F. N. Iandola *et al.*, "SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and <0.5MB model size," *arXiv preprint arXiv:1602.07360*, 2016.
- [61] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in CVPR, 2016.
- [62] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *ICLR*, 2019.
- [63] L. Van der Maaten and G. Hinton, "Visualizing data using t-SNE," Journal of Machine Learning Research, 2008.