

# SRNet: Improving Generalization in 3D Human Pose Estimation with a Split-and-Recombine Approach

Ailing Zeng<sup>1,2\*</sup>, Xiao Sun<sup>2</sup>, Fuyang Huang<sup>1</sup>, Minhao Liu<sup>1</sup>, Qiang Xu<sup>1</sup>, and Stephen Lin<sup>2</sup>

<sup>1</sup> The Chinese University of Hong Kong

<sup>2</sup> Microsoft Research Asia

alzeng, fyhuang, mhliu, qxu@cse.cuhk.edu.hk  
xias, stevelin@microsoft.com

**Abstract.** Human poses that are rare or unseen in a training set are challenging for a network to predict. Similar to the long-tailed distribution problem in visual recognition, the small number of examples for such poses limits the ability of networks to model them. Interestingly, *local* pose distributions suffer less from the long-tail problem, i.e., local joint configurations within a rare pose may appear within other poses in the training set, making them less rare. We propose to take advantage of this fact for better generalization to rare and unseen poses. To be specific, our method splits the body into local regions and processes them in separate network branches, utilizing the property that a joint’s position depends mainly on the joints within its local body region. Global coherence is maintained by recombining the global context from the rest of the body into each branch as a low-dimensional vector. With the reduced dimensionality of less relevant body areas, the training set distribution within network branches more closely reflects the statistics of *local* poses instead of global body poses, without sacrificing information important for joint inference. The proposed split-and-recombine approach, called *SRNet*, can be easily adapted to both single-image and temporal models, and it leads to appreciable improvements in the prediction of rare and unseen poses.

**Keywords:** Human pose estimation, 2D to 3D, long-tailed distribution

## 1 Introduction

Human pose estimation is a longstanding computer vision problem with numerous applications, including human-computer interaction, augmented reality, and computer animation. For predicting 3D pose, a common approach is to first estimate the positions of keypoints in the 2D image plane, and then lift these keypoints into 3D. The first step typically leverages the high performance of

---

\* The work is done when Ailing Zeng is an intern at Microsoft Research Asia.

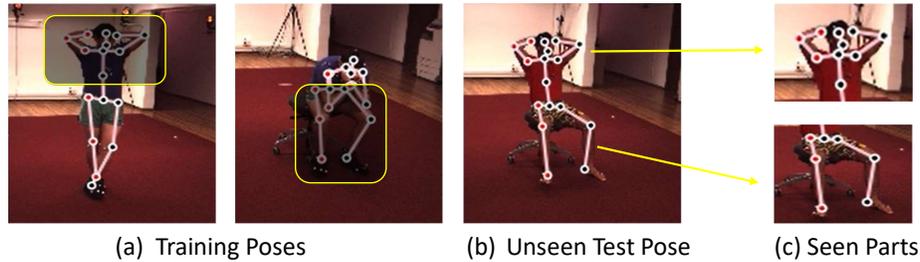


Fig. 1: An unseen test pose (b) may be decomposed into local joint configurations (c) that appear in poses that exist in the training set (a). Our method takes advantage of this property to improve estimation of rare and unseen poses.

existing 2D human pose estimation algorithms [4, 21, 28]. For the second stage, a variety of techniques have been proposed, based on structural or kinematic body models [2, 6, 7, 14, 31], learned dependencies and relationships among body parts [7, 22], and direct regression [18].

Besides algorithm design, an important factor in the performance of a machine learning system is its training data. A well-known issue in pose estimation is the difficulty of predicting poses that are rare or unseen in the training set. Since few examples of such poses are available for training, it is hard for the network to learn a model that can accurately infer them. Better generalization to such poses has been explored by augmenting the training set with synthetically generated images [27, 3, 29, 32, 19, 12]. However, the domain gap that exists between real and synthesized images may reduce their efficacy. Different viewpoints of existing training samples have also been simulated to improve generalization to other camera positions [7], but this provides only a narrow range of pose augmentations.

In this work, we propose to address the rare/unseen pose problem through a novel utilization of the data in the original training set. Our approach is based on the observation that rare poses at the global level are composed of local joint configurations that are generally less rare in the training set. For example, a bicycling pose may be uncommon in the dataset, but the left leg configuration of this pose may resemble the left legs of other poses in the dataset, such as stair climbing and marching. Many instances of a local pose configuration may thus exist in the training data among different global poses, and they could be leveraged for learning local pose. Moreover, it is possible to reconstruct unseen poses as a combination of local joint configurations that are presented in the dataset. For example, an unseen test pose may be predicted from the upper body of one training pose and the lower body of another, as illustrated in Fig. 1.

Based on this observation, we design a network structure that splits the human body into local groups of joints that have strong inter-relationships within each group and relatively weak dependencies on joints outside the group. Each

group is processed in a separate network branch. To account for the weak dependencies that exist with the rest of the body, low-dimensional global context is computed from the other branches and recombined into the branch. The dimensionality reduction of less-relevant body areas within the global context decreases their impact on the feature learning for local joints within each branch. At the same time, accounting for some degree of global context can avoid local pose estimates that are incoherent with the rest of the body. With this split-and-recombine approach, called *SRNet*, generalization performance is enhanced to effectively predict global poses that are rare or absent from the training set.

In extensive comparisons to state-of-the-art techniques, SRNet exhibits competitive performance on single-frame input and surpasses them on video input. More importantly, we show that SRNet can elevate performance considerably on rare/unseen poses. Moreover, we conduct various ablation studies to validate our approach and to examine the impact of different design choices.

## 2 Related Work

Extensive research has been conducted on reconstructing 3D human pose from 2D joint predictions. In the following, we briefly review methods that are closely related to our approach.

***Leveraging Local Joint Relations*** Many recent works take advantage of the physical connections that exist between body joints to improve feature learning and joint prediction. As these connections create strong inter-dependencies and spatial correlations between joints, they naturally serve as paths for information sharing [14, 35, 2, 5] and for encoding kinematic [7] and anatomical [6, 8, 23, 31] constraints.

The structure of these connections defines a locality relationship among joints. More closely connected joints have greater inter-dependency, while distantly connected joints have less and indirect dependence via the joints that lie on the paths between them. These joint relationships have been modeled hierarchically, with feature learning that starts within local joint groups and then expands to account for all the joint groups together at a global level [22]. Alternatively, these relationships have been represented in a graph structure, with feature learning and pose estimation conducted in a graph convolutional network (GCN) [35, 2, 5]. Within a GCN layer, dependencies are explicitly modeled between connected joints and expand in scope to more distant joints through the stacking of layers.

For both hierarchical and GCN based techniques, feature learning within a local group of joints can be heavily influenced by joints outside the group. In contrast, the proposed approach SRNet restricts the impact of non-local joints on local feature learning through dimensionality reduction of the global context, which facilitates the learning of local joint configurations without ignoring global pose coherence. By doing so, SRNet can achieve better generalization to rare and unseen poses in the training data.

**Generalization in Pose Estimation** A common approach to improve generalization in pose estimation is to generate more training images through data augmentation. Approaches to augmentation have included computer graphics rendering [3, 29], image-based synthesis [27], fitting 3D models to people and deforming them to generate new images [25], and changing the background, clothing and occluders in existing images [19]. These methods can produce a large amount of data for training, but the gap in realism between artificially constructed images and actual photographs may limit their effectiveness. Since our technique obtains improved training data distributions by focusing on local pose regions instead of global poses, it does not involve image manipulation or synthesis, thereby maintaining the realism of the original training set.

Other works have explored generalization to different viewpoints [7, 30] and to in-the-wild scenes [34, 8], which are orthogonal to our goal of robustly estimating rare and unseen poses.

**Robustness to Long-tailed Distributions** In visual recognition, there exists a related problem of long-tailed training set distributions, where classes in the distribution tail have few examples. As these few examples have little impact on feature learning, the recognition performance for tail classes is often poor. Recent approaches to this problem include metric learning to enforce inter-class margins [9] and meta-learning that learns to regress many-shot model parameters from few-shot model parameters [33]. Such techniques improve the discriminability of tail classes from other classes, but they are not compatible with the problem of keypoint localization in pose estimation.

### 3 Method

To address the issue of rare and unseen poses, our approach is to decompose global pose estimation into a set of local pose estimation problems. For this

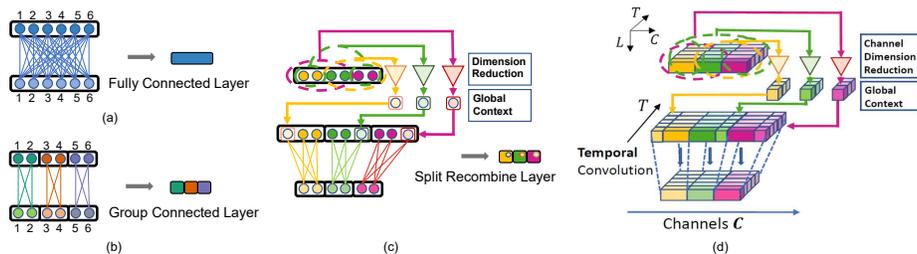


Fig. 2: Illustration of (a) a fully connected layer, (b) group connected layer, (c) our split-and-recombine layer, and (d) our convolution layer for temporal models. The four types of layers can be stacked to form different network structures as shown in Figure 3.

strategy to be effective, the local problems must be defined in a manner that allows the feature learning of a local region to primarily reflect the statistical distribution of its local poses, yet account for other local regions such that the final overall pose estimate is globally coherent.

In the following, we start with the most common baseline for lifting 2D joints to 3D, and then describe how to modify this baseline to follow our strategy.

**Fully-Connected Network Baseline** For lifting 2D keypoints to 3D, a popular yet simple baseline is to use a fully-connected network (FCN) consisting of several layers [18]. Given the 2D keypoint detections  $\mathcal{K} = \{K_i | i, \dots, N\} \in \mathbb{R}^{2N}$  in the image coordinate system, the FCN estimates the 3D joint locations  $\mathcal{J} = \{J_i | i, \dots, N\} \in \mathbb{R}^{3N}$  in the camera coordinate system with the origin at the root joint  $J_0$ . Formally, a fully-connected network layer can be expressed as

$$\mathbf{f}^{l+1} = \Theta^l \mathbf{f}^l \quad (1)$$

where  $\Theta^l \in \mathbb{R}^{D^{l+1} \times D^l}$  is the fully-connected weight matrix to be learned, and  $D^l$  is the feature dimension for the  $l^{th}$  layer, namely  $\mathbf{f}^l \in \mathbb{R}^{D^l}$ . Batch normalization and ReLU activation are omitted for brevity. For the input and output layers, their feature dimensions are  $D^1 = 2N$  and  $D^{L+1} = 3N$ , respectively.

It can be noted that in this FCN baseline, each output joint and each intermediate feature is connected to all of the input joints indiscriminately, allowing the prediction of an output joint to be overfitted to the positions of distant joints with little relevance. In addition, all the output joints share the same set of features entering the final layer and have only this single linear layer ( $\Theta_i^L$ ) to determine a solution particular to each joint.

**Body Partitioning into Local Pose Regions** In turning global pose estimation into several local pose estimation problems, a suitable partitioning of the human body into local pose regions is needed. A local pose region should contain joints whose positions are heavily dependent on one another, but less so on joints outside the local region.

Here, we adopt the partitioning used in [22], where the body is divided into left/right arms, left/right legs, and torso. These parts have distinctive and coordinated behaviors such that the joint positions within each group are highly correlated. In contrast, joint positions between groups are significantly less related.

To accommodate this partitioning, the FCN layers are divided into groups. Formally,  $\mathcal{G}_g^l$  represents the feature/joint indexes of the  $g^{th}$  group at layer  $l$ . Specifically, for the first input layer,

$$\mathcal{G}_0^1 = [\text{joint indices of the right arm}], \quad (2)$$

and for the intermediate feature layers, we have

$$\mathcal{G}_0^l = [\text{feature indices of the right leg at the } l_{th} \text{ layer}]. \quad (3)$$

Then, a *group connected layer* for the  $g^{th}$  group is expressed as

$$\mathbf{f}^{l+1}[\mathcal{G}_g^{l+1}] = \Theta_g^l \mathbf{f}^l[\mathcal{G}_g^l]. \quad (4)$$

In a *group connected layer*, the connections between joint groups are removed. This “local connectivity” structure is commonly used in convolutional neural networks to capture spatially local patterns [13]. The features learned in a group depend on the statistical distribution of the training data for only its local pose region. In other words, this local feature is learned independently of the pose configurations of other parts, and thus it generalizes well to any global pose that includes similar local joint configurations.

However, a drawback of group connected layers is that the status of the other body parts is completely unknown when inferring the local pose. As a result, the set of local inferences may not be globally coherent, leading to low performance. There is thus a need to account for global information while largely preserving local feature independence.

***SRNet: Incorporating Low-Dimensional Global Context*** To address this problem, we propose to incorporate *Low-Dimensional Global Context* (LDGC) in a group connected layer. It coarsely represents information from the less relevant joints, and is brought back to the local group in a manner that limits disruption to the local pose modeling while allowing the local group to account for non-local dependencies. This split-and-recombine approach for global context can be expressed as the following modification of Eq. 4:

$$\mathbf{f}^{l+1}[\mathcal{G}_g^{l+1}] = \Theta_g^l (\mathbf{f}^l[\mathcal{G}_g^l] \circ \mathcal{M} \mathbf{f}^l[\mathcal{G}^l \setminus \mathcal{G}_g^l]) \quad (5)$$

where  $\mathbf{f}^l[\mathcal{G}^l \setminus \mathcal{G}_g^l]$  is the global context for the  $g^{th}$  group.  $\mathcal{M}$  is a mapping function that defines how the global context is represented. Special cases of the mapping function are  $\mathcal{M} = Identity$ , equivalent to a fully-connected layer, and  $\mathcal{M} = Zero$ , which is the case for a group connected layer. The mapped global context is recombined with local features by an operator  $\circ$ , typically concatenation for an *FCN*.

The mapping function  $\mathcal{M}$  acts as a gate controlling the information passed from the non-local joints. If the gate is wide open (*FCN*), the local feature learning will account for the exact joint positions of other body parts, weakening the ability to model the local joint configurations of rare global poses. However, if the gate is closed (*Group*), the local feature learning receives no knowledge about the state of other body parts and may lose global coherence. We argue that the key to achieving high performance and strong generalization to rare poses is to learn a low-dimensional representation of the global context, namely

$$\mathcal{M} = \Gamma_g^l \in \mathbb{R}^{H \times (D^l - D_g^l)} \quad (6)$$

where  $D_g^l$  is the feature dimension for the  $g^{th}$  group at layer  $l$ .  $\Gamma_g^l$  is a weight matrix that maps the global context  $\mathbf{f}^l[\mathcal{G}^l \setminus \mathcal{G}_g^l]$  of dimensions  $D^l - D_g^l$  to a small number of dimensions  $H$ .

In our experiments, we empirically evaluate different design choices for the mapping function  $\mathcal{M}$  and the combination operator  $\circ$ .

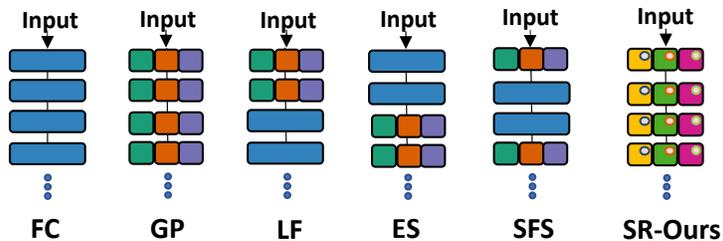


Fig. 3: Illustration of fully connected (FC), group connected (GP), late fusion (LF), early split (ES), split-fuse-split (SPS) and split-and-recombine (SR) models. The components of each layer are shown in Figure 2.

**Network Structure** With our split-and-recombine approach for processing local pose regions and global context, the *SRNet* model can be illustrated as shown at the right side of Fig. 3. This network follows the structure of a group connected network, with the body joints split into separate groups according to their local pose region. Added to this network are connections to each group in a layer from the other groups in the preceding layer. As illustrated in Fig. 2 (c,d), these connections representing the global context reduce the dimensionality of the context features and recombine them with the local features in each group. These inputs are mapped into outputs of the original feature dimensions.

