

Invariant Teacher and Equivariant Student for Unsupervised 3D Human Pose Estimation

Chenxin Xu, Siheng Chen*, Maosen Li, Ya Zhang*

Cooperative Medianet Innovation Center, Shanghai Jiao Tong University
{xcxwakaka, sihengc, maosen_li, ya_zhang}@sjtu.edu.cn

Abstract

We propose a novel method based on teacher-student learning framework for 3D human pose estimation without any 3D annotation or side information. To solve this unsupervised-learning problem, