

3D Human Pose Estimation with Relational Networks

Sungheon Park
sungheonpark@snu.ac.kr

Nojun Kwak
nojunk@snu.ac.kr

Department of Transdisciplinary Studies
Seoul National University
Seoul, Korea

Abstract

In this paper, we propose a novel 3D human pose estimation algorithm from a single image based on neural networks. We adopted the structure of the relational networks in order to capture the relations among different body parts. In our method, each pair of different body parts generates features, and the average of the features from all the pairs are used for 3D pose estimation. In addition, we propose a dropout method that can be used in relational modules, which inherently imposes robustness to the occlusions. The proposed network achieves state-of-the-art performance for 3D pose estimation in Human 3.6M dataset, and it effectively produces plausible results even in the existence of missing joints.

1 Introduction

Human pose estimation (HPE) is a fundamental task in computer vision, which can be adopted to many applications such as action recognition, human behavior analysis, virtual reality and so on. Estimating 3D pose of human body joints from 2D joint locations is an under-constrained problem. However, since human joints are connected by rigid bodies, the search space of 3D pose is limited to the range of joints. Therefore, it is able to learn 3D structures from 2D positions, and numerous studies on 2D-to-3D mapping of human body have been conducted. Recently, Martinez et al. [19] proved that a simple fully connected neural network that accepts raw 2D positions as an input gives surprisingly accurate results. Inspired by this result, we designed a network that accepts 2D positions of joints as inputs and generates 3D positions based on them.

Human body can be divided into arms, legs, a head, and a torso, each of which has distinctive behaviors and movements. We designed the network so that it learns the relations among different body parts. The relational modules for the neural networks proposed in [24] provided a way to learn relations between the components within a neural network architecture. We adopt this relational modules for 3D HPE with a little modification. Specifically, the body joints are divided into several groups, and the relations between them are learned via relational networks. The features from all pairs of groups are averaged to generate the feature vectors which are used for 3D pose regression. We found this simple structure outperforms the baseline which uses a fully connected network. Moreover, we propose a method that can impose robustness to the missing points during the training. The proposed method, named

as relational dropout, randomly drops one of the pair features when they are averaged, which simulates the case that certain groups of joints are missing during the training. To capture the relations among joints within a group, we also designed a hierarchical relational network, which further allows robustness to wrong 2D joint inputs. Lastly, we discovered that the proposed structure of the network modified from [19] and the finetuning schemes improve the performance of HPE. The proposed method achieved state-of-the-art performance in 3D HPE on Human 3.6M dataset [13], and the network can robustly estimate 3D poses even when multiple joints are missing using the proposed relational dropout scheme.

The rest of the paper is organized as follows. Some of the related works are introduced in Section 2 and the proposed method is explained in Section 3. Section 4 shows experimental results and finally, Section 5 concludes the paper.

2 Related Work

Early stage studies on 3D HPE from RGB images used hand-crafted features such as local shape context [10], histogram of gradients [23, 22], or segmentation results [12]. From those features, 3D poses were retrieved via regression using relevance vector machine [10], randomized trees [27], structured SVMs [12], KD-trees [57] or Bayesian non-parametric models [60].

Recent advancements in neural networks boosted the performance of 3D HPE. Li and Chan [17] firstly applied convolutional neural networks (CNNs) to 3D HPE. Since then, various models that capture structured representation of human bodies have been combined with CNNs using, for instance, denoising autoencoders [63], maximum-margin cost function [18], and pose priors from 3D data [9, 15, 28].

It has been proven that 2D pose information acts a crucial role for 3D pose estimation. Park et al. [24] directly propagated 2D pose estimation results to the 3D pose estimation part in a single CNN. Pavlakos et al. [26] proposed a volumetric representation that gradually increases the resolution of the depth from heatmaps of 2D pose. Mehta et al. [20] similarly regressed the position of each coordinate using heatmaps. There are a couple of works that directly regress 3D pose from an image using constraints on human joints [52] or combining weakly-supervised learning [40]. Tome et al. [65] lifted 2D pose heatmaps to 3D pose via probabilistic pose models. Tekin et al. [54] combined features from both RGB images and 2D pose heatmaps which were used for 3D pose estimation.

