

# 3D HUMAN POSE REGRESSION USING GRAPH CONVOLUTIONAL NETWORK

Soubarna Banik\*    Alejandro Mendoza García†    Alois Knoll\*

\* Technical University of Munich, Department of Informatics, Munich, Germany

† reFit Systems, Munich, Germany

## ABSTRACT

3D human pose estimation is a difficult task, due to challenges such as occluded body parts and ambiguous poses. Graph convolutional networks encode the structural information of the human skeleton in the form of an adjacency matrix, which is beneficial for better pose prediction. We propose one such graph convolutional network named PoseGraphNet for 3D human pose regression from 2D poses. Our network uses an adaptive adjacency matrix and kernels specific to neighbor groups. We evaluate our model on the Human3.6M dataset which is a standard dataset for 3D pose estimation. Our model’s performance is close to the state-of-the-art, but with much fewer parameters. The model learns interesting adjacency relations between joints that have no physical connections, but are behaviorally similar.

**Index Terms**— 3D Human pose regression, Graph Convolutional Networks, CNN

## 1. INTRODUCTION

3D human pose estimation (HPE) directly from RGB images is a challenging task. Compared to RGB images 3D point cloud data has inherent depth information, which is useful for 3D pose prediction but is computationally expensive to work with [1, 2]. Recent approaches [3, 4, 5] have split the task into two parts: predicting 2d pose from an image, and then lifting them to 3D poses. Restricting the 3D pose to the camera space allowed the 2D to 3D pose task to be learnable using only supervised learning techniques [6].

It has been shown that additional information such as structural constraints of the skeleton improves the performance further [7, 5]. In some earlier work, the structural information is provided to the system by a separate statistical model [3] or by a dictionary of basic poses [7]. These types of structural models need to be learned separately. As the learned structural model is dependent on the training dataset, it also restricts the number of possible 3D poses.

Recent advances in graph convolutional networks (GCN) and their use in pose estimation look promising [5, 8, 9]. GCNs provide a lightweight solution for representing the structure of a human body without any additional overhead.

Zhao et al. [8] propose a semantic GCN for encoding both local and global relationships among different body joints for 2D to 3D pose regression. Cai et al. [5] use a spatio-temporal GCN to capture the temporal constraints in addition to the structural constraints. They also capture multi-scale features by using a U-Net architecture with residual connections. Hope-Net [9] uses an adaptive adjacency matrix and a graph U-Net architecture for hand-object pose estimation.

In this paper, we propose a graph convolutional network named PoseGraphNet for 2D to 3D pose regression. Our model is similar to [6], but uses GCN instead of a fully connected network. It uses adaptive adjacency matrices and kernels specific to neighbor groups. Our method outperforms the base model [6] on Human3.6m dataset [10] and is close to the performance of the state-of-the-art graph-based approaches [5] despite having much fewer parameters.

## 2. POSEGRAPHNET

We propose a graph convolution-based network, named PoseGraphNet for 3D pose regression. Figure 1 describes the overall pipeline and the architecture of our proposed PoseGraphNet model. We use a baseline 2D pose detection model named Cascaded Pyramid Network (CPN) [11] which takes an RGB image as input and generates a heatmap for every joint, indicating the location of the joint in the image. We extract the 2D joint coordinates from the heatmaps and pass them to our PoseGraphNet model. The output of our model is the 3D position of the joints relative to the root joint. The input joint coordinates are in image space and the outputs are in camera coordinate space.

**Graph Convolutional Networks:** GCN is a neural network that works directly on graphs and can learn features of graph-structured data. In graph-based pose estimation, the human skeleton is represented as a graph,  $G = (V, E)$ , where  $V$  denotes the set of  $N$  body joints and  $E$  denotes the set of edges signifying the connections between the body joints. The edges are represented by means of an adjacency matrix,  $A \in \mathbb{R}^{N \times N}$ . The graph consists of node-level features and  $A$  indicates, how the features of a node’s neighbors and itself will be aggregated to form the output of the layer.

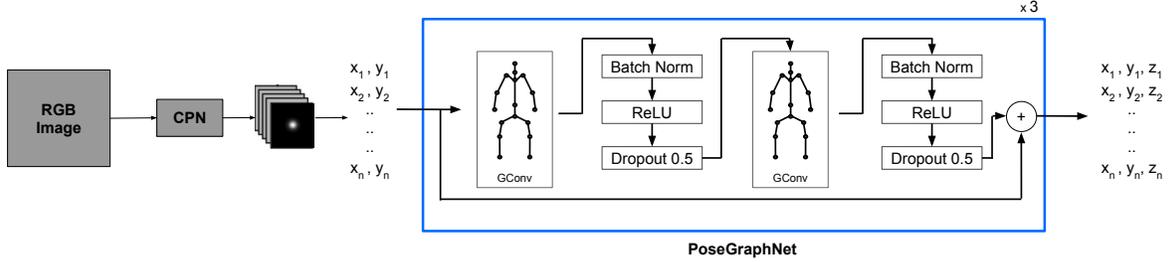


Fig. 1: Overview of the 3D pose regression algorithm. The proposed PoseGraphNet is highlighted by the blue rectangle.

## 2.1. Architecture

Our proposed PoseGraphNet is based on [6] by Martinez et al. We choose this model in comparison to the more recent, U-shaped graph-based models [5, 9], as the feed-forward network of [6] has less computational complexity and performs well without utilizing any structural information of the skeleton. We also want to find out how much the structural information and graph convolution alone help in improving the performance compared to that of [6]. PoseGraphNet consists of 3 residual blocks, where each block is made of two units of a graph convolution layer with 768 feature channels, batch normalization, a ReLU activation, and dropout. The block input is added to the block output with a residual connection. An initial unit of graph convolution of 768 channels and a final graph convolution layer of 3 output channels are not depicted in Figure 1 due to space constraints.

**Adaptive adjacency matrix:** In addition to learning the trainable weights of the graph convolution layers, we also learn the adjacency matrix as in [9]. A pre-defined adjacency matrix  $A$  representing the default physical structure of the human body allows only equal contribution by every neighbor. We observe, for most actions some body joints such as the trunk are relatively stable and do not move rapidly compared to the joints at the extremities. Therefore, a neighbor joint situated at a body extremity should have a different contribution compared to a neighbor joint that is closer to the stable body trunk. An adaptive adjacency matrix allows learning different contribution rates for different types of neighbors in addition to learning physically missing but behaviorally similar connections.

**Kernels for different neighbor groups:** In addition to the adaptive adjacency matrix, we learn different weights for different types of neighbors following [5], resulting in the extraction of distinct features for different neighbor groups. We categorize the neighboring nodes into three groups - a) self, b) parent (closer to the pelvis joint), and c) child (away from the pelvis joint). Output features of hidden layer  $i$  is computed as,

$$H_i = \sigma(\tilde{A}^{self} H_{i-1} W_i^{self} + \tilde{A}^{parent} H_{i-1} W_i^{parent} + \tilde{A}^{child} H_{i-1} W_i^{child}), \quad (1)$$

where  $\sigma$  is the activation function,  $H_{i-1} \in \mathbb{R}^{N \times F_{i-1}}$  is the input feature matrix,  $F_i$  is the number of features in layer  $i$  and  $W_i^{self}, W_i^{parent}, W_i^{child} \in \mathbb{R}^{F_{i-1} \times F_i}$  are the trainable weight matrices of layer  $i$ , for respective neighbor groups.  $\tilde{A}^{self}, \tilde{A}^{parent}$  and  $\tilde{A}^{child}$  are the adjacency matrices of the respective neighbor groups, normalized symmetrically with their corresponding degree matrices. Note, all three adjacency matrices are adaptive and learned together with the weight matrices.

## 3. RESULTS

### 3.1. Implementation Details

In our experiment, the 2D inputs are in the image coordinate system and are normalized so that image width is in the range  $[-1, +1]$ , maintaining the image aspect ratio. For the 3D outputs, we use the camera coordinate system, instead of an arbitrary world coordinate system. As pointed out in [6, 5, 8], it is unrealistic to train a model to infer 3D joint locations in an arbitrary coordinate space, as there can be multiple possible solutions for the same 2D input. The camera frame makes the 2D to 3D joint prediction task coherent across different camera views. Also, we are interested in the relative positions of the joints with respect to the root joint, which is the pelvis joint. Hence, we zero-center the output joints by subtracting the root joint.

### 3.2. Training Details

We train our model using the Adam optimizer [14] for 70 epochs, using a batch size of 256, a dropout rate of 0.5, and with a starting learning rate of 0.001. We use a learning rate scheduler and reduce the learning rate by a factor of 0.1 when the validation loss saturates. The network is initialized using Xavier initialization [15]. We do not use any data augmentation. We minimize the mean squared error (MSE) loss between the predicted joints and the ground truth joint,  $L(J^{pred}, J^{gt}) = \sum_{i=1}^N \|J_i^{pred} - J_i^{gt}\|^2$ , where  $J^{pred}$  and  $J^{gt}$  are the predicted and ground truth 3D joints respectively. The hip joint is excluded from evaluation as it remains zero throughout.

RHip	LHip	Spine	Thorax	Neck	LSho.	RSho.	LKnee	RKnee	Head	LElb.	RFoot	RElb.	LFoot	RWrist	LWrist	Avg.
22.3	23.5	35.6	43.9	53.6	55.0	55.2	56.7	57.3	60.3	79.5	79.6	83.0	91.6	106.9	108.1	59.5

**Table 1:** Mean Per Joint Position Error (MPJPE) (Protocol #1) in millimeter between predicted and ground-truth poses on Human 3. 6m Dataset at joint level, in increasing order. The last column reports the average error.

Protocol #1	Dir.	Disc.	Eat	Greet	Phone	Photo	Pose	Purch.	Sit	SitD.	Smoke	Wait	WalkD.	Walk	WalkT.	Avg.
Martinez et al. [6] ICCV'17	51.8	56.2	58.1	59.0	69.5	78.4	55.2	58.1	74.0	94.6	62.3	59.1	65.1	49.5	52.4	62.9
Cai et al. [5] ICCV'19	<b>46.5</b>	<b>48.8</b>	<b>47.6</b>	<b>50.9</b>	<b>52.9</b>	61.3	48.3	<b>45.8</b>	<b>59.2</b>	64.4	<b>51.2</b>	<b>48.4</b>	53.5	<b>39.2</b>	<b>41.2</b>	<b>50.6</b>
Zhao et al. [8] CVPR'19	47.3	60.7	51.4	60.5	61.1	<b>49.9</b>	<b>47.3</b>	68.1	86.2	<b>55.0</b>	67.8	61.0	<b>42.1</b>	60.6	45.3	57.6
Pavlo et al [12] CVPR'19	47.1	50.6	49.0	51.8	53.6	61.4	49.4	47.4	59.3	67.4	52.4	49.5	55.3	39.5	42.7	51.8
Pavlakos et al. [13] CVPR'18	48.5	54.4	54.4	52.0	59.4	65.3	49.9	52.9	65.8	71.1	56.6	52.9	60.9	44.7	47.8	56.2
Ours	51.0	55.3	54.0	54.6	62.4	76.0	51.6	52.7	79.3	87.1	58.4	56.0	61.8	48.1	44.1	59.5
Protocol #2	Dir.	Disc.	Eat	Greet	Phone	Photo	Pose	Purch.	Sit	SitD.	Smoke	Wait	WalkD.	Walk	WalkT.	Avg.
Martinez et al. [6] ICCV'17	39.5	43.2	46.4	47.0	51.0	56.0	41.4	40.6	56.5	69.4	49.2	45.0	49.5	38.0	43.1	47.7
Cai et al. [5] ICCV'19	36.8	<b>38.7</b>	38.2	41.7	40.7	46.8	37.9	35.6	47.6	<b>51.7</b>	<b>41.3</b>	<b>36.8</b>	<b>42.7</b>	31.0	<b>34.7</b>	40.2
Pavlo et al [12] CVPR'19	36.0	<b>38.7</b>	<b>38.0</b>	41.7	<b>40.1</b>	<b>45.9</b>	<b>37.1</b>	<b>35.4</b>	<b>46.8</b>	53.4	41.4	36.9	43.1	<b>30.3</b>	34.8	<b>40.0</b>
Pavlakos et al. [13] CVPR'18	<b>34.7</b>	39.8	41.8	<b>38.6</b>	42.5	47.5	38.0	36.6	50.7	56.8	42.6	39.6	43.9	32.1	36.5	41.8
Ours	38.4	43.1	42.9	44.0	47.8	56.0	39.3	39.8	61.8	67.1	46.1	43.4	48.4	40.7	35.1	46.4

**Table 2:** Mean Per Joint Position Error (MPJPE) in millimeter between predicted and ground-truth 3D poses on Human 3.6m Dataset under Protocol #1 and Protocol #2. The lower the error, the better. The best scores are highlighted in bold.

Model	Parameters	Throughput	MPJPE (mm)
Cai et al. [5]	3.77M	95.53	50.6
Martinez et al. [6]	4.29M	911.92	62.9
Ours	1.25M	207.29	59.5

**Table 3:** Number of parameters, Throughput (forward passes/second with batch size 1), and MPJPE (Protocol-1) of models. Higher throughput is better.

### 3.3. Dataset

We evaluate our proposed approach on the Human3.6M dataset [10], a standard dataset for 3D human pose estimation. The dataset contains 3.6 million RGB images of 7 subjects, from 4 different viewpoints performing 15 actions such as walking, eating, sitting, discussing, etc. 2D locations in image space and 3D ground truth positions in camera space of 17 body joints are available. As per the standard protocol, we use images of all 4 viewpoints of subjects 1, 5, 6, 7, and 8 for training, and subjects 9 and 11 for evaluation.

### 3.4. Evaluation Metrics

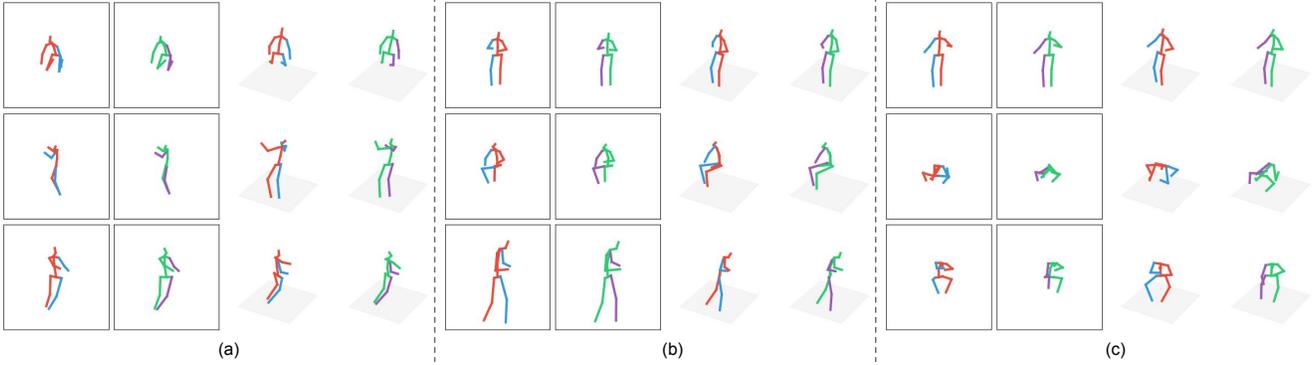
We evaluate the PoseGraphNet using mean per joint position error (MPJPE), which computes the mean Euclidean distance of the predicted 3D joints to the ground truth joints in millimeters. In previous work MPJPE is computed under two protocols [6, 5, 8]. As per *Protocol #1*, the error is computed after aligning the root joints. In *Protocol #2* the 3D pose is further aligned with the ground truth using a rigid transformation (Procrustes alignment), before computing the error.