*SRNet* balances feature learning dependency between the most related local region and the less related global context. For demonstrating the efficacy of this split-and-recombine strategy, we describe several straightforward modifications to the FC network, illustrated in Fig. 3, that will serve as baselines in our experiments:

- *Late Fusion (LF)*, which postpones feature sharing among groups by cutting their connections in the earlier layers. This structure is similar to that used in [22], which first learns features within each group and fuses them later with stacked fully connected layers. Formally, the *LF* baseline is defined as

$$\mathbf{f}^{l+1} = \begin{cases} \text{Cat}\{\mathbf{f}^l[\mathcal{G}_g^l] | g = 1, \dots, G\}, & \text{if } l < L_{fuse} \\ \Theta^l \mathbf{f}^l, & \text{otherwise.} \end{cases} \quad (7)$$

- *Early Split (ES)*, which aims to provide more expressive features for each group, by cutting the connections between groups in the latter layers. Formally,

$$\mathbf{f}^{l+1} = \begin{cases} \Theta^l \mathbf{f}^l, & \text{if } l < L_{split} \\ \text{Cat}\{\mathbf{f}^l[\mathcal{G}_g^l] | g = 1, \dots, G\}, & \text{otherwise.} \end{cases} \quad (8)$$

- *Group Connected (GP)*, which is the standard group connected network. *Late Fusion* degenerates to this form when  $L_{fuse} = L$ , and *Early Split* does when  $L_{split} = 1$ .

- *Split-Fuse-Split (SFS)*, where the middle  $L_{link}$  layers of the network are fully-connected and the connections between joint groups are cut in the first and last layers.

The differences among FCN, these baselines, and SRNet are summarized in Table 1. The table also highlights differences of SRNet from GCN [35] and its modification LCN [5], specifically in controlling the feature dependencies between local regions and global context.

Method	local inference	local features	info passing	dimension control
FCN	weak(share feature)	no	yes	no
Group	yes	yes	no	no
Late Fusion [22]	weak(share feature)	yes	yes	no
Early Split	yes	no	yes	no
Split-Fuse-Split	yes	yes	yes	no
GCN [35]	no(share weight)	yes(overlapped group)	yes	no
LCN [5]	yes	yes(overlapped group)	yes	no
Ours (SRNet)	yes	yes	yes	<b>yes</b>

Table 1: Different network structures used for 2D to 3D pose estimation.

***SR Convolution for Temporal Models*** Temporal information in 2D to 3D mapping is conventionally modeled with Recurrent Neural Networks (RNNs) [2, 14]. Recently, Pavllo et al. [23] introduce an alternative temporal convolutional model that stacks all the spatial information of a frame into the channel dimensions and replaces fully-connected operations in the spatial domain by convolutions in the temporal domain. The temporal convolution enables parallel processing of multiple frames and brings greater accuracy, efficiency, and simplicity.

As illustrated in Fig. 2 (d), our split-and-recombine modification to the fully connected layer can be easily adapted to the temporal convolution model by applying the same split-and-recombine strategy to the channels during convolution. Specifically, the group connected operations for local joint groups are replaced with corresponding temporal convolutions for local feature learning, where the features from other joint groups undergo an extra convolution for dimension reduction in channels and are concatenated back with the local joint group as global context. We call this split-and-recombine strategy in the channel dimensions during convolution the *SR convolution*.

## 4 Datasets and Rare-Pose Evaluation Protocols

### 4.1 Datasets and Evaluation Metrics

Our approach is validated on two popular benchmark datasets:

- *Human3.6M* [11] is a large benchmark widely-used for 3D human pose estimation. It consists of 3.6 million video frames from four camera viewpoints with 15

activities. Accurate 3D human joint locations are obtained from motion capture devices. Following convention, we use the mean per joint position error (MPJPE) for evaluation, as well as the Procrustes Analysis MPJPE (PA-MPJPE).

– *MPI-INF-3DHP* [19, 20] is a more challenging 3D pose dataset, containing not only constrained indoor scenes but also complex outdoor scenes. Compared to Human3.6M, it covers a greater diversity of poses and actions. For evaluation, we follow common practice [31, 5, 32, 34] by using the Percentage of Correct Keypoints (PCK) with a threshold of 150mm and the Area Under Curve (AUC) for a range of PCK thresholds.

## 4.2 Evaluation Protocols

The most widely-used evaluation protocol [18, 32, 10, 23, 5] on Human3.6M uses five subjects (S1, S5, S6, S7, S8) as training data and the other two subjects (S9, S11) as testing data. We denote it as the *Subject Protocol*. However, the rare pose problem is not well examined in this protocol since each subject is asked to perform a fixed series of actions in a similar way.

To better demonstrate the generalization of our method to rare/unseen poses, we use two other protocols, introduced in [5]. The first is the *Cross Action Protocol* that trains on only one of the 15 actions in the Human3.6M dataset and tests on all actions. The second is the *Cross Dataset Protocol* that applies the model trained on Human3.6M to the test set of MPI-INF-3DHP. In addition, we propose a new protocol called *Rare Pose Protocol*.

*Rare Pose Protocol*. In this protocol, we use all the poses of subjects S1, S5, S6, S7, S8 for training but only use a subset of poses from subjects S9, S11 for testing. The subset is selected as the rarest poses in the testing set. To identify rare poses, we first define pose similarity ( $PS$ ) as

$$PS(\mathcal{J}, \mathcal{I}) = \frac{1}{N} \sum_i^N \exp\left(-\frac{\|J_i - I_i\|^2}{2\sigma_i^2}\right) \quad (9)$$

where  $\|J_i - I_i\|$  is the Euclidean distance between corresponding joints of two poses. To compute  $PS$ , we pass  $\Delta J_i$  through an unnormalized Gaussian with a standard deviation  $\sigma_i$  controlling falloff. This yields a pose similarity that ranges between 0 and 1. Perfectly matched poses will have  $PS = 1$ , and if all joints are distant by more than the standard deviation  $\sigma_i$ ,  $PS$  will be close to 0. The occurrence of a pose  $\mathcal{J}$  in a pose set  $\mathbb{O}$  is defined as the average similarity between itself and all the poses in  $\mathbb{O}$ :

$$OCC_{\mathbb{O}}(\mathcal{J}) = \frac{1}{M} \sum_{\mathcal{I}}^{\mathbb{O}} PS(\mathcal{J}, \mathcal{I}) \quad (10)$$

where  $M$  is the total number of poses in the pose set  $\mathbb{O}$ . A rare pose will have a low occurrence value with respect to other poses in the set. We select the  $R\%$  of poses with the lowest occurrence values in the testing set for evaluation.