While 3D pose estimation from images have shown impressive performance, there is another approach that infers a 3D pose directly from the result of 2D pose estimation. It usually has a two-stage procedure: 1) 2D pose estimation using CNNs and 2) 3D pose inference via neural networks using the estimated 2D pose. Chen and Ramanan [5] found that a non-parametric nearest neighbor model that estimates a 3D pose from a 2D pose showed comparable performance when the precise 2D pose information is provided. Moreno-Noguer [21] proposed a neural network that outputs 3D Euclidean distance matrices from 2D inputs. Martinez et al. [19] proposed a simple neural network that directly regresses a 3D pose from raw 2D joint positions. The network consists of two residual modules [8] with batch normalization [11] and dropout [51]. The method showed state-of-the-art performance despite its simple structure. The performance has been further improved by recent works. Fang et al. [6] proposed a pose grammar network that incorporates a set of knowledge learned from human body, which was designed as a bidirectional recurrent neural network. Yang et al. [66] used adversarial learning to implicitly learn geometric configuration of human body. Cha et al. [2]

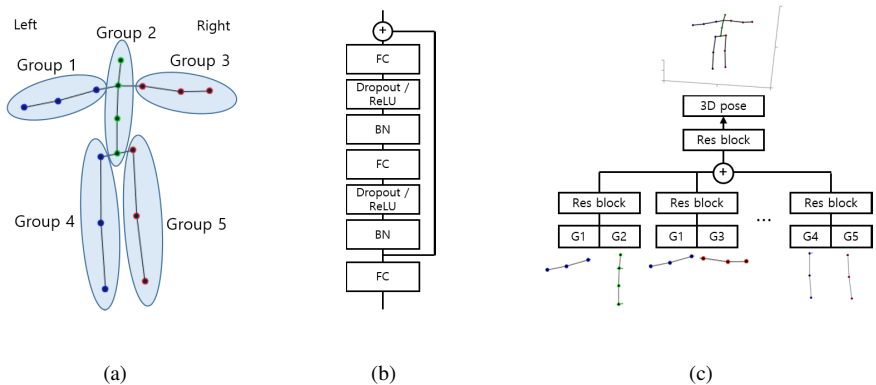


Figure 1: (a) Group configurations used in this paper. We divided 16 2D input joints to non-overlapping 5 groups each of which corresponds to left/right arms, left/right legs and a torso. (b) The residual module used in this paper. We adopted the structure suggested in [9]. (c) The structure of the RN for 3D HPE. Features extracted from all pairs of groups are averaged to produce features for pose estimation. Each Resblock in the figure has the same structure shown in (b).

developed a consensus algorithm that generates a 3D pose from multiple partial hypotheses which are based on a non-rigid structure from motion algorithm [16]. The method is similar to our method in that they divided the body joints into multiple groups. However, our proposed method integrates the features of all groups within the network rather than generating a 3D pose from each group as in Cha et al. [4].

There are a few approaches that exploit temporal information using various methods such as overcomplete dictionaries [68, 69], 3D CNNs [7], sequence-to-sequence networks [10], and multiple-view settings [25]. In this paper, we focus on the case that both training and testing are conducted on a single image.

3 Methods

3.1 Relational Networks for 3D Human Pose Estimation

Relation networks (RN) proposed in [29] consists of two parts, one that does relational reasoning and the other that performs a task-specific inference. The output of the RN is formulated as follows:

$$RN(O) = f\left(\sum_{(i,j)} g(o_i, o_j)\right), \quad (1)$$

where f and g are functions that are represented as corresponding neural networks, and $O = \{o_1, \dots, o_n\}$ is the set of objects. Pairs of different objects o_i, o_j are fed to the network g , and the relation of all pairs are summed together to generate features that capture relational information.

We adopt the concept and the structure of the RN to 3D human pose estimation. The network proposed in this paper takes $2n_{2D}$ -dimensional vectors as inputs and outputs $3(n_{3D} - 1)$ -dimensional vectors where n_{2D} and n_{3D} are the number of 2D and 3D joints respectively.

For 2D inputs, we used (x, y) coordinates of detected joints in RGB images whereas relative positions of (x, y, z) coordinates from the root joint are estimated for 3D pose estimation. In the original RN [29], a neural network module that generates a pairwise relation, $g(\cdot)$, shares weights across all pairs of objects. This weight sharing makes the network learn order-invariant relations. However, this scheme is not applied to our 2D-to-3D regression of human pose as the following reasons. While original RN tries to capture the holistic relations that does not depend on the position of the objects or order of pairs, the groups on human body represent different parts where order of pairs matters. For instance, if the 2D positions of the left arm and the right arm are switched, the 3D pose should also be changed accordingly. However, the relational features generated will be the same for both cases if the order of pair is not considered. For these reasons, we did not use weight sharing for relational models. The 3D HPE algorithm proposed in this paper is formulated as

$$S_{3D}(S_{2D}) = f\left(\frac{1}{n_p} \sum_{(i,j)} g_{i,j}(G_i, G_j)\right), \quad (2)$$

where n_p is the number of pairs, S_{3D} , S_{2D} represents 3D and 2D shape of human body joints respectively, and G_i corresponds to the subsets of 2D input joints belonging to group i . We divide the input 2D joints to non-overlapping five groups as illustrated in Fig. 1(a). Total 16 joints are given as an input to the proposed network. Each joint group contains 3 or 4 joints, which we designed so that each group has a small range of variations. Each group represents a different part of a human body in this configuration. In other words, the groups contain joints from left/right arms, left/right legs, or the rest (a head and a torso). Thus, the relational network captures how different body parts are related with each other. All pairs of (i, j) such that $i < j$ are fed to the network and generates features of the same dimension. The mean of the relational features is passed to the next network module that is denoted as $f(\cdot)$ in Eq. 2. We empirically found that using the mean of the relational features instead of the sum stabilizes training.