### 3.5. Quantative results

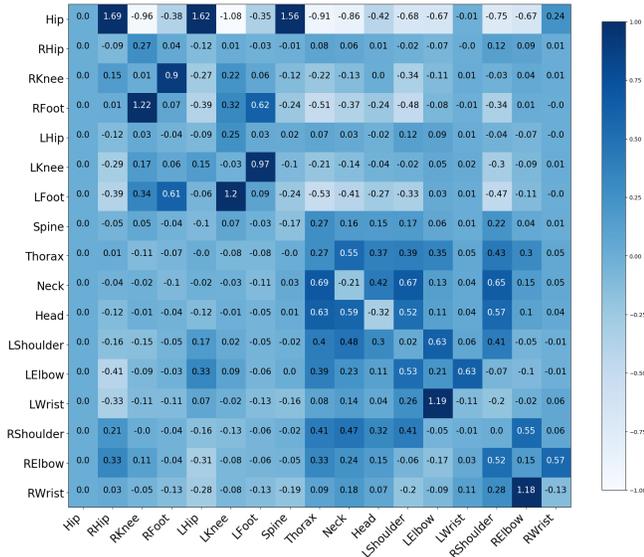
The average MPJPE score, as well as joint level scores for Protocol #1, are reported in Table 1. Most of the error occurs for joints at the body extremities such as head, elbow, foot, and wrist. Our model achieves 59.5 mm of mean error as per Protocol #1, particularly 20.7, 22.8, and 42.0 mm in x, y, and z axis respectively. As can be seen, the majority of the error is in the z or the depth axis.

Table 2 compares the performance of PoseGraphNet with the state-of-the-art 3D pose estimation methods. We only consider approaches that operate on a single input frame for this comparison. The mean MPJPE score for both protocols is reported in Table 2. The mean scores for individual actions are also reported. Though our method could not outperform Cai et al. [5] in Protocol #1 and Pavlo et al. [12] in Protocol #2, it achieved a better result than Martinez et al. [6]. Though our model is similar to [6], as we utilize the skeletal structure by using GCN, the performance improves.

In practice, the 2D to 3D pose regression works in succession to the 2D pose detection model. So the 2D to 3D model should be minimum in size and have high throughput. Our motivation for this research is to develop a model that is accurate, fast and small enough to be used in a real-world application such as reFit's Gamo [16]. We compare our method with the best performing state-of-the-art method [5] and also with Martinez et al.'s approach [6] in terms of the number of parameters and throughput in Table 3. Though our model's performance lags behind the state-of-the-art [5] by  $\approx 9mm$ , it has around  $1/3^{rd}$  the number of trainable parameters compared to the other two methods. We compute the throughput by the number of forward passes in one second with batch size 1. We compute this in evaluation mode in a machine with GeForce RTX 2060 GPU, AMD Ryzen 9 3900X 12-core pro-



**Fig. 2:** Qualitative results on the test set of Human3.6M dataset. For each column (a),(b), and (c), from left to right: 2D ground truth pose in image space, predicted 2D pose by CPN [11], 3D ground truth pose in camera space, predicted 3D pose. In ground truth 2D pose and 3D pose, left and right limbs are marked in blue and red respectively, and in the predicted pose, in purple and green respectively.



**Fig. 3:** Learned adjacency matrix  $\tilde{A}^{child}$ . Dark blue indicates strong relation.

processor, and 32GB RAM. As it is evident from Table 3, our model is accurate enough, even with relatively fewer parameters while maintaining a higher throughput than the state-of-the-art [5].

### 3.6. Qualitative Results

Figure 2 shows qualitative results on 9 randomly selected samples from the test set of Human3.6M dataset. For most of the cases, PoseGraphNet can detect the 3D pose correctly. It correctly detects the left and right joints, even in cases of the back view of the person (column (b), row 1 from top). Also,

error from the 2D pose detection stage propagates to the 2D to 3D pose part, as can be seen in few examples in Figure 2 (column (c), row 3).

Figure 3 shows the learned adjacency matrix  $\tilde{A}^{child}$  of PoseGraphNet. The adjacency values indicate interesting relationships between the joints. For example, the right and left feet are strongly related (0.62), though they are not physically connected. However, the right and left wrists have a weak relation (0.11). This also follows from intuition, because hands move more independently of each other compared to legs. It also learns a weak relation (0.21) between the right shoulder and the right hip, suggesting movement of the right shoulder will have a small effect on the right hip. As we did not enforce any restriction on the adjacency matrix,  $\tilde{A}^{child}$  contains connections irrespective of its original initialization of child connections.