## 5 Experiments

### 5.1 Ablation Study

All our ablation studies are evaluated on the Human3.6M dataset. Both the *Subject Protocol* and our new *Rare Pose Protocol* are used for evaluation. MPJPE serves as the evaluation metric. *Please refer to the supplementary material for details on data pre-processing and training settings.*

**Effect of the SRNet** In Table 2, we compare our *SR* network with the baseline networks *FC*, *GP*, *LF*, *ES*, *SFS* described in Section 3. By default, the mapped dimensions  $H$  is set to 1 and the combination operator  $\circ$  is set to multiplication.

It is found that both the *LF* and *ES* baselines are superior to the *FC* and *GP* baselines. A more detailed comparison is shown in Figure 4, with different splitting configurations for each. In Figure 4, both *LF* and *ES* start on the left with the *FC* network ( $L_{fuse} = 0$  and  $L_{split} = 6$ ). When the fully connected layers are gradually replaced by group connected layers, from the beginning of the network for *LF* ( $L_{fuse} \uparrow$ ) or the end for *ES* ( $L_{split} \downarrow$ ), testing error decreases at first. This indicates that *local feature learning* and *local inference* are each helpful. However, replacing all the fully connected layers with group layers leads to a sharp performance drop. This indicates that a local joint region still needs information from the other regions to be effective.

Combining the *LF* and *ES* baselines to create *SFS* yields further improvements from sharing the merits of both. These two simple modifications already improve upon *FC* by a large margin. Specifically, it improves upon *FC* by 7.4mm (relative 15.8%), 16.8mm (relative 22.1%) and 20.9mm (relative 23.5%) on the *Subject*, *Rare Pose 20%* and *Rare Pose 10%* protocols, respectively.

*SRNet* with its use of low-dimensional global context performs significantly better than the *SFS* baseline. Specifically, it improves upon *SFS* by 2.8mm (relative 7.1%), 10.6mm (relative 17.9%) and 14.5mm (relative 21.2%) on the *Subject*, *Rare Pose 20%* and *Rare Pose 10%* protocols, respectively. Note that the improvement is especially large on rare poses, indicating the stronger generalization ability of *SRNet* to such poses. Figure 5 breaks down the results of the *FC*, *SFS* and *SR* networks for different degrees of pose rareness. It can be seen that the performance improvements of *SRNet* increase for rarer poses.

Protocol	GP	FC	LF	ES	SFS	Ours (SR)
Subject (100%)	62.7	<u>46.8</u>	39.8	42.0	<u>39.4</u> ↓7.4	<b>36.6</b> ↓2.8(7.1%)
Rare Pose (20%)	89.1	<u>76.0</u>	60.1	62.5	<u>59.2</u> ↓16.8	<b>48.6</b> ↓10.6(17.9%)
Rare Pose (10%)	98.9	<u>89.1</u>	69.9	73.5	<u>68.2</u> ↓20.9	<b>53.7</b> ↓14.5(21.3%)

Table 2: Comparing the *SR* network to different baseline networks under the *Subject* protocol and the *Rare Pose* protocol (with 10% and 20% of the rarest poses). MPJPE is used as the evaluation metric. The improvements of *SFS* from *FC*, and of *SR* from *SFS*, are shown as subscripts.



Fig. 4: Training loss and test error (MPJPE in mm) of the baseline networks (a) *LF*, (b) *ES*, and (c) *SFS* with respect to different splitting configurations. Rare Pose Protocol 20% is used for evaluation.

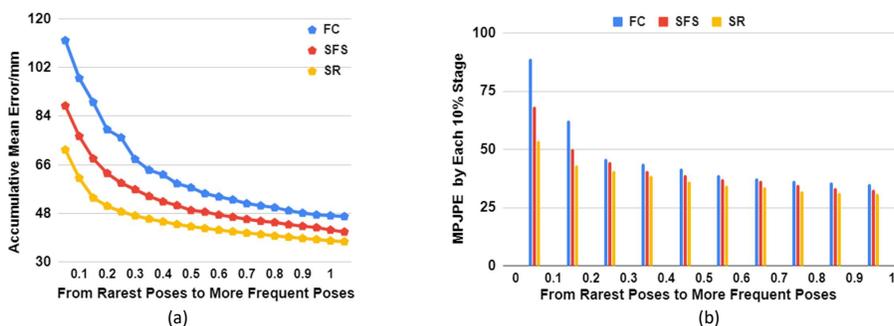


Fig. 5: Testing errors of *FC*, *SFS* and our *SR* networks on poses with different degrees of rareness. The horizontal axis represents the top  $R\%$  of rarest poses. (a) Cumulative mean error. (b) Mean error at 10% intervals.

**Low-Dimensional Representation for Global Context** Table 3 presents the performance of SRNet with different dimensions  $H$  for the global context. To readily accommodate different dimensions, the combination operator  $\circ$  is set to concatenation. When  $H = 0$ , *SRNet* degenerates into the *GP* baseline. It can be seen that the *GP* baseline without any global context obtains the worst performance. This shows the importance of bringing global context back to local groups. Also, it is found that a lower dimension of global context leads to better results. The best result is achieved by keeping *only a single dimension* for global context. This indicates that a low-dimensional representation of the global context can better account for the information of less relevant joints while preserving effective local feature learning.

**Effect of Grouping Strategy** We compare the results of using different numbers of local groups in Table 4. The corresponding sets of joints for each group are shown in Figure 6. More joint groups leads to greater local correlation of joints within a group. It is shown that the performance first improves with more groups, especially in the *Rare Pose* case. However, with more than five groups,

Dimension $H$	0(GP)	1	3	10	25% $D_g^l$	50% $D_g^l$	100% $D_g^l$	200% $D_g^l$
Subject (100%)	62.7	<b>38.3</b>	41.9	42.9	44.6	45.3	45.7	45.8
Rare Pose (20%)	89.1	<b>49.1</b>	52.2	53.6	56.5	63.8	70.1	78.5

Table 3: Effect of using different dimension  $H$  for the proposed global context. 100% corresponds to the same dimensionality as the local feature.

Group Num	1(FC)	2	3	5	6	8	17	Shuffled Groups	5	4	3	2	0(Ours)
Subj.(100%)	46.8	41.4	37.7	<b>36.6</b>	41.1	46.7	91.5	Subject (100%)	53.8	49.7	45.9	43.4	<b>36.6</b>
Rare(20%)	76.0	55.9	51.0	<b>48.6</b>	55.6	60.6	115.3						

Table 4: Mean testing error with different group numbers for *Subject* and *Rare Pose* 20% protocols.

Table 5: Mean testing error with respect to the number of shuffled groups (5 in total).

the performance drops due to weakening of local features when there are fewer than three joints in a group. Moreover, to show that separation into *local* groups is important, we randomly shuffle the joints among a subset of groups, with five groups in total, as illustrated in Figure 6 (f). In Table 5, it is shown that the performance decreases when randomly shuffling joints among more groups. This indicates that a strong physical relationship among joints in a group is a prerequisite for learning effective local features.

## 5.2 Comparison with State-of-The-Art Methods

We compare with the state of the art in three different experiment settings: cross-dataset on MPI-INF-3DHP, single-frame model on Human3.6M, and temporal model on Human3.6M. MPJPE is used as the evaluation metric on Human3.6M. *Please refer to the supplementary material for more results that use PA-MPJPE as the evaluation metric.*

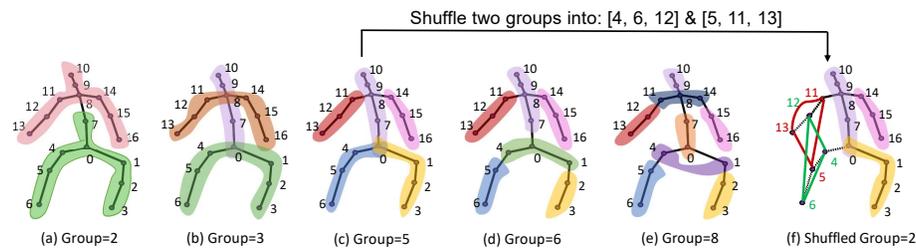


Fig. 6: Different numbers of groups and their corresponding sets of joints. Joints in the same group are shown in the same color. A random shuffle setting is also shown and evaluated in Table 5.

**Cross-Dataset Results on MPI-INF-3DHP.** In this setting, we follow common practice [16, 32, 5, 34, 36] by training the model on the Human3.6M training set and evaluating on the MPI-INF-3DHP test set to examine cross-dataset generalization ability. PCK and AUC are used as evaluation metrics, for which higher is better. In Table 6, our approach achieves the best cross-data generalization performance. It improves upon the state-of-the-art [5] by a large margin (4.9% on PCK and 19.3% on AUC), indicating superior generalization ability.

Method	Martinez [18]	Mehta [19]	Luo [16]	Biswas [1]	Yang [34]	Zhou [36]	Wang [32]	Ci [5]	Ours
Outdoor	31.2	58.8	65.7	67.4	-	72.7	-	77.3	<b>80.3</b> $\uparrow$ 3.9%
PCK	42.5	64.7	65.6	65.8	69.0	69.2	71.2	74.0	<b>77.6</b> $\uparrow$ 4.9%
AUC	17.0	31.7	33.2	31.2	32.0	32.5	33.8	36.7	<b>43.8</b> $\uparrow$ 19.3%

Table 6: Cross-dataset results on MPI-INF-3DHP. All models are trained on Human3.6M and tested on the MPI-INF-3DHP test set.

Method	Direct	Discuss	Eat	Greet	Phone	Photo	Pose	Purcha.	Sit	SitD	Smoke	Wait	WalkD	Walk	WalkT	Avg.
Martinez [18]	127.3	104.6	95.0	116.1	95.5	117.4	111.4	125.0	116.9	93.6	111.0	126.0	131.0	106.3	140.5	114.5
Ci [5]	102.8	88.6	79.1	99.3	80.0	91.5	93.2	89.6	90.4	<b>76.6</b>	89.6	102.1	108.8	90.8	118.9	93.4
Ours	<b>92.3</b>	<b>71.4</b>	<b>71.8</b>	<b>86.4</b>	<b>66.8</b>	<b>79.1</b>	<b>82.5</b>	<b>86.6</b>	<b>88.9</b>	93.4	<b>66.1</b>	<b>83.0</b>	<b>74.4</b>	<b>90.0</b>	<b>97.8</b>	<b>82.0</b>

Table 7: *Cross Action* results compared with *Fully Connected Network* and *Locally Connected Network* on Human3.6M. Smaller values are better.

Method	Luvizon[17]	Martinez[18]	Park[22]	Wang [32]	Zhao[35]	Ci[5]	Pavlo [23]	Cai[2]	Ours
Subject (100%)	64.1	62.9	58.6	58.0	57.6	52.7	51.8	50.6	<b>49.9</b>

Table 8: Comparison on *single-frame 2D pose detection input* in terms of mean per-joint position error (MPJPE). Best result in bold.

**Single-Frame Model Results on Human3.6M.** In this setting, we first compare with the state-of-the-art method [5] using the *Cross Action Protocol*. Table 7 shows that our method yields an overall improvement of 11.4mm (relative 12.2% improvement) over [5] and performs better on 93% of the actions, indicating strong generalization ability on unseen actions. We also compare our model to previous works using the standard *Subject Protocol* in Table 8 and Table 9, which use 2D keypoint detection and 2D ground truth as inputs, respectively. Under the standard protocol, our model surpasses the state-of-the-art

Method	Martinez [18]	Pham [24]	Biswas[1]	Zhao [35]	Wang [32]	Ci [5]	Ours
Subject (100%)	45.5	42.4	42.8	43.8	37.6	36.4	<b>33.9</b>

Table 9: Comparison on *single-frame 2D ground truth pose input* in terms of MPJPE. With ground truth 2D pose as input, the upper bound of these methods is explored.

in each case, namely [2] for 2D keypoint detection and [5] for 2D ground truth as input.

**Temporal Model Results on Human3.6M** In Table 10, our approach achieves the new state-of-the-art with either 2D keypoint detection or 2D ground truth (with  $\nabla$ ) as input. Specifically, SRNet improves upon [23] from 37.2mm to 32.0mm (relative 14.0% improvement) with 2D ground truth input and from 46.8mm to 44.8mm (relative 4.3% improvement) with 2D keypoint detection input. Besides, SRNet has around one fifth parameters 3.61M of [23](16.95M) with 243 frame poses as input.

Method	Direct	Discuss	Eat	Greet	Phone	Photo	Pose	Purcha.	Sit	SitD	Smoke	Wait	WalkD	Walk	WalkT	Avg.
Hossain et al. [26]	48.4	50.77	57.2	55.2	63.1	72.6	53.0	51.7	66.1	80.9	59.0	57.3	62.4	46.6	49.6	58.3
Lee et al. [14]	<b>40.2</b>	49.2	47.8	52.6	50.1	75.0	50.2	43.0	55.8	73.9	54.1	55.6	58.2	43.3	43.3	52.8
Cai et al. [2]	44.6	47.4	45.6	48.8	50.8	59.0	47.2	43.9	57.9	61.9	49.7	46.6	51.3	37.1	39.4	48.8
Pavlo et al. [23]	45.2	46.7	43.3	45.6	48.1	55.1	44.6	44.3	57.3	65.8	47.1	44.0	49.0	32.8	33.9	46.8
Lin et al. [15]	42.5	<b>44.8</b>	<b>42.6</b>	44.2	48.5	57.1	<b>42.6</b>	41.4	56.5	64.5	47.4	43.0	48.1	33.0	35.1	46.6
<i>Ours</i>	46.6	47.1	43.9	<b>41.6</b>	<b>45.8</b>	<b>49.6</b>	46.5	<b>40.0</b>	<b>53.4</b>	<b>61.1</b>	<b>46.1</b>	<b>42.6</b>	<b>43.1</b>	<b>31.5</b>	<b>32.6</b>	<b>44.8</b>
Hossain et al. [26] $\nabla$	35.2	40.8	37.2	37.4	43.2	44.0	38.9	35.6	42.3	44.6	39.7	39.7	40.2	32.8	35.5	39.2
Lee et al. [14] $\nabla$	32.1	36.6	34.3	37.8	44.5	49.9	40.9	36.2	44.1	45.6	35.3	35.9	37.6	30.3	35.5	38.4
Pavlo et al. [23] $\nabla$	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	37.2
<i>Ours</i> $\nabla$	<b>34.8</b>	<b>32.1</b>	<b>28.5</b>	<b>30.7</b>	<b>31.4</b>	<b>36.9</b>	<b>35.6</b>	<b>30.5</b>	<b>38.9</b>	<b>40.5</b>	<b>32.5</b>	<b>31.0</b>	<b>29.9</b>	<b>22.5</b>	<b>24.5</b>	<b>32.0</b>

Table 10: Comparison on *Temporal Pose input* in terms of mean per-joint position error (MPJPE). Below the double line,  $\nabla$  indicates use of 2D ground truth pose as input, which is examined to explore the upper bound of these methods. Best results in bold.

## 6 Conclusion

In this paper, we proposed SRNet, a split-and-recombine approach that improves generalization performance in 3D human pose estimation. The key idea is to design a network structure that splits the human body into local groups of joints and recombines a low-dimensional global context for more effective learning. Experimental results show that SRNet outperforms state-of-the-art techniques, especially for rare and unseen poses.