We used ResNet structures proposed in [9] for neural network modules that are used for relation extraction and 3D pose estimation. The structure of a single module is illustrated in Fig. 1(b). A fully connected layer is firstly applied to increase the input dimension to that of a feature vector. Then, a residual network consisting of two sets of batch normalization [10], dropout [5], a ReLU activation function, and a fully connected layer is applied. The overall structure of the proposed network for 3D HPE is illustrated in Fig. 1(c).

It can be advantageous if we are able to capture the relations of pairs of individual joints. However, in this case, there are total $\frac{n_{2D}(n_{2D}-1)}{2}$ pairs which makes the network quite large. Instead, we designed a hierarchical relational network in which relations between two joints in a group are extracted within the group. The feature of each group G_k is generated as

$$G_k = \frac{1}{n_{pk}} \sum_{(i,j)} g_{i,j}^k(P_i, P_j), \quad (3)$$

where n_{pk} is the number of pairs in group k , and P_i, P_j correspond to 2D joints that belong to group k . The generated features are used as an input to the next relational network which is formulated as Eq. 2. Empirically, we observe that the hierarchical representation does not outperform a single level relational network, but the structure is advantageous if the relational dropout is applied as described in Section 3.2.

3.2 Relational Dropout

In this section, we propose a regularization method, which we call ‘*relational dropout*’, that can be applied to relational networks. Similar to dropout [6], we randomly drop the relational feature vectors that contain information on a certain group. In this paper, we restrict the number of dropping element to be at most 1. Thus, when the number of groups is n_G , among the $\frac{n_G(n_G-1)}{2}$ pairs, $n_G - 1$ relational feature vectors are dropped and replaced with zero vectors when relational dropout is applied. After the mean of the feature vectors are calculated, it is divided by the portion of non-dropping vectors to maintain the scale of the feature vector as in the general dropout method. Concretely, when group k is selected to be dropped, the formulation becomes

$$S_{3D}(S_{2D}|\text{drop} = k) = f\left(\frac{1}{n_p - n_G + 1} \sum_{(i \neq k, j \neq k)} g_{i,j}(G_i, G_j)\right). \quad (4)$$

Dropping features of a certain group simulates the case that the 2D points belonging to the dropping group are missing. Hence, the network learns to estimate the 3D pose not only when all the 2D joints are visible but also when some of them are invisible. The relational dropout is applied with the probability of p_{drop} during the training. Since at most one group is dropped, the combinational variability of missing joints is limited. To alleviate the problem, we applied the proposed relational dropout to hierarchical relational networks. In this case, we are able to simulate the case when a certain joint in a group is missing and to simulate various combinations of missing joints. At test time, we simply apply relational dropout to the groups that contain missing points.

3.3 Implementation Details

For the networks used in the experiments, the pose estimator $f(\cdot)$ in the relational networks has fully connected layers of 2,048 dimensions with a dropout probability of 0.5. For the modules $g_{i,j}(\cdot)$ that generates relational feature vector of the pairs G_i and G_j , 1,024 dimensional fully connected layers with a dropout probability of 0.25 are used. Lastly, for the hierarchical relational networks, the modules that generate relations from the pairs of 2D joints consist of 256 dimensional fully connected layers with a dropout probability of 0.1. When the relational dropout is applied during the training, p_{drop} is set to 0.2 for the case that one of the groups of joints is dropped, and it is set to 0.1 when the relational dropout is applied to the hierarchical relational units to drop a single joint.

We used stacked hourglass network [2] to infer 2D joint positions from training and test images. We finetuned the network pre-trained on MPII human pose dataset [4] using the frames of Human3.6M dataset. Mean subtraction is the only pre-processing applied to both 2D and 3D joint positions.

The proposed network is trained using ADAM optimizer [14] with a starting learning rate of 0.001. The batch size is set to 128, and the learning rate is halved for every 20,000 iterations. The network is trained for 100,000 iterations.

As a final note, we found that finetuning the trained model to each sequence of Human 3.6M dataset improves the estimation performance. During the finetuning, batch normalization statistics are fixed and the dropout probability is set to 0.5 in all modules.

