3D Human Pose Estimation Using Möbius Graph Convolutional Networks

Niloofar Azizi¹, Horst Possegger¹, Emanuele Rodolà², and Horst Bischof¹

¹ Graz University of Technology, Graz, Austria {azizi, possegger, bischof}@icg.tugraz.at
² Sapienza University of Rome, Rome, Italy rodola@di.uniroma1.it

Abstract. 3D human pose estimation is fundamental to understanding human behavior. Recently, promising results have been achieved by graph convolutional networks (GCNs), which achieve state-of-the-art performance and provide rather light-weight architectures. However, a major limitation of GCNs is their inability to encode all the transformations between joints explicitly. To address this issue, we propose a novel spectral GCN using the Möbius transformation (Möbius-GCN). In particular, this allows us to directly and explicitly encode the transformation between joints, resulting in a significantly more compact representation. Compared to even the lightest architectures so far, our novel approach requires 90–98% fewer parameters, *i.e.* our lightest MöbiusGCN uses only 0.042M trainable parameters. Besides the drastic parameter reduction, explicitly encoding the transformation of joints also enables us to achieve state-of-the-art results. We evaluate our approach on the two challenging pose estimation benchmarks, Human3.6M and MPI-INF-3DHP, demonstrating both state-of-the-art results and the generalization capabilities of MöbiusGCN.

1 Introduction

Estimating 3D human pose helps to analyze human motion and behavior, thus enabling high-level computer vision tasks such as action recognition [30], sports analysis [49, 64], augmented and virtual reality [15]. Although human pose estimation approaches already achieve impressive results in 2D, this is not sufficient for many analysis tasks, because several 3D poses can project to exactly the same 2D pose. Thus, knowledge of the third dimension can significantly improve the results on the high-level tasks.

Estimating 3D human joint positions, however, is challenging. On the one hand, there are only very few labeled datasets because 3D annotations are expensive. On the other hand, there are self-occlusions, complex joint inter-dependencies, small and barely visible joints, changes in appearance like clothing and lighting, and the many degrees of freedom of the human body.

To solve 3D human pose estimation, some methods utilize multi-views [50, 70], synthetic datasets [46], or motion [25, 55]. For improved generalization, however, we follow the most common line of work and estimate 3D poses given only the 2D estimate of a single RGB image as input, similar to [27, 33, 45]. First, we compute 2D pose joints



Fig. 1: Our MöbiusGCN accurately learns the transformation (particularly the rotation) between joints by leveraging the Möbius transformation given estimated 2D joint positions from a single RGB image. The spectral GCN puts a scalar-valued function on each node of the graph and compares the graph signal with the graph filter on the eigenspace of the graph Laplacian matrix and returns it to the spatial domain. We define the Möobius transformation as our scalar-valued function. Möbius transformation is mainly the composition of the rotation function and translation functions. Thus, it is capable of encoding the transformation between the human body joints.

given RGB images using an off-the-shelf architecture. Second, we approximate the 3D pose of the human body using the estimated 2D joints.

With the advent of deep learning methods, the accuracy of 3D human pose estimation has significantly improved, *e.g.* [33, 45]. Initially, these improvements were driven by CNNs (Convolutional Neural Networks). However, these assume that the input data is stationary, hierarchical, has a grid-like structure, and shares local features across the data domain. The convolution operator in the CNN assumes that the nodes have fixed neighbor positions and a fixed number of neighbor nodes. Therefore, CNNs are not applicable to graph-structured data. The input to 3D pose estimation from 2D joint positions, however, is graph-structured data. Thus, to handle this irregular nature of the data, GCNs (Graph Convolutional Networks) have been proposed [4].

GCNs are able to achieve state-of-the-art performance for 2D-to-3D human pose estimation with comparably few parameters, *e.g.* [71]. Nevertheless, to the best of our knowledge, none of the previous GCN approaches explicitly models the inter-segmental angles between joints. Learning the inter-segmental angle distribution explicitly along with the translation distribution, however, leads to encoding better feature representations. Thus, we present a novel spectral GCN architecture, *MöbiusGCN*, to accurately learn the transformation between joints and to predict 3D human poses given 2D joint positions from a single RGB image. To this end, we leverage the Möbius transformation on the eigenvalue matrix of the graph Laplacian. Previous GCNs applied for estimating the 3D pose of the human body are defined in the real domain, *e.g.* [71]. Our Möbius-GCN operates in the complex domain, which allows us to encode all the transformations (*i.e.* inter-segmental angles and translation) between nodes simultaneously (Figure 1).

An enriched feature representation achieved by encoding the transformation distribution between joints using a Möbius transformation provides us with a compact model. A light DNN architecture makes the network independent of expensive hardware setup, enabling the use of mobile phones and embedded devices at inference time. This can be achieved by our compact MöbiusGCN architecture.

Due to a large number of weights that need to be estimated, fully-supervised stateof-the-art approaches need an enormous amount of annotated data, where data annotation is both time-consuming and requires expensive setup. Our MöbiusGCN, on the contrary, requires only a tiny fraction of the model parameters, which allows us to achieve competitive results with significantly fewer annotated data.

We summarize our main contributions as follows:

- We introduce a novel spectral GCN architecture leveraging the Möbius transformation to explicitly encode the pose, in terms of inter-segmental angles and translations between joints.
- We achieve state-of-the-art 3D human pose estimation results, despite requiring only a fraction of the model parameters (*i.e.* 2–9% of even the currently lightest approaches).
- Our light-weight architecture and the explicit encoding of transformations lead to state-of-the-art performance compared to other semi-supervised methods, by training only on a reduced dataset given estimated 2D human joint positions.

2 Related Work

Human 3D Pose Estimation. The classical approaches addressing the 3D human pose estimation task are usually based on hand-engineered features and leverage prior assumptions, *e.g.* using motion models [57] or other common heuristics [18, 48]. Despite good results, their major downside is the lack of generality.

Current state-of-the-art approaches in computer vision, including 3D human pose estimation, are typically based on DNNs (Deep Neural Networks), *e.g.* [24, 29, 31, 52]. To use these architectures, it is assumed that the statistical properties of the input data have locality, stationarity, and multi-scalability [16], which reduces the number of parameters.

Although DNNs achieve state-of-the-art in spaces governed by Euclidean geometry, a lot of the real-world problems are of a non-Euclidean nature. For these problem classes, GCNs have been introduced. There are two types of GCNs: spectral GCN and spatial GCN. Spectral GCNs rely on the Graph Fourier Transform, which analyzes the graph signals in the vector space of the graph Laplacian matrix. The second category, spatial GCN, is based on feature transformations and neighborhood aggregation on the graph. Well-known spatial GCN approaches include Message Passing Neural Networks [12] and GraphSAGE [14].

For 3D human pose estimation, GCNs achieve competitive results with comparably few parameters. Pose estimation with GCNs has been addressed, *e.g.* in [27, 66, 71]. Xu and Takano [66] proposed Graph Stacked Hourglass Networks (GraphSH), in which graph-structured features are processed across different scales of human skeletal representations. Liu et al. [27] investigated different combinations of feature transformations and neighborhood aggregation of spatial features. They also showed the benefits of using separate weights to incorporate a node's self-information. Zhao et al. [71] proposed semantic GCN (SemGCN), which currently represents the lightest architecture

(0.43M). The key idea is to learn the adjacency matrix, which lets the architecture encode the graph's semantic relationships between nodes. In contrast to SemGCN, we can further reduce the number of parameters by an order of magnitude (0.042M) by explicitly encoding the transformation between joints. The key ingredient to this significant reduction is the Möbius transformation.

Möbius Transformation. The Möbius transformation has been used in neural networks as an activation function [32, 39], in hyperbolic neural networks [11], for data augmentation [73], and for knowledge graph embedding [37]. Our work is the first to introduce Möbius transformations for spectral graph convolutional networks. To utilize the Möbius transformation, we have to design our neural network in the complex domain. The use of complex numbers (analysis in polar coordinates) to harness phase information along with the signal amplitude is well established in signal processing [32]. By applying the Möbius transformation, we let the architecture encode the transformations (*i.e.* inter-segmental angle and translation) between joints explicitly, which leads to a very compact architecture.

Handling Rotations. Learning the rotation between joints in skeletons has been investigated previously for 3D human pose estimation [2, 40, 75]. Learning the rotation using Euler angles or quaternions, however, has obvious issues like discontinuities [53, 77]. Continuous functions are easier to learn in neural networks [77]. Zhou et al. [77] tackle the discontinuity by lifting the problem to 5 and 6 dimensions. Another direction of research focuses on designing DNNs for inverse kinematics with the restricted assumption of putting joint angles in a specific range to avoid discontinuities. However, learning the full range of rotations is necessary for many real-world problems [77]. Our MöbiusGCN is continuous by definition and thus, allows us to elegantly encode rotations.

Data Reduction. A major benefit of light architectures is that they require smaller datasets to train. Semi-supervised methods require a small subset of annotated data and a large set of unannotated data for training the network. These methods are actively investigated in different domains, considering the difficulty of providing annotated datasets. Several semi-supervised approaches for 3D human pose estimation benefit from applying more constraints over the possible space solutions by utilizing multiview approaches using RGB images from different cameras [35, 50, 51, 63, 69]. These methods need expensive laboratory setups to collect synchronized multi-view data.

Pavllo et al. [45] phrase the loss over the back-projected estimated 3D human pose to 2D human pose space conditioned on time. Tung et al. [61] use generative adversarial networks to reduce the required annotated data for training the architecture. Iqbal et al. [19] relax the constraints using weak supervision; they introduce an end-to-end architecture that estimates 2D pose and depth independently, and uses a consistency loss to estimate the pose in 3D.

Our compact MöbiusGCN achieves competitive state-of-the-art results with only scarce training samples. MöbiusGCN does not require any multi-view setup or temporal information. Further, it does not rely on large unlabeled datasets. It just requires a small annotated dataset to train. In contrast, the previous semi-supervised methods require complicated architectures and a considerable amount of unlabeled data during the training phase.

3 Spectral Graph Convolutional Network



Fig. 2: The complete pipeline of the proposed MöbiusGCN architecture; The output of the off-the-shelf stacked hourglass architecture [38], *i.e.* estimated 2D joints of the human body, is the input to the MöbiusGCN architecture. The MöbiusGCN architecture locally encodes the transformation between the joints of the human body. SVD is the singular value decomposition of the normalized Laplacian matrix. Function g_{θ} is the Möbius transformation applied on the eigenvalues of the eigenvalue matrix independently. x is the graph signal and ω are the learnable parameters, both in the complex domain.

3.1 Graph Definitions

Let $\mathcal{G}(V, E)$ represent a graph consisting of a finite set of N vertices, $V = \{v_1, \ldots, v_N\}$, and a set of M edges $E = \{e_1, \ldots, e_M\}$, with $e_j = (v_i, v_k)$ where $v_i, v_k \in V$. The graph's adjacency matrix $\mathbf{A}_{N \times N}$ contains 1 in case two vertices are connected and 0 otherwise. $\mathbf{D}_{N \times N}$ is a diagonal matrix where \mathbf{D}_{ii} is the degree of vertex v_i . A graph is directed if $(v_i, v_k) \neq (v_k, v_i)$, otherwise it is an undirected graph. For an undirected graph, the adjacency matrix is symmetric. The non-normalized graph Laplacian matrix is defined as $\mathbf{L} = \mathbf{D} - \mathbf{A}$, and can be normalized to $\mathbf{\bar{L}} = \mathbf{I} - \mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-\frac{1}{2}}$, where \mathbf{I} is the identity matrix. $\mathbf{\bar{L}}$ is real, symmetric, and positive semi-definite. Therefore, it has N ordered, real, and non-negative eigenvalues $\{\lambda_i : i = 1, \ldots, N\}$ and corresponding orthonormal eigenvectors $\{\mathbf{u}_i : i = 1, \ldots, N\}$.

A signal **x** defined on the nodes of the graph is a vector $\mathbf{x} \in \mathbb{R}^N$, where its *i*-th component represents the function value at the *i*-th vertex in V. Similarly, $\mathbf{X} \in \mathbb{R}^{N \times d}$ is called a d-dimensional graph signal on \mathcal{G} [56].

3.2 Graph Fourier Transform

Graph signals $\mathbf{x} \in \mathbb{R}^N$ admit a graph Fourier expansion $\mathbf{x} = \sum_{i=1}^N \langle \mathbf{u}_i, \mathbf{x} \rangle \mathbf{u}_i$, where $\mathbf{u}_i, i = 1, \dots, N$ are the eigenvectors of the graph Laplacian [56]. Eigenvalues and eigenvectors of the graph Laplacian matrix are analogous to frequencies and sinusoidal basis functions in the classical Fourier series expansion.

3.3 Spectral Graph Convolutional Network

Spectral GCNs [5] build upon the graph Fourier transform. Let x be the graph signal and y be the graph filter on graph \mathcal{G} . The graph convolution $*_{\mathcal{G}}$ can be defined as

$$\mathbf{x} *_{\mathcal{G}} \mathbf{y} = \mathbf{U}(\mathbf{U}^{\top} \mathbf{x} \odot \mathbf{U}^{\top} \mathbf{y}) \tag{1}$$

where the matrix U contains the eigenvectors of the normalized graph Laplacian and \odot is the Hadamard product. This can also be written as

$$\mathbf{x} *_{\mathcal{G}} g_{\theta} = \mathbf{U} g_{\theta}(\mathbf{\Lambda}) \mathbf{U}^{\top} \mathbf{x}, \tag{2}$$

where $g_{\theta}(\Lambda)$ is a diagonal matrix with the parameter $\theta \in \mathbb{R}^N$ as a vector of Fourier coefficients.

3.4 Spectral Graph Filter

Based on the corresponding definition of g_{θ} in Eq. (2), spectral GCNs can be classified into spectral graph filters with smooth functions and spectral graph filters with rational functions.

Spectral Graph Filter with Smooth Functions. Henaff et al. [16] proposed defining $g_{\theta}(\Lambda)$ to be a smooth function (smoothness in the frequency domain corresponds to the spatial decay), to address the localization problem.

Defferrard et al. [9] proposed defining the function g_{θ} in such a way to be directly applicable over the Laplacian matrix to address the computationally costly Laplacian matrix decomposition and multiplication with the eigenvector matrix in Eq. (2).

Kipf and Welling [21] defined $g_{\theta}(\mathbf{L})$ to be the Chebychev polynomial by assuming all the eigenvalues in the range of [-1, 1]. Computing the polynomials of the Chebychev polynomial, however, is computationally expensive. Also, considering polynomials with higher orders causes overfitting. Therefore, Kipf and Welling [21] approximated the Chebychev polynomial with its first two orders.

Spectral Graph Filter with Rational Functions. Fractional spectral GCNs, unlike polynomial spectral GCNs, can model sharp changes in the frequency response [3]. Levie et al. [23] put the eigenvalues of the Laplacian matrix on the unit circle by applying the Cayley transform on the Laplacian matrix with a learned parameter, named spectral coefficient, that lets the network focus on the most useful frequencies.

Our proposed MöbiusGCN is also a fractional GCN which applies the Möbius transformation on the eigenvalue matrix of the normalized Laplacian matrix to encode the transformations between joints.

4 MöbiusGCN

A major drawback of previous spectral GCNs is that they do not encode the transformation distribution between nodes explicitly. We address this by applying the Möbius transformation function over the eigenvalue matrix of the decomposed Laplacian matrix. This simultaneous encoding of the rotation and translation distribution in the complex domain leads to better feature representations and fewer parameters in the network.

The input to our first block of MöbiusGCN are the joint positions in 2D Euclidean space, given as $\mathcal{J} = \{J_i \in \mathbb{R}^2 | i = 1, \dots, \kappa\}$, which can be computed directly from the image. Our goal is then to predict the corresponding 3D Euclidean joint positions $\hat{\mathcal{Y}} = \{\hat{\mathcal{Y}}_i \in \mathbb{R}^3 | i = 1, \dots, \kappa\}.$

We leverage the structure of the input data, which can be represented by a connected, undirected and unweighted graph. The input graphs are fixed and share the same topological structure, which means the graph structure does not change, and each training and test example differs only in having different features at the vertices. In contrast to pose estimation, tasks like protein-protein interaction [62] are not suitable for our MöbiusGCN, because there the topological structure of the input data can change across samples.

4.1 **Möbius Transformation**

The general form of a Möbius transformation [32] is given by $f(z) = \frac{az+b}{cz+d}$ where $a, b, c, d, z \in \mathbb{C}$ satisfy $ad - bc \neq 0$. The Möbius transformation can be expressed as the composition of simple transformations. Specifically, if $c \neq 0$, then:

- $f_1(z) = z + d/c$ defines translation by d/c,
- $f_2(z) = 1/z$ defines inversion and reflection with respect to the real axis, $f_3(z) = \frac{bc-ad}{c^2}z$ defines homothety and rotation,
- $f_4(z) = z + a/c$ defines the translation by a/c.

These functions can be composed to form the Möbius transformation

$$f(z) = f_4 \circ f_3 \circ f_2 \circ f_1(z) = \frac{az+b}{cz+d}$$

where \circ denotes the composition of two functions f and g as $f \circ q(z) = f(q(z))$.

The Möbius transformation is analytic everywhere except at the pole $z = -\frac{d}{a}$. Since a Möbius transformation remains unchanged by scaling with a coefficient [32], we normalize it to yield the determinant 1. We observed that in our gradient-based optimization setup, the Möbius transformation in each node converges into the fixed points. In particular, the Möbius transformation can have two fixed points (loxodromic), one fixed point (parabolic or circular), or no fixed point. The fixed points can be computed by solving $\frac{az+b}{cz+d} = z$, which gives $\gamma_{1,2} = \frac{a-d+\sqrt{(a-d)^2-4bc}}{2c}$.

MöbiusGCN 4.2

To predict the 3D human pose, we explicitly encode the local transformations between joints, where each joint corresponds to a node in the graph. To do so, we define $g_{\theta}(\Lambda)$ in Eq. (2) to be the Möbius transformation applied to the Laplacian eigenvalues, resulting in the following fractional spectral graph convolutional network

$$\mathbf{x} *_{\mathcal{G}} g_{\theta}(\mathbf{\Lambda}) = \mathbf{U} \operatorname{M"obius}(\mathbf{\Lambda}) \mathbf{U}^{\top} \mathbf{x} = \sum_{i=0}^{N-1} \operatorname{M"obius}_{i}(\lambda_{i}) \mathbf{u}_{i} \mathbf{u}_{i}^{\top} \mathbf{x}, \quad (3)$$

where

$$\text{M\"obius}_i(\lambda_i) = \frac{a_i \lambda_i + b_i}{c_i \lambda_i + d_i},\tag{4}$$

with $a_i, b_i, c_i, d_i, \lambda_i \in \mathbb{C}$.

Applying the Möbius transformation over the Laplacian matrix places the signal in the complex domain. To return back to the real domain, we sum it up with its conjugate

$$\mathbf{Z} = 2\Re\{w \mathbf{U} \text{ M\"obius}(\mathbf{\Lambda}) \mathbf{U}^{\top} \mathbf{x}\},\tag{5}$$

where w is the shared complex-valued learnable weight to encode different transformation features. This causes the number of learned parameters to be reduced by a factor equal to the number of joints (nodes of the graph). The inter-segmental angles between joints are encoded by learning the rotation functions between neighboring nodes.

We can easily generalize this definition to the graph signal matrix $\mathbf{X} \in \mathbb{C}^{N \times d}$ with d input channels (*i.e.* a d-dimensional feature vector for every node) and $\mathbf{W} \in \mathbb{C}^{d \times F}$ feature maps. This defines a MöbiusGCN block

$$\mathbf{Z} = \sigma(2\Re\{\mathbf{U}\operatorname{M\"obius}(\mathbf{\Lambda})\mathbf{U}^{\top}\mathbf{X}\mathbf{W}\} + \mathbf{b}), \tag{6}$$

where $\mathbf{Z} \in \mathbb{R}^{N \times F}$ is the convolved signal matrix, σ is a nonlinearity (*e.g.* ReLU [36]), and b is a bias term.

To encode enriched and generalized joint transformation feature representations, we make the architecture deep by stacking several blocks of MöbiusGCN. Stacking these blocks yields our complete architecture for 3D pose estimation, as shown in Figure 2.

To apply the Möbius transformation over the matrix of eigenvalues of the Laplacian matrix, we encode the weights of the Möbius transformation for each eigenvalue in four diagonal matrices A, B, C, D and compute

$$\mathbf{U}\operatorname{M\ddot{o}bius}(\mathbf{\Lambda})\mathbf{U}^{\top} = \mathbf{U}(\mathbf{A}\mathbf{\Lambda} + \mathbf{B})(\mathbf{C}\mathbf{\Lambda} + \mathbf{D})^{-1}\mathbf{U}^{\top}.$$
(7)

4.3 Why MöbiusGCN is a Light Architecture

As a direct consequence of applying the Möbius transformation for the graph filters in polar coordinates, the filters in each block can encode the inter-segmental angle features between joints in addition to the translation features explicitly. By applying the Möbius transformation, the graph filter scales and rotates the eigenvectors of the Laplacian matrix in the graph Fourier transform simultaneously. This leads to learning better feature representations and thus, yields a more compact architecture.

For a better understanding, consider the following analogy with the classical Fourier transform: While it can be hard to construct an arbitrary signal by a linear combination of basis functions with real coefficients (i.e., the signal is built just by changing the amplitudes of the basis functions), it is significantly easier to build a signal by using *complex* coefficients, which change both the phase and amplitude.

In previous spectral GCNs, specifically [21], the Chebychev polynomials are only able to scale the eigenvectors of the Laplacian matrix, in turn requiring both more parameters and additional nonlinearities to encode the rotation distribution between joints implicitly.

4.4 Discontinuity

Our model encodes the transformation between joints in the complex domain by learning the parameters of the normalized Möbius transformation. In the definition of the Möbius transformation, if $ad - bc \neq 0$, then the Möbius transformation is an injective function and thus, continuous by the definition of continuity for neural networks given by [77]. MöbiusGCN does not suffer from discontinuities in representing intersegmental angles, in contrast to Euler angles or quaternions. Additionally, this leads to significantly fewer parameters in our architecture.

Protocol #1	# Param.	Dir.	Disc.	Eat	Greet	Phone	Photo	Pose	Purch.	Sit	SitD.	Smoke	Wait	WalkD.	Walk	WalkT.	Average
Martinez et al. [33]	4.2M	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
Tekin et al. [59]	n/a	54.2	61.4	60.2	61.2	79.4	78.3	63.1	81.6	70.1	107.3	69.3	70.3	74.3	51.8	63.2	69.7
Sun et al. [58]	n/a	52.8	54.8	54.2	54.3	61.8	67.2	53.1	53.6	71.7	86.7	61.5	53.4	61.6	47.1	53.4	59.1
Yang et al. [68]	n/a	51.5	58.9	50.4	57.0	62.1	65.4	49.8	52.7	69.2	85.2	57.4	58.4	43.6	60.1	47.7	58.6
Hossain and Little [17]	16.9M	48.4	<u>50.7</u>	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
Fang et al. [10]	n/a	50.1	54.3	57.0	57.1	66.6	73.3	53.4	55.7	72.8	88.6	60.3	57.7	62.7	47.5	50.6	60.4
Pavlakos et al. [43]	n/a	48.5	54.4	54.5	52.0	<u>59.4</u>	65.3	49.9	52.9	65.8	71.1	56.6	52.9	60.9	44.7	47.8	56.2
SemGCN [71]	0.43M	48.2	60.8	51.8	64.0	64.6	53.6	51.1	67.4	88.7	57.7	73.2	65.6	48.9	64.8	51.9	60.8
Sharma et al. [54]	n/a	48.6	54.5	54.2	55.7	62.2	72.0	50.5	54.3	70.0	78.3	58.1	55.4	61.4	45.2	49.7	58.0
GraphSH [66] *	3.7M	45.2	49.9	47.5	<u>50.9</u>	54.9	66.1	48.5	46.3	59.7	71.5	51.4	48.6	53.9	39.9	44.1	51.9
Ours (HG)	<u>0.16M</u>	<u>46.7</u>	60.7	47.3	50.7	64.1	<u>61.5</u>	46.2	45.3	67.1	80.4	<u>54.6</u>	<u>51.4</u>	55.4	<u>43.2</u>	48.6	<u>52.1</u>
Ours (HG)	0.04M	52.5	61.4	47.8	53.0	66.4	65.4	<u>48.2</u>	<u>46.3</u>	71.1	84.3	57.8	52.3	57.0	45.7	50.3	54.2
Liu et al. [27] (GT)	4.2M	36.8	<u>40.3</u>	33.0	36.3	37.5	45.0	39.7	34.9	40.3	<u>47.7</u>	37.4	38.5	38.6	29.6	32.0	37.8
GraphSH [66] (GT)	3.7M	35.8	38.1	31.0	<u>35.3</u>	35.8	43.2	<u>37.3</u>	<u>31.7</u>	38.4	45.5	35.4	36.7	36.8	27.9	30.7	35.8
SemGCN [71] (GT)	0.43M	37.8	49.4	37.6	40.9	45.1	41.4	40.1	48.3	50.1	42.2	53.5	44.3	40.5	47.3	39.0	43.8
Ours (GT)	0.16M	31.2	46.9	32.5	31.7	41.4	<u>44.9</u>	33.9	30.9	49.2	55.7	35.9	36.1	37.5	<u>29.07</u>	33.1	36.2
Ours (GT)	0.04M	33.6	48.5	34.9	34.8	46.0	49.5	36.7	33.7	50.6	62.7	38.9	40.3	41.4	33.1	36.3	40.0

Table 1: Quantitative comparisons w.r.t. MPJPE (in mm) on Human3.6M [18] under Protocol #1. Best in bold, second-best underlined. In the upper part, all methods use stacked hourglass (HG) 2D estimates [38] as inputs, except for [66] (which uses CPN [7], indicated by *). In the lower part, all methods use the 2D ground truth (GT) as input.

5 Experimental Results

5.1 Datasets and Evaluation Protocols

We use the publicly available motion capture dataset Human3.6M [18]. It contains 3.6 million images produced by 11 actors performing 15 actions. Four different calibrated RGB cameras are used to capture the subjects during training and test time. Same as previous works, *e.g.* [33, 43, 54, 58, 59, 66, 71], we use five subjects (S1, S5, S6, S7, S8) for training and two subjects (S9 and S11) for testing. Each sample from the different camera views is considered independently. We also use MPI-INF-3DHP dataset [34] to test the generalizability of our model. MPI-INF-3DHP contains 6 subjects for testing in three different scenarios: studio with a green screen (GS), studio without green screen (noGS), and outdoor scene (Outdoor). Note that for experiments on MPI-INF-3DHP we also only trained on Human3.6M.

Following [33, 58, 59, 66, 71], we use the MPJPE protocol, referred to as Protocol #1. MPJPE is the mean per joint position error in millimeters between predicted joint positions and ground truth joint positions after aligning the pre-defined root joints

(*i.e.* the pelvis joint). Note that some works (*e.g.* [27, 45]) use the P-MPJPE metric, which reports the error after a rigid transformation to align the predictions with the ground truth joints. We explicitly select the standard MPJPE metric as it is more challenging and also allows for a fair comparison to previous related works. For the MPI-INF-3DHP test set, similar to previous works [28, 66], we use the percentage of correct 3D keypoints (3D PCK) within a 150 mm radius [34] as evaluation metric.

5.2 Implementation Details

2D Pose Estimation. The inputs to our architecture are the 2D joint positions estimated from the RGB images for all four cameras independently. Our method is independent of the off-the-shelf architecture used for estimating 2D joint positions. Similar to previous works [33, 71], we use the stacked hourglass architecture [38] to estimate the 2D joint positions. The hourglass architecture is an autoencoder architecture that stacks the encoder-decoder with skip connections multiple times. Following [71], the stacked hourglass network is first pre-trained on the MPII [1] dataset and then fine-tuned on the Human3.6M [18] dataset. As described in [45], the input joints are scaled to image coordinates and normalized to [-1, 1].

3D Pose Estimation. The ground truth 3D joint positions in the Human3.6M dataset are given in world coordinates. Following previous works [33, 71], we transform the joint positions to the camera space given the camera calibration parameters. Similar to previous works [33, 71], to make the architecture trainable, we chose a predefined joint (the pelvis joint) as the center of the coordinate system. We do not use any augmentations throughout all our experiments.

We trained our architecture using Adam [20] with an initial learning rate of 0.001 and used mini-batches of size 64. The learning rate is dropped with a decay rate of 0.5 when the loss on the validation set saturates. The architecture contains seven Möbius-GCN blocks, where each block, except the first and the last block with the input and the output channels 2 and 3 respectively, contains either 64 channels (leading to 0.04M parameters) or 128 channels (leading to 0.16M parameters). We initialized the weights using the Xavier method [13]. During the test phase, the scale of the outputs is calibrated by forcing the sum of the length of all 3D bones to be equal to a canonical skeleton [42, 74, 76]. To help the architecture differentiate between different 3D poses with the same 2D pose, similar to Poier et al. [47], we provide the center of mass of the subject to the architecture as an additional input. Same as [33, 71], we predict 16 joints (*i.e.* without the 'Neck/Nose' joint).

Also, as in previous works [27, 33, 42, 71], our network predicts the normalized locations of 3D joints. We did all our experiments on an NVIDIA GeForce RTX 2080 GPU using the PyTorch framework [41]. For the loss function, same as previous works e.g. [33, 45], we use the mean squared error (MSE) between the 3D ground truth joint locations \hat{Y} and our predictions $\hat{\hat{Y}}$, *i.e.*

$$\mathcal{L}(\mathcal{Y}, \hat{\mathcal{Y}}) = \sum_{i=1}^{\kappa} (\mathcal{Y}_i - \hat{\mathcal{Y}}_i)^2.$$
(8)

Complex-valued MöbiusGCN. In complex-valued neural networks, the data and the weights are represented in the complex domain. A complex function is holomorphic

(complex-differentiable) if not only their partial derivatives exist but also satisfy the Cauchy-Riemann equations. Complex neural networks have different applications, *e.g.* Wolter and Yao [65] proposed a complex-valued recurrent neural network which helps solving the exploding/vanishing gradients problem in RNNs. Complex-valued neural networks are easier to optimize than real-valued neural networks and have richer representational capacity [60]. Considering the Liouville theorem [32], designing a fully complex differentiable (holomorphic) neural network is hard as only constant functions are both holomorph and bounded. Nevertheless, it was shown that in practice full complex differentiability of complex neural networks is not necessary [60].

In complex-valued neural networks, the complex convolution operator is defined as

$$\mathbf{W} * \mathbf{h} = (\mathbf{A} * \mathbf{x} - \mathbf{B} * \mathbf{y}) + i(\mathbf{B} * \mathbf{x} + \mathbf{A} * \mathbf{y}),$$

where $\mathbf{W} = \mathbf{A} + i\mathbf{B}$ and $\mathbf{h} = \mathbf{x} + i\mathbf{y}$. A and B are real matrices and x and y are real vectors. We also apply the same operators on our graph signals and graph filters. The PyTorch framework [41] utilizes Wirtinger calculus [22] for backpropagation, which optimizes the real and imaginary partial derivatives independently.



Fig. 3: Qualitative results of MöbiusGCN on Human3.6M [18].

5.3 Fully-supervised MöbiusGCN

In the following, we compare the results of MöbiusGCN in a fully-supervised setup with the previous state-of-the-art for 3D human pose estimation on the Human3.6M and MPI-INF-3DHP datasets. For this, we use a) estimated 2D poses using the stacked hourglass architecture (HG) [38] as input and b) the 2D ground truth (GT).

Comparisons on Human3.6M. Table 1 shows the comparison of our MöbiusGCN to the state-of-the-art methods under Protocol #1 on Human3.6M dataset.

By setting the number of channels to 128 in each block of the MöbiusGCN (0.16M parameters), given estimated 2D joint positions (HG), we achieve an average MPJPE of 52.1mm over all actions and test subjects. Using the ground truth (GT) 2D joint

positions as input, we achieve an MPJPE of 36.2mm. These results are on par with the state-of-the-art, *i.e.* GraphSH [66], which achieves 51.9mm and 35.8mm, respectively. Note, however, that MöbiusGCN drastically reduces the number of training parameters by up to 96% (0.16M vs. 3.7M).

Reducing the number of channels to 64, we still achieve impressive results considering the lightness of our architecture (*i.e.* only 0.042M parameters). Compared to GraphSH [66], we reduce the number of parameters by 98.9% (0.042M vs 3.7M) and still achieve notable results, *i.e.* MPJPE of 40.0mm vs. 35.8mm (using 2D GT inputs) and 54.2mm vs. 51.9mm (using the 2D HG inputs). Note that GraphSH [66] is the only approach which uses a better 2D pose estimator as input (CPN [7] instead of HG [38]). Nevertheless, our MöbiusGCN (with 0.16M parameters) achieves competitive results. Furthermore, MöbiusGCN outperforms the previously lightest architecture SemGCN [71], *i.e.* 40.0mm vs 43.8mm (using 2D GT inputs) and 54.2mm vs 60.8mm (using 2D HG input), although we require only 9.7% of their number of parameters (0.042M vs. 0.43M). Figure 3 shows qualitative results of our MöbiusGCN with 0.16M parameters on unseen subjects of the Human3.6M dataset given the 2D ground truth (GT) as input.

Method	# Parameters	GS	noGS	Outdoor	All(PCK)
Martinez et al. [33]	4.2M	49.8	42.5	31.2	42.5
Mehta et al. [34]	n/a	70.8	62.3	58.8	64.7
Luo et al. [28]	n/a	71.3	59.4	65.7	65.6
Yang et al. [68]	n/a	-	-	-	69.0
Zhou et al. [76]	n/a	71.1	64.7	72.7	69.2
Ci et al. [8]	n/a	74.8	70.8	77.3	74.0
Zhou et al. [72]	n/a	75.6	71.3	80.3	75.3
GraphSH [66]	3.7M	81.5	81.7	75.2	80.1
Ours	0.16M	<u>79.2</u>	<u>77.3</u>	83.1	80.0

Table 2: Results on the MPI-INF-3DHP test set [34]. Best in bold, second-best underlined.

Comparisons on MPI-INF-3DHP. The quantitative results on MPI-INF-3DHP [34] are shown in Table 2. Although we train MöbiusGCN only on the Human3.6M [18] dataset and our architecture is lightweight, the results indicate our strong generalization capabilities to unseen datasets, especially for the most challenging outdoor scenario. Figure 4 shows some qualitative results on unseen self-occlusion examples from the test set of MPI-INF-3DHP dataset with MöbiusGCN trained only on Human3.6M.



Fig. 4: Qualitative self-occlusion results of MöbiusGCN on MPI-INF-3DHP [34] (trained only on Human3.6m).

Though MöbiusGCN for 3D human pose estimation has comparably fewer parameters, it is computationally expensive, *i.e.* $O(n^3)$, both in the forward and backward pass due to the decomposition of the Laplacian matrix. In practice, however, this is not a concern for human pose estimation because of the small human pose graphs, *i.e.* ~ 20 nodes. More specifically, a single forward pass takes on average only 0.001 s.

Method	# Parameters MPJPE					
Liu et al. [27]	4.20M	37.8				
GraphSH [66]	3.70M	35.8				
Liu et al. [27]	1.05M	40.1				
GraphSH [66]	0.44M	39.2				
SemGCN [71]	0.43M	43.8				
Yan et al. [67]	0.27M	57.4				
Veličković et al. [62]	0.16M	82.9				
Chebychev-GCN	0.08M	110.6				
Ours	0.66M	33.7				
Ours	0.16M	36.2				
Ours	0.04M	40.0				

Table 3: Supervised quantitative comparison between GCN architectures on Human3.6M [18] under Protocol #1. Best in bold, second-best underlined. All methods use 2D ground truth as input.

Comparison to Previous GCNs. Table 3 shows our performance in comparison to previous GCN architectures. Besides significantly reducing the number of required parameters, applying the Möbius transformation also allows us to leverage better feature representations. Thus, MöbiusGCN can outperform all other light-weight GCN architectures. It even achieves better results (36.2mm *vs.* 39.2mm) than the light-weight version of the state-of-the-art GraphSH [66], which requires 0.44M parameters.

We also compare our proposed spectral GCN with the vanilla spectral GCN, *i.e.* Chebychev-GCN [21]. Each block of Chebychev-GCN is the real-valued spectral GCN from [21]. We use 7 blocks, similar to our MöbiusGCN, with 128 channels each. Our complex-valued MöbiusGCN with only 0.04M clearly outperforms the Chebychev-GCN [21] with 0.08M parameters (40.0mm *vs.* 110.6mm). This highlights the representational power of our MöbiusGCN in contrast to vanilla spectral GCNs.

Method	Temp	MV	Input	MPJPE
Rhodin et al. [50]	X	\checkmark	RGB	131.7
Pavlakos et al. [44]	\checkmark	\checkmark	RGB	110.7
Chen et al. [6]	X	\checkmark	HG	91.9
Li et al. [26]	\checkmark	X	RGB	88.8
Ours (0.16M)	×	X	HG	82.3
Iqbal et al. [19]	X	\checkmark	GT	62.8
Ours (0.16M)	X	X	GT	62.3

Table 4: Semi-supervised quantitative comparison on Human3.6M [18] under Protocol #1. Temp, MV, GT, and HG stand for temporal, multi-view, ground-truth, and stacked hourglass as 2D pose input respectively. Best in bold, second-best underlined.

5.4 MöbiusGCN with Reduced Dataset

A major practical limitation with training neural network architectures is to acquire sufficiently large and accurately labeled datasets. Semi-supervised methods try to address this by combining fewer labeled samples with large amounts of unlabeled data. Another benefit of MöbiusGCN is that we require fewer training samples. Having a better feature representation in MöbiusGCN leads to a light architecture and therefore, requires less training samples.

To demonstrate this, we train MöbiusGCN with a limited number of samples. In particular, we use only one subject to train MöbiusGCN and do not need any unlabeled data. Table 4 compares the MöbiusGCN to the semi-supervised approaches [6, 19, 26, 44, 50], which were trained using both labeled and unlabeled data. As can be seen, MöbiusGCN performs favorably: we achieve an MPJPE of 82.3mm (given 2D HG inputs) and an MPJPE of 62.3mm (using the 2D GT input). In contrast to previous works, we neither utilize other subjects as weak supervision nor need large unlabeled datasets during training.

As shown in Table 4, MöbiusGCN also outperforms methods which rely on multiview cues [6, 50] or leverage temporal information [26]. Additionally, we achieve better results to [19], even though, in contrast to this approach, we do not incorporate multiview information or require extensive amounts of unlabeled data during training.

Table 5 analyzes the effect of increasing the number of training samples. As can be seen, our MöbiusGCN only needs to train on three subjects to perform on par with SemGCN [71].

Subject	MöbiusGCN	MöbiusGCN	SemGCN [71]
S1	56.4 ₅₀	<u>62.3</u> 10	63.9 ₅₀
S1, S5	44.5 ₅₀	47.9 ₁₀	58.750
S1, S5, S6	42.3 ₅₀	43.1_{10}	48.550
All	36.2	36.2	43.8
# Parameters	0.16M	0.16M	0.43M

Table 5: Evaluating the effects of using fewer training subjects on Human3.6M [18] under Protocol #1 (given 2D GT inputs). The first 3 experimental results are after **10** and **50** (full convergence) training epochs for MöbiusGCN and SemGCN, respectively. Best in bold, second-best underlined.

6 Conclusion and Discussion

We proposed a novel rational spectral GCN (MöbiusGCN) to predict 3D human pose estimation by encoding the transformation between joints of the human body, given the human body joint positions in 2D. Our proposed method achieves state-of-the-art result accuracy while preserving the compactness of the model with lower number of parameters than the most compact model existing in the literature (lower number of parameters by an order of magnitude). We verified the generalizability of our model on the MPI-INF-3DHP dataset, where we achieve state-of-the-art results on the most challenging in-the-wild (outdoor) scenario.

Our proposed simple and light-weight architecture requires less data for training. This allows us to outperform the previous lightest architecture by just training our model with three subjects on the Human3.6M dataset. We also showed promising results of our architecture in comparison to previous state-of-the-art semi-supervised architectures despite not using any temporal or multi-view information or large unlabeled datasets.

Acknowledgement. This research was funded by the Austrian Research Promotion Agency (FFG) under project no. 874065 and the ERC grant no. 802554 (SPECGEO).

References

- Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2D Human Pose Estimation: New Benchmark and State of the Art Analysis. In *CVPR*, 2014. 10
- [2] Carlos Barrón and Ioannis A Kakadiaris. Estimating anthropometry and pose from a single uncalibrated image. *Comput. Vis. Image Underst.*, 81(3):269–284, 2001.
 4
- [3] Filippo Maria Bianchi, Daniele Grattarola, Lorenzo Livi, and Cesare Alippi. Graph Neural Networks with Convolutional Arma Filters. *IEEE TPAMI*, 2021. (Early access article). 6
- [4] Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. Geometric Deep Learning: Going Beyond Euclidean Data. *IEEE Signal Process. Mag.*, 34(4):18–42, 2017. 2
- [5] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. Spectral Networks and Locally Connected Networks on Graphs. In *ICLR*, 2014. 6
- [6] Xipeng Chen, Kwan-Yee Lin, Wentao Liu, Chen Qian, and Liang Lin. Weaklysupervised Discovery of Geometry-aware Representation for 3D Human Pose Estimation. In *CVPR*, 2019. 13, 14
- [7] Yilun Chen, Zhicheng Wang, Yuxiang Peng, Zhiqiang Zhang, Gang Yu, and Jian Sun. Cascaded Pyramid Network for Multi-Person Pose Estimation. In *CVPR*, 2018. 9, 12
- [8] Hai Ci, Chunyu Wang, Xiaoxuan Ma, and Yizhou Wang. Optimizing Network Structure for 3D Human Pose Estimation. In *ICCV*, 2019. 12
- [9] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional Neural Networks on Graphs with Fast Localized Spectral Filtering. In *NeurIPS*, 2016. 6
- [10] Hao-Shu Fang, Yuanlu Xu, Wenguan Wang, Xiaobai Liu, and Song-Chun Zhu. Learning Pose Grammar to Encode Human Body Configuration for 3D Pose Estimation. In AAAI, 2018. 9
- [11] Octavian Ganea, Gary Becigneul, and Thomas Hofmann. Hyperbolic Neural Networks. In *Proc. NeurIPS*, 2018. 4
- [12] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural Message Passing for Quantum Chemistry. In *ICML*, 2017. 3
- [13] Xavier Glorot and Yoshua Bengio. Understanding the Difficulty of Training Deep Feedforward Neural Networks. In AISTATS, 2010. 10
- [14] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive Representation Learning on Large Graphs. In *NeurIPS*, 2017. 3
- [15] Xintong Han, Zuxuan Wu, Zhe Wu, Ruichi Yu, and Larry S Davis. Viton: An Image-based Virtual Try-on Network. In CVPR, 2018. 1
- [16] Mikael Henaff, Joan Bruna, and Yann LeCun. Deep Convolutional Networks on Graph-structured Data. arXiv preprint arXiv:1506.05163, 2015. 3, 6
- [17] Mir Rayat Imtiaz Hossain and James J Little. Exploiting Temporal Information for 3D Human Pose Estimation. In ECCV, 2018. 9

- 16 Azizi et al.
- [18] 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 TPAMI*, 36(7):1325–1339, 2014. 3, 9, 10, 11, 12, 13, 14
- [19] Umar Iqbal, Pavlo Molchanov, and Jan Kautz. Weakly-Supervised 3D Human Pose Learning via Multi-view Images in the Wild. In *CVPR*, 2020. 4, 13, 14
- [20] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. In *ICLR*, 2015. 10
- [21] Thomas N Kipf and Max Welling. Semi-supervised Classification with Graph Convolutional Networks. In *ICLR*, 2017. 6, 8, 13
- [22] Ken Kreutz-Delgado. The Complex Gradient Operator and the CR-calculus. arXiv preprint arXiv:0906.4835, 2009. 11
- [23] Ron Levie, Federico Monti, Xavier Bresson, and Michael M Bronstein. Cayleynets: Graph Nonvolutional Neural Networks with Complex Rational Spectral Filters. *IEEE Trans. Signal Process*, 67(1):97–109, 2018. 6
- [24] Chen Li and Gim Hee Lee. Generating Multiple Hypotheses for 3D Human Pose Estimation with Mixture Density Network. In *CVPR*, 2019. **3**
- [25] Wenhao Li, Hong Liu, Runwei Ding, Mengyuan Liu, and Pichao Wang. Lifting Transformer for 3D Human Pose Estimation in Video. *arXiv preprint arXiv:2103.14304*, 2021. 1
- [26] Zhi Li, Xuan Wang, Fei Wang, and Peilin Jiang. On Boosting Single-frame 3D Human Pose Estimation via Monocular Videos. In *ICCV*, 2019. 13, 14
- [27] Kenkun Liu, Rongqi Ding, Zhiming Zou, Le Wang, and Wei Tang. A Comprehensive Study of Weight Sharing in Graph Networks for 3D Human Pose Estimation. In *ECCV*, 2020. 1, 3, 9, 10, 13
- [28] Chenxu Luo, Xiao Chu, and Alan Yuille. A Fully Convolutional Network for 3D Human Pose Estimation. In *BMVC*, 2018. 10, 12
- [29] Dingli Luo, Songlin Du, and Takeshi Ikenaga. Multi-task Neural Network with Physical Constraint for Real-time Multi-person 3D Pose Estimation from Monocular Camera. *Multimed. Tools. Appl.*, 80:27223–27244, 2021. 3
- [30] Diogo C Luvizon, David Picard, and Hedi Tabia. 2D/3D Pose Estimation and Action Recognition Using Multitask Deep Learning. In *CVPR*, 2018. 1
- [31] Xiaoxuan Ma, Jiajun Su, Chunyu Wang, Hai Ci, and Yizhou Wang. Context Modeling in 3D Human Pose Estimation: A Unified Perspective. In CVPR, 2021. 3
- [32] Danilo P Mandic and Vanessa Su Lee Goh. Complex-valued Nonlinear Adaptive Filters: Noncircularity, Widely Linear and Neural Models. John Wiley & Sons, 2009. 4, 7, 11
- [33] Julieta Martinez, Rayat Hossain, Javier Romero, and James J Little. A Simple Yet Effective Baseline for 3D Human Pose Estimation. In *ICCV*, 2017. 1, 2, 9, 10, 12
- [34] D. Mehta, H. Rhodin, D. Casas, P. Fua, O. Sotnychenko, W. Xu, and C. Theobalt. Monocular 3D Human Pose Estimation In The Wild Using Improved CNN Supervision. In 3DV, 2017. 9, 10, 12
- [35] Rahul Mitra, Nitesh B Gundavarapu, Abhishek Sharma, and Arjun Jain. Multiview-consistent Semi-supervised Learning for 3D Human Pose Estimation. In CVPR, 2020. 4

- [36] Vinod Nair and Geoffrey E Hinton. Rectified Linear Units Improve Restricted Boltzmann Machines. In Proc. ICML, 2010. 8
- [37] Mojtaba Nayyeri, Sahar Vahdati, Can Aykul, and Jens Lehmann. 5* Knowledge Graph Embeddings with Projective Transformations. In AAAI, 2021. 4
- [38] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked Hourglass Networks for Human Pose Estimation. In ECCV, 2016. 5, 9, 10, 11, 12
- [39] Necati Özdemir, Beyza B İskender, and Nihal Yılmaz Özgür. Complex-valued Neural Network with Möbius Activation Function. *Commun Nonlinear*, 16:4698– 4703, 2011. 4
- [40] Vasu Parameswaran and Rama Chellappa. View Independent Human Body Pose Estimation from a Single Perspective Image. In *CVPR*, 2004. 4
- [41] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *NeurIPS*, 2019. 10, 11
- [42] Georgios Pavlakos, Xiaowei Zhou, Konstantinos G Derpanis, and Kostas Daniilidis. Coarse-to-fine Volumetric Prediction for Single-image 3D Human Pose. In *CVPR*, 2017. 10
- [43] Georgios Pavlakos, Xiaowei Zhou, and Kostas Daniilidis. Ordinal Depth Supervision for 3D Human Pose Estimation. In CVPR, 2018. 9
- [44] Georgios Pavlakos, Nikos Kolotouros, and Kostas Daniilidis. Texturepose: Supervising Human Mesh Estimation with Texture Consistency. In *ICCV*, 2019. 13, 14
- [45] Dario Pavllo, Christoph Feichtenhofer, David Grangier, and Michael Auli. 3D Human Pose Estimation in Video with Temporal Convolutions and Semi-supervised Training. In CVPR, 2019. 1, 2, 4, 10
- [46] Xi Peng, Zhiqiang Tang, Fei Yang, Rogerio S Feris, and Dimitris Metaxas. Jointly Optimize Data Augmentation and Network Training: Adversarial Data Augmentation in Human Pose Estimation. In CVPR, 2018. 1
- [47] Georg Poier, David Schinagl, and Horst Bischof. Learning Pose Specific Representations by Predicting Different Views. In CVPR, 2018. 10
- [48] Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh. Reconstructing 3D Human Pose from 2D Image Landmarks. In ECCV, 2012. 3
- [49] Konstantinos Rematas, Ira Kemelmacher-Shlizerman, Brian Curless, and Steve Seitz. Soccer on Your Tabletop. In CVPR, 2018. 1
- [50] Helge Rhodin, Mathieu Salzmann, and Pascal Fua. Unsupervised Geometryaware Representation for 3D Human Pose Estimation. In *ECCV*, 2018. 1, 4, 13, 14
- [51] Helge Rhodin, Jörg Spörri, Isinsu Katircioglu, Victor Constantin, Frédéric Meyer, Erich Müller, Mathieu Salzmann, and Pascal Fua. Learning Monocular 3D Human Pose Estimation from Multi-view Images. In CVPR, 2018. 4
- [52] István Sárándi, Timm Linder, Kai O Arras, and Bastian Leibe. MeTRAbs: Metric-Scale Truncation-Robust Heatmaps for Absolute 3D Human Pose Estimation. *IEEE Trans. Biom. Behav. Identity Sci.*, 3(1):16–30, 2020. 3

- 18 Azizi et al.
- [53] Ashutosh Saxena, Justin Driemeyer, and Andrew Y Ng. Learning 3D Object Orientation from Images. In *ICRA*, 2009. 4
- [54] Saurabh Sharma, Pavan Teja Varigonda, Prashast Bindal, Abhishek Sharma, and Arjun Jain. Monocular 3D Human Pose Estimation by Generation and Ordinal Ranking. In *ICCV*, 2019. 9
- [55] Matthew Shere, Hansung Kim, and Adrian Hilton. Temporally Consistent 3D Human Pose Estimation Using Dual 360deg Cameras. In *ICCV*, 2021. 1
- [56] David I Shuman, Sunil K Narang, Pascal Frossard, Antonio Ortega, and Pierre Vandergheynst. The Emerging Field of Signal Processing on Graphs: Extending High-dimensional Data Analysis to Networks and Other Irregular Domains. *IEEE Signal Process. Mag.*, 30(3):83–98, 2013. 5
- [57] Cristian Sminchisescu. 3D Human Motion Analysis in Monocular Video Techniques and Challenges. In AVSS, 2006. 3
- [58] Xiao Sun, Jiaxiang Shang, Shuang Liang, and Yichen Wei. Compositional Human Pose Regression. In *ICCV*, 2017. 9
- [59] Bugra Tekin, Pablo Marquez-Neila, Mathieu Salzmann, and Pascal Fua. Learning to Fuse 2D and 3D Image Cues for Monocular Body Pose Estimation. In *ICCV*, 2017. 9
- [60] Chiheb Trabelsi, Olexa Bilaniuk, Ying Zhang, Dmitriy Serdyuk, Sandeep Subramanian, Joao Felipe Santos, Soroush Mehri, Negar Rostamzadeh, Yoshua Bengio, and Christopher J Pal. Deep Complex Networks. In *ICLR*, 2018. 11
- [61] Hsiao-Yu Fish Tung, Adam W Harley, William Seto, and Katerina Fragkiadaki. Adversarial Inverse Graphics Networks: Learning 2D-to-3D Lifting and Imageto-image Translation from Unpaired Supervision. In *ICCV*, 2017. 4
- [62] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph Attention Networks. In *ICLR*, 2018. 7, 13
- [63] Bastian Wandt, Marco Rudolph, Petrissa Zell, Helge Rhodin, and Bodo Rosenhahn. CanonPose: Self-supervised Monocular 3D Human Pose Estimation in the Wild. In *CVPR*, 2021. 4
- [64] Jianbo Wang, Kai Qiu, Houwen Peng, Jianlong Fu, and Jianke Zhu. AI Coach: Deep Human Pose Estimation and Analysis for Personalized Athletic Training Assistance. In ACM-MM, 2019. 1
- [65] Moritz Wolter and Angela Yao. Complex Gated Recurrent Neural Networks. In *NeurIPS*, 2018. 11
- [66] Tianhan Xu and Wataru Takano. Graph Stacked Hourglass Networks for 3D Human Pose Estimation. In CVPR, 2021. 3, 9, 10, 12, 13
- [67] Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial Temporal Graph Convolutional Networks for Skeleton-based Action Recognition. In AAAI, 2018. 13
- [68] Wei Yang, Wanli Ouyang, Xiaolong Wang, Jimmy Ren, Hongsheng Li, and Xiaogang Wang. 3D Human Pose Estimation in the Wild by Adversarial Learning. In *CVPR*, 2018. 9, 12
- [69] Yuan Yao, Yasamin Jafarian, and Hyun Soo Park. Monet: Multiview Semisupervised Keypoint Detection via Epipolar Divergence. In *ICCV*, 2019. 4
- [70] Zhe Zhang, Chunyu Wang, Weichao Qiu, Wenhu Qin, and Wenjun Zeng. Ada-Fuse: Adaptive Multiview Fusion for Accurate Human Pose Estimation in the Wild. *IJCV*, 129:703–718, 2021. 1

- [71] Long Zhao, Xi Peng, Yu Tian, Mubbasir Kapadia, and Dimitris N. Metaxas. Semantic Graph Convolutional Networks for 3D Human Pose Regression. In *CVPR*, 2019. 2, 3, 9, 10, 12, 13, 14
- [72] Kun Zhou, Xiaoguang Han, Nianjuan Jiang, Kui Jia, and Jiangbo Lu. Hemlets pose: Learning part-centric heatmap triplets for accurate 3D human pose estimation. In *ICCV*, 2019. 12
- [73] Sharon Zhou, Jiequan Zhang, Hang Jiang, Torbjörn Lundh, and Andrew Y Ng. Data Augmentation with Möbius Transformations. *Mach. learn.: sci. technol.*, 2 (2):025016, 2021. 4
- [74] 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 TPAMI*, 41(4): 901–914, 2018. 10
- [75] Xingyi Zhou, Xiao Sun, Wei Zhang, Shuang Liang, and Yichen Wei. Deep Kinematic Pose Regression. In ECCV, 2016. 4
- [76] Xingyi Zhou, Qixing Huang, Xiao Sun, Xiangyang Xue, and Yichen Wei. Towards 3D Human Pose Estimation in the Wild: a Weakly-supervised Approach. In *ICCV*, 2017. 10, 12
- [77] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the Continuity of Rotation Representations in Neural Networks. In *CVPR*, 2019. 4, 9