## References

1. Biswas, S., Sinha, S., Gupta, K., Bhowmick, B.: Lifting 2d human pose to 3d: A weakly supervised approach. In: 2019 International Joint Conference on Neural Networks (IJCNN). pp. 1–9. IEEE (2019)
2. Cai, Y., Ge, L., Liu, J., Cai, J., Cham, T.J., Yuan, J., Thalmann, N.M.: Exploiting spatial-temporal relationships for 3d pose estimation via graph convolutional networks. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2272–2281 (2019)
3. Chen, W., Wang, H., Li, Y., Su, H., Wang, Z., Tu, C., Lischinski, D., Cohen-Or, D., Chen, B.: Synthesizing training images for boosting human 3d pose estimation. In: 2016 Fourth International Conference on 3D Vision (3DV). pp. 479–488. IEEE (2016)
4. Chen, Y., Wang, Z., Peng, Y., Zhang, Z., Yu, G., Sun, J.: Cascaded pyramid network for multi-person pose estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 7103–7112 (2018)
5. Ci, H., Wang, C., Ma, X., Wang, Y.: Optimizing network structure for 3d human pose estimation. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2262–2271 (2019)
6. Dabral, R., Mundhada, A., Kusupati, U., Afaque, S., Sharma, A., Jain, A.: Learning 3d human pose from structure and motion. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 668–683 (2018)
7. Fang, H.S., Xu, Y., Wang, W., Liu, X., Zhu, S.C.: Learning pose grammar to encode human body configuration for 3d pose estimation. In: Thirty-Second AAAI Conference on Artificial Intelligence (2018)
8. Habibie, I., Xu, W., Mehta, D., Pons-Moll, G., Theobalt, C.: In the wild human pose estimation using explicit 2d features and intermediate 3d representations. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 10905–10914 (2019)
9. Huang, C., Li, Y., Loy, C.C., Tang, X.: Learning deep representation for imbalanced classification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2016)
10. Huang, F., Zeng, A., Liu, M., Lai, Q., Xu, Q.: Deepfuse: An imu-aware network for real-time 3d human pose estimation from multi-view image. arXiv preprint arXiv:1912.04071 (2019)
11. Ionescu, C., Papava, D., Olaru, V., Sminchisescu, C.: Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **36**(7), 1325–1339 (jul 2014)
12. Jahangiri, E., Yuille, A.L.: Generating multiple diverse hypotheses for human 3d pose consistent with 2d joint detections. In: Proceedings of the IEEE International Conference on Computer Vision Workshops. pp. 805–814 (2017)
13. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: Advances in neural information processing systems. pp. 1097–1105 (2012)
14. Lee, K., Lee, I., Lee, S.: Propagating lstm: 3d pose estimation based on joint interdependency. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 119–135 (2018)
15. Lin, J., Lee, G.H.: Trajectory space factorization for deep video-based 3d human pose estimation. arXiv preprint arXiv:1908.08289 (2019)

16. Luo, C., Chu, X., Yuille, A.: Orinet: A fully convolutional network for 3d human pose estimation. arXiv preprint arXiv:1811.04989 (2018)
17. Luvizon, D.C., Picard, D., Tabia, H.: 2d/3d pose estimation and action recognition using multitask deep learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5137–5146 (2018)
18. Martinez, J., Hossain, R., Romero, J., Little, J.J.: A simple yet effective baseline for 3d human pose estimation. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2640–2649 (2017)
19. Mehta, D., Rhodin, H., Casas, D., Fua, P., Sotnychenko, O., Xu, W., Theobalt, C.: Monocular 3d human pose estimation in the wild using improved cnn supervision. In: 2017 International Conference on 3D Vision (3DV). pp. 506–516. IEEE (2017)
20. Mehta, D., Sridhar, S., Sotnychenko, O., Rhodin, H., Shafiei, M., Seidel, H.P., Xu, W., Casas, D., Theobalt, C.: Vnect: Real-time 3d human pose estimation with a single rgb camera. ACM Transactions on Graphics (TOG) **36**(4), 1–14 (2017)
21. Newell, A., Yang, K., Deng, J.: Stacked hourglass networks for human pose estimation. In: European conference on computer vision. pp. 483–499. Springer (2016)
22. Park, S., Kwak, N.: 3d human pose estimation with relational networks. arXiv preprint arXiv:1805.08961 (2018)
23. Pavllo, D., Feichtenhofer, C., Grangier, D., Auli, M.: 3d human pose estimation in video with temporal convolutions and semi-supervised training. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 7753–7762 (2019)
24. Pham, H.H., Salmane, H., Khoudour, L., Crouzil, A., Zegers, P., Velastin, S.A.: A unified deep framework for joint 3d pose estimation and action recognition from a single rgb camera. arXiv preprint arXiv:1907.06968 (2019)
25. Pishchulin, L., Jain, A., Andriluka, M., Thorm ahlen, T., Schiele, B.: Articulated people detection and pose estimation: Reshaping the future. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2012)
26. Rayat Imtiaz Hossain, M., Little, J.J.: Exploiting temporal information for 3d human pose estimation. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 68–84 (2018)
27. Rogez, G., Schmid, C.: Mocap-guided data augmentation for 3d pose estimation in the wild. In: Advances in neural information processing systems. pp. 3108–3116 (2016)
28. Sun, K., Xiao, B., Liu, D., Wang, J.: Deep high-resolution representation learning for human pose estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5693–5703 (2019)
29. Varol, G., Romero, J., Martin, X., Mahmood, N., Black, M.J., Laptev, I., Schmid, C.: Learning from synthetic humans. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 109–117 (2017)
30. Véges, M., Varga, V., Lőrincz, A.: 3d human pose estimation with siamese equivariant embedding. Neurocomputing **339**, 194–201 (2019)
31. Wandt, B., Rosenhahn, B.: Repnet: Weakly supervised training of an adversarial reprojection network for 3d human pose estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 7782–7791 (2019)
32. Wang, L., Chen, Y., Guo, Z., Qian, K., Lin, M., Li, H., Ren, J.S.: Generalizing monocular 3d human pose estimation in the wild. arXiv preprint arXiv:1904.05512 (2019)
33. Wang, Y.X., Ramanan, D., Hebert, M.: Learning to model the tail. In: Conference on Neural Information Processing Systems (2017)

34. Yang, W., Ouyang, W., Wang, X., Ren, J., Li, H., Wang, X.: 3d human pose estimation in the wild by adversarial learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5255–5264 (2018)
35. Zhao, L., Peng, X., Tian, Y., Kapadia, M., Metaxas, D.N.: Semantic graph convolutional networks for 3d human pose regression. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3425–3435 (2019)
36. Zhou, X., Huang, Q., Sun, X., Xue, X., Wei, Y.: Towards 3d human pose estimation in the wild: a weakly-supervised approach. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 398–407 (2017)