Method	Direct	Discuss	Eat	Greet	Phone	Photo	Pose	Purchase
Pavlakos <i>et al.</i> [26]	67.4	71.9	66.7	69.1	72.0	77.0	65.0	68.3
Tekin <i>et al.</i> [52]	54.2	61.4	60.2	61.2	79.4	78.3	63.1	81.6
Zhou <i>et al.</i> [40]	54.8	60.7	58.2	71.4	62.0	65.5	53.8	55.6
Martinez <i>et al.</i> [19]	51.8	56.2	58.1	59.0	69.5	78.4	55.2	58.1
Fang <i>et al.</i> [8]	50.1	54.3	57.0	57.1	66.6	73.3	53.4	55.7
Cha <i>et al.</i> [9]	48.4	52.9	55.2	53.8	62.8	73.3	52.3	52.2
Yang <i>et al.</i> [44]	51.5	58.9	50.4	57.0	62.1	65.4	49.8	52.7
FC baseline	50.5	54.5	52.4	56.7	62.2	74.0	55.2	52.0
RN-hier	49.9	53.9	52.8	56.6	60.8	76.1	54.3	51.3
RN	49.7	54.0	52.0	56.4	60.9	74.1	53.4	51.1
RN-FT	49.4	54.3	51.6	55.0	61.0	73.3	53.7	50.0

Method	Sit	SitDown	Smoke	Wait	WalkD	Walk	WalkT	Avg
Pavlakos <i>et al.</i> [26]	83.7	96.5	71.7	65.8	74.9	59.1	63.2	71.9
Tekin <i>et al.</i> [52]	70.1	107.3	69.3	70.3	74.3	51.8	63.2	69.7
Zhou <i>et al.</i> [40]	75.2	111.6	64.2	66.1	51.4	63.2	55.3	64.9
Martinez <i>et al.</i> [19]	74.0	94.6	62.3	59.1	65.1	49.5	52.4	62.9
Fang <i>et al.</i> [8]	72.8	88.6	60.3	57.7	62.7	47.5	50.6	60.4
Cha <i>et al.</i> [9]	71.0	89.9	58.2	53.6	61.0	43.2	50.0	58.8
Yang <i>et al.</i> [44]	69.2	85.2	57.4	58.4	60.1	43.6	47.7	58.6
FC baseline	70.0	90.8	58.7	56.8	60.4	46.3	52.2	59.7
RN-hier	68.5	90.9	58.5	56.4	59.3	45.5	50.0	59.2
RN	69.3	90.4	58.1	56.4	59.5	45.6	50.6	59.0
RN-FT	68.5	88.7	58.6	56.8	57.8	46.2	48.6	58.6

Table 1: MPJPE (in mm) on Human 3.6M dataset under Protocol 1.

4 Experimental Results

We used Human 3.6M dataset [13] to validate the proposed algorithm. The dataset is the largest dataset for 3D HPE, and it consists of 15 action sequences which were performed by 7 different persons. Following the previous works, we used 5 subjects (S1, S5, S6, S7, S8) for training and 2 subjects (S9, S11) for testing. Mean per-joint position error (MPJPE) is used as an evaluation metric. We reported MPJPE for two types of alignments: aligning the root joints of the estimated pose and the ground truth pose denoted as *Protocol 1*, and aligning via Procrustes analysis including scaling, rotation, and translation denoted as *Protocol 2*. The proposed method is compared to the recently proposed methods that estimates 3D pose from a single image [9, 8, 19, 21, 26, 52, 56, 40].

To compare the performance of the proposed algorithm to the network that does not use relational networks, we designed a baseline network containing only fully connected layers. The baseline network consists of two consecutive ResBlocks of 2,048 dimensions. Dropout with probability of 0.5 is applied.

The MPJPE of various algorithms using Protocol 1 is provided in Table 1. It can be seen that the baseline network already outperforms most of the existing methods, which validates the superiority of the proposed residual modules. The relational networks are trained without applying relational dropouts. The proposed relational network (*RN*) gains 0.7 mm improvements over the baseline on average, and it is further improved when the network is finetuned on each sequence (*RN-FT*), which achieves state-of-the-art performance. Therefore, it is verified that capturing relations between different groups of joints improves the pose estimation performance despite its simpler structure and training procedures than the compared methods. Hierarchical relational networks (*RN-hier*) does not outperform *RN* al-

Method	Direct	Discuss	Eat	Greet	Phone	Photo	Pose	Purchase
Moreno-Noguer [24]	66.1	61.7	84.5	73.7	65.2	67.2	60.9	67.3
Martinez et al. [14]	39.5	43.2	46.4	47.0	51.0	56.0	41.4	40.6
Fang et al. [8]	38.2	41.7	43.7	44.9	48.5	55.3	40.2	38.2
Cha et al. [9]	39.6	41.7	45.2	45.0	46.3	55.8	39.1	38.9
Yang et al. [36]	26.9	30.9	36.3	39.9	43.9	47.4	28.8	29.4
FC baseline	43.3	45.7	44.2	48.0	51.0	56.8	44.3	41.1
RN-hier	42.5	44.9	44.2	47.4	49.1	57.4	43.9	40.5
RN	42.4	45.2	44.2	47.5	49.5	56.4	43.0	40.5
RN-FT	38.3	42.5	41.5	43.3	47.5	53.0	39.3	37.1

Method	Sit	SitDown	Smoke	Wait	WalkD	Walk	WalkT	Avg
Moreno-Noguer [24]	103.5	74.6	92.6	69.6	71.5	78.0	73.2	74.0
Martinez et al. [14]	56.5	69.4	49.2	45.0	49.5	38.0	43.1	47.7
Fang et al. [8]	54.5	64.4	47.2	44.3	47.3	36.7	41.7	45.7
Cha et al. [9]	55.0	67.2	45.9	42.0	47.0	33.1	40.5	45.7
Yang et al. [36]	36.9	58.4	41.5	30.5	29.5	42.5	32.2	37.7
FC baseline	57.0	68.8	49.2	45.3	50.5	38.2	45.0	48.9
RN-hier	56.7	68.5	48.5	44.7	49.4	37.0	43.1	48.1
RN	56.8	68.4	48.4	44.7	49.8	37.6	44.1	48.2
RN-FT	54.1	64.3	46.0	42.0	44.8	34.7	38.7	45.0

Table 2: MPJPE on Human 3.6M dataset under Protocol 2.

Method	Protocol 1				Protocol 2			
	None	Rand 2	L Arm	R Leg	None	Rand 2	L Arm	R Leg
Moreno-Noguer [24]	-	-	-	-	74.0	106.8	109.4	100.2
FC baseline	59.7	256.1	213.9	222.7	48.9	192.3	153.8	155.7
FC-drop	68.6	241.6	98.1	90.6	52.3	159.7	82.0	70.2
RN	59.0	540.2	314.1	332.8	48.2	280.7	225.8	214.1
RN-drop	59.3	218.7	73.8	70.6	45.5	145.3	62.7	55.0
RN-hier-drop	59.7	65.9	74.5	70.4	45.6	51.4	63.0	55.2

Table 3: MPJPE on Human 3.6M dataset with various types of missing joints.

though it has bigger number of parameters than *RN*. We conjecture the reason to be that it is hard to capture the useful relations in a small number of joints which leads to output poorer features than the ones using the raw 2D positions.

The MPJPE using the alignment Protocol 2 is provided in Table 2. When shape aligning via Procrustes analysis is applied, our method RN-FT showed superior performance to the existing methods except [36].

Next, we discuss the effectiveness of the relational dropout for the case of missing joints. MPJPE for all sequences with various types of missing joints are measured and provided in Table 3. We simulated 3 types of missing joints following [24], which are 2 random joints (Rand 2), left arm (L Arm), and right leg (R Leg). We consider 3 missing joints for the latter 2 cases including shoulder or hip joints. Note that [24] used different training schemes for experiments on missing joints where six subjects were used for training. For the baseline method that can be applied to the fully connected network, we assign zero to the value of input 2D joints with the probability of 0.1, which is denoted as *FC-drop*. It imposes robustness to the missing joints compared to the *FC baseline* in which random drop is not applied. When relational dropout is applied to the relational network (*RN-drop*), the model outperforms *FC-drop* in all cases. The model successfully estimates 3D pose when one of the groups in the relational network is missing. Therefore, it shows smaller MPJPE when the left arm or the right leg is not visible. However, when two joints belonging to different