## 4. CONCLUSION

We present a graph convolution network named PoseGraphNet for 3D human pose regression from 2D pose. Our model is an improved version of the fully connected model proposed in [6]. Introduction of the skeletal structural information into the model in the form of the adjacency matrix and graph convolution improves the prediction performance over [6]. Our model’s performance is close to that of the state-of-the-art [5], while having only  $1/3^{rd}$  the number of parameters and double the throughput. Additionally, employing an adaptive adjacency matrix allowed us to learn specific adjacency relations between joints, that cannot be handcrafted.

**Acknowledgements:** The authors acknowledge the financial support by the Federal Ministry of Education and Research of Germany for this project (Start MTI/16SV8115).

## 5. REFERENCES

- [1] Gyeongsik Moon, “V2V-PoseNet: Voxel-to-Voxel Prediction Network for Accurate 3D Hand and Human Pose Estimation from a Single Depth Map,” in *CVPR*, 2018.
- [2] Christian Zimmermann, Tim Welschehold, Christian Dornhege, Wolfram Burgard, and Thomas Brox, “3D Human Pose Estimation in RGBD Images for Robotic Task Learning,” in *ICRA*, 2018.
- [3] Denis Tome and Chris Russell, “Lifting from the Deep: Convolutional 3D Pose Estimation from a Single Image,” in *CVPR*, 2017.
- [4] Mir Rayat, Imtiaz Hossain, and James J Little, “Exploiting temporal information for 3D human pose estimation,” in *ECCV*, 2018.
- [5] Yujun Cai, Lihao Ge, Jun Liu, Jianfei Cai, Tat-Jen Cham, Junsong Yuan, and Nadia Magnenat Thalmann, “Exploiting Spatial-temporal Relationships for 3D Pose Estimation via Graph Convolutional Networks,” in *ICCV*, 2019.
- [6] Julieta Martinez, Rayat Hossain, Javier Romero, and James J Little, “A simple yet effective baseline for 3d human pose estimation,” in *ICCV*, 2017.
- [7] Manuel J Marin-Jimenez, Francisco J Romero-Ramirez, Rafael Munoz-Salinas, and Rafael Medina-Carnicer, “3d human pose estimation from depth maps using a deep combination of poses,” *Journal of Visual Communication and Image Representation*, 2018.
- [8] Long Zhao, Xi Peng, Yu Tian, Mubbasir Kapadia, and Dimitris N Metaxas, “Semantic Graph Convolutional Networks for 3D Human Pose Regression,” in *CVPR*, 2019.
- [9] Bardia Doosti, Shujon Naha, Majid Mirbagheri, and David J. Crandall, “Hope-Net: A graph-based model for hand-Object pose estimation,” *CVPR*, 2020.
- [10] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu, “Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments,” *PAMI*, 2013.
- [11] Yilun Chen, Zhicheng Wang, Yuxiang Peng, Zhiqiang Zhang, Gang Yu, and Jian Sun, “Cascaded Pyramid Network for Multi-person Pose Estimation,” *CVPR*, 2018.
- [12] Dario Pavllo, Eth Zürich, Christoph Feichtenhofer, David Grangier, Google Brain, and Michael Auli, “3D human pose estimation in video with temporal convolutions and semi-supervised training,” in *CVPR*, 2019.
- [13] Georgios Pavlakos, Xiaowei Zhou, and Kostas Daniilidis, “Ordinal Depth Supervision for 3D Human Pose Estimation,” *CVPR*, 2018.
- [14] Diederik P. Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” in *ICLR*, 2015.
- [15] Xavier Glorot and Yoshua Bengio, “Understanding the difficulty of training deep feedforward neural networks,” in *Proceedings of the 13th international conference on artificial intelligence and statistics. JMLR Workshop and Conference Proceedings*, 2010.
- [16] reFit Systems, “reFit Gamo - Digital therapy system for individual orthopedic rehabilitation,” 2021.