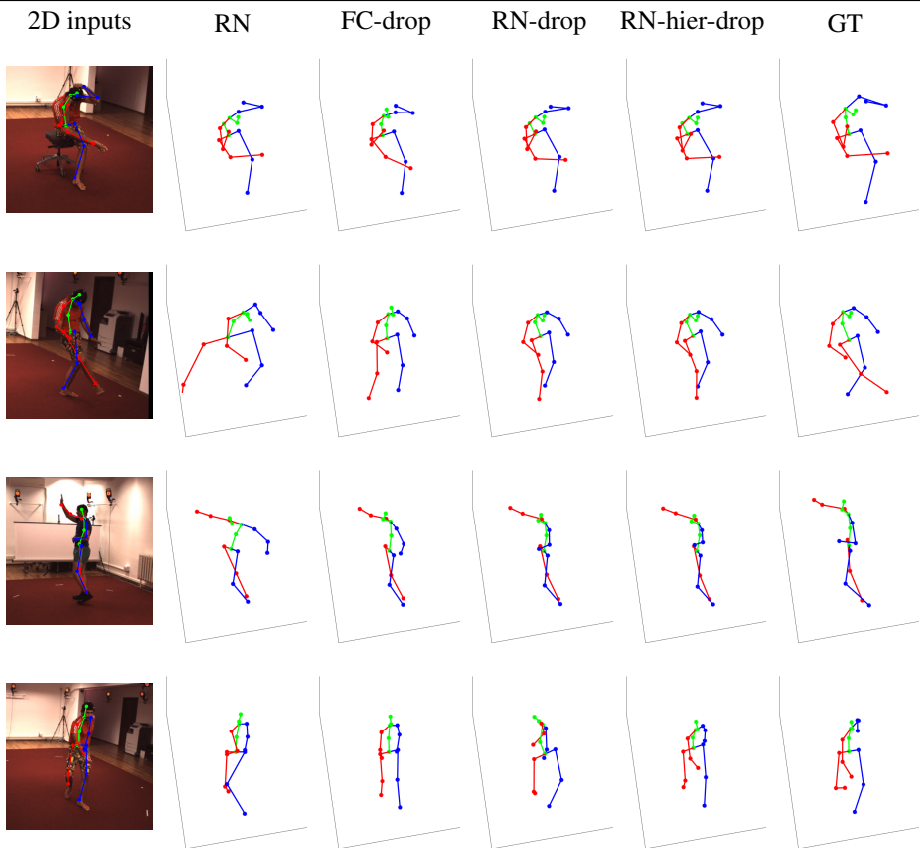


Figure 2: Qualitative results on Human 3.6M dataset in various cases of missing joints. For the 2D pose detection results, visible joints are marked as \bullet , and missing joints are marked as \times . Five groups are denoted as green (torso), red (right arm/leg) and blue (left arm/leg).

groups are missing, the two groups are dropped at the same time, which is not simulated during the training. Thus, *RN-drop* shows poor performance for the case that random two joints are missing. This problem can be handled when relational dropout is applied to the hierarchical relational network. When one joint is missing in a group, relational dropout is applied to hierarchical relational unit within the group. In the case that two or more joints are missing in a group, relational dropout is applied to the group. This model (*RN-hier-drop*) showed impressive performance in all types of missing joints. Another advantage of the relational dropout is that it does not degrade the performance of the case of all-visible joints. It can be inferred that the robustness on missing joints increases as various combinations of missing joints are simulated during the training.

Qualitative results on Human 3.6M dataset are provided in Figure 2. Each row simulates different cases of missing joints, none, right leg, left arm, and random 2 joints. The results of *RN*, *FC-drop*, *RN-drop*, *RN-hier-drop* is displayed with ground truth poses. When all joints are visible, all models generate similar poses that are close to the ground truth. On the other hand, *RN* generates inaccurate poses when 2D inputs contain missing points. *RN-drop* provides more accurate results than *FC-drop*, but the model fails when joints of two different

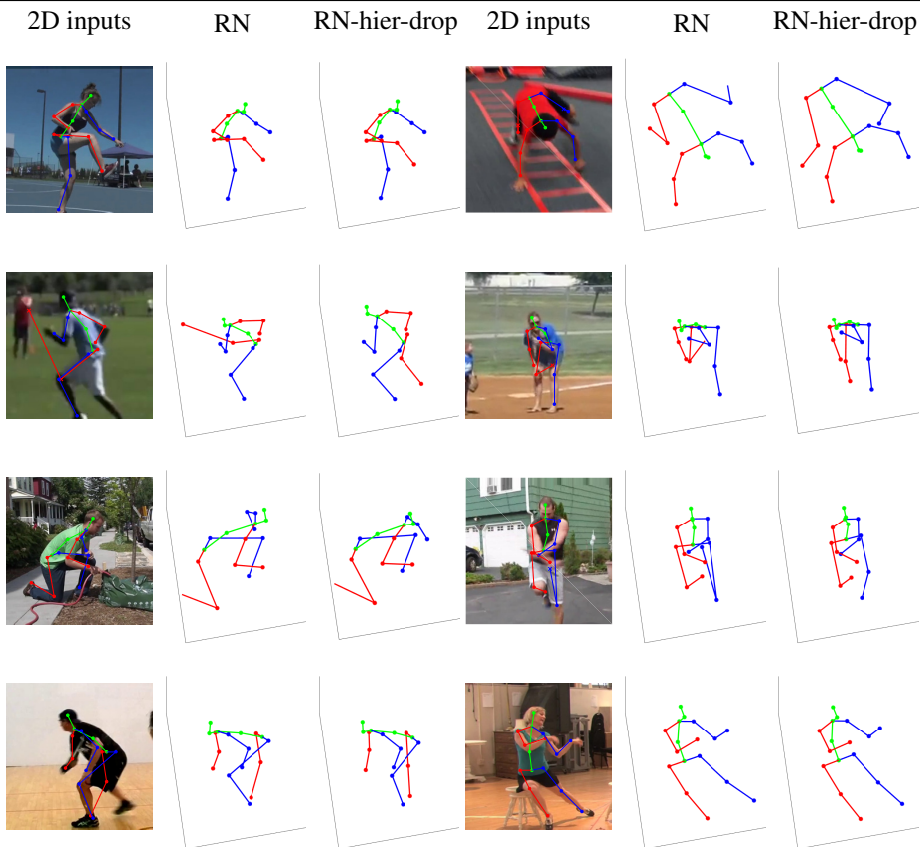


Figure 3: Qualitative results on MPII pose dataset.

groups are missing. It can be seen that *RN-hier-drop* outputs 3D poses that are similar to the ground truth poses in all cases. More results can be found on the supplementary materials.

Lastly, we displayed qualitative results on real world images. We used MPII human pose dataset [2] which is designed for 2D human pose estimation. 3D pose estimation results for the relational network (*RN*) and the hierarchical relational network with relational dropouts (*RN-hier-drop*) are provided in Figure 3. We first generate 2D pose results for the images and the joints whose maximum heatmap value is less than 0.4 are treated as missing joints for *RN-hier-drop*. As it can be seen in the second and third rows of Figure 2, *RN-hier-drop* generates more plausible poses than *RN* when some 2D joints are wrongly detected. The last row shows failure cases which contain noisy 2D inputs or an unfamiliar 3D pose that is not provided during the training.

5 Conclusion

In this paper, we propose a novel method for 3D human pose estimation. The relational network designed for 3D pose estimation showed state-of-the-art performance despite its simple structure. We also proposed the relational dropout which is fitted for the relational network. The relational dropout successfully impose the robustness to the missing points

while maintaining the performance of the original network. The proposed network is flexible in that it allows lots of variations in terms of its structure, group organization, and the policy of the relational dropout. The relational dropout can also be applied to other tasks that use relational networks.

Acknowledgments

This work was supported by Next-Generation Information Computing Development Program through the National Research Foundation of Korea (2017M3C4A7077582).

References

- [1] Ankur Agarwal and Bill Triggs. Recovering 3d human pose from monocular images. *IEEE transactions on pattern analysis and machine intelligence*, 28(1):44–58, 2006.
- [2] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2014.
- [3] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J Black. Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In *European Conference on Computer Vision*, pages 561–578. Springer, 2016.
- [4] Geonho Cha, Minsik Lee, Jungchan Cho, and Songhwai Oh. Deep pose consensus networks. *arXiv preprint arXiv:1803.08190*, 2018.
- [5] Ching-Hang Chen and Deva Ramanan. 3d human pose estimation = 2d pose estimation + matching. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [6] Haoshu Fang, Yuanlu Xu, Wenguan Wang, Xiaobai Liu, and Song-Chun Zhu. Learning knowledge-guided pose grammar machine for 3d human pose estimation. *arXiv preprint arXiv:1710.06513*, 2017.
- [7] Agne Grincunaite, Amogh Gudi, Emrah Tasli, and Marten den Uyl. Human pose estimation in space and time using 3d cnn. In *European Conference on Computer Vision*, pages 32–39. Springer, 2016.
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European Conference on Computer Vision*, pages 630–645. Springer, 2016.
- [10] Mir Rayat Imtiaz Hossain and James J Little. Exploiting temporal information for 3d pose estimation. *arXiv preprint arXiv:1711.08585*, 2017.

- [11] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 448–456, Lille, France, 07–09 Jul 2015. PMLR.
- [12] Catalin Ionescu, Fuxin Li, and Cristian Sminchisescu. Latent structured models for human pose estimation. In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pages 2220–2227. IEEE, 2011.
- [13] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. 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, 2014.
- [14] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [15] Christoph Lassner, Javier Romero, Martin Kiefel, Federica Bogo, Michael J Black, and Peter V Gehler. Unite the people: Closing the loop between 3d and 2d human representations. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [16] Minsik Lee, Jungchan Cho, and Songhwai Oh. Consensus of non-rigid reconstructions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4670–4678, 2016.
- [17] Sijin Li and Antoni B Chan. 3d human pose estimation from monocular images with deep convolutional neural network. In *Asian Conference on Computer Vision*, pages 332–347. Springer, 2014.
- [18] Sijin Li, Weichen Zhang, and Antoni B Chan. Maximum-margin structured learning with deep networks for 3d human pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2848–2856, 2015.
- [19] Julieta Martinez, Rayat Hossain, Javier Romero, and James J. Little. A simple yet effective baseline for 3d human pose estimation. In *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
- [20] Dushyant Mehta, Srinath Sridhar, Oleksandr Sotnychenko, Helge Rhodin, Mohammad Shafiei, Hans-Peter Seidel, Weipeng Xu, Dan Casas, and Christian Theobalt. Vnect: Real-time 3d human pose estimation with a single rgb camera. *ACM Transactions on Graphics (TOG)*, 36(4):44, 2017.
- [21] Francesc Moreno-Noguer. 3d human pose estimation from a single image via distance matrix regression. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1561–1570. IEEE, 2017.
- [22] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *European Conference on Computer Vision*, pages 483–499. Springer, 2016.

- [23] Katsunori Onishi, Tetsuya Takiguchi, and Yasuo Ariki. 3d human posture estimation using the hog features from monocular image. In *Pattern Recognition, 2008. ICPR 2008. 19th International Conference on*, pages 1–4. IEEE, 2008.
- [24] Sunghoon Park, Jihye Hwang, and Nojun Kwak. 3d human pose estimation using convolutional neural networks with 2d pose information. In Gang Hua and Hervé Jégou, editors, *Computer Vision – ECCV 2016 Workshops*, pages 156–169, Cham, 2016. Springer International Publishing.
- [25] Georgios Pavlakos, Xiaowei Zhou, Konstantinos G. Derpanis, and Kostas Daniilidis. Harvesting multiple views for marker-less 3d human pose annotations. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [26] Georgios Pavlakos, Xiaowei Zhou, Konstantinos G Derpanis, and Kostas Daniilidis. Coarse-to-fine volumetric prediction for single-image 3d human pose. In *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*, pages 1263–1272. IEEE, 2017.
- [27] Grégory Rogez, Jonathan Rihan, Srikumar Ramalingam, Carlos Orrite, and Philip HS Torr. Randomized trees for human pose detection. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8. IEEE, 2008.
- [28] Gregory Rogez, Philippe Weinzaepfel, and Cordelia Schmid. Lcr-net: Localization-classification-regression for human pose. In *CVPR 2017-IEEE Conference on Computer Vision & Pattern Recognition*, 2017.
- [29] Adam Santoro, David Raposo, David G. T. Barrett, Mateusz Malinowski, Razvan Pascanu, Peter Battaglia, and Timothy P. Lillicrap. A simple neural network module for relational reasoning. *CoRR*, abs/1706.01427, 2017.
- [30] Marta Sanzari, Valsamis Ntouskos, and Fiora Pirri. Bayesian image based 3d pose estimation. In *European Conference on Computer Vision*, pages 566–582. Springer, 2016.
- [31] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- [32] Xiao Sun, Jiayang Shang, Shuang Liang, and Yichen Wei. Compositional human pose regression. In *The IEEE International Conference on Computer Vision (ICCV)*, volume 2, 2017.
- [33] Bugra Tekin, Isinsu Katircioglu, Mathieu Salzmann, Vincent Lepetit, and Pascal Fua. Structured prediction of 3d human pose with deep neural networks. In *Proceedings of the British Machine Vision Conference (BMVC)*, pages 130.1–130.11. BMVA Press, September 2016.
- [34] Bugra Tekin, Pablo Marquez-Neila, Mathieu Salzmann, and Pascal Fua. Learning to fuse 2d and 3d image cues for monocular body pose estimation. In *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.

- [35] Denis Tome, Christopher Russell, and Lourdes Agapito. Lifting from the deep: Convolutional 3d pose estimation from a single image. *CVPR 2017 Proceedings*, pages 2500–2509, 2017.
- [36] Wei Yang, Wanli Ouyang, Xiaolong Wang, Jimmy Ren, Hongsheng Li, and Xiaogang Wang. 3d human pose estimation in the wild by adversarial learning. *arXiv preprint arXiv:1803.09722*, 2018.
- [37] Hashim Yasin, Umar Iqbal, Bjorn Kruger, Andreas Weber, and Juergen Gall. A dual-source approach for 3d pose estimation from a single image. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [38] Xiaowei Zhou, Menglong Zhu, Spyridon Leonardos, Konstantinos G Derpanis, and Kostas Daniilidis. Sparseness meets deepness: 3d human pose estimation from monocular video. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4966–4975, 2016.
- [39] Xiaowei Zhou, Menglong Zhu, Georgios Pavlakos, Spyridon Leonardos, Konstantinos G Derpanis, and Kostas Daniilidis. Monocap: Monocular human motion capture using a cnn coupled with a geometric prior. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.
- [40] Xingyi Zhou, Qixing Huang, Xiao Sun, Xiangyang Xue, and Yichen Wei. Towards 3d human pose estimation in the wild: a weakly-supervised approach. In *IEEE International Conference on Computer Vision*, 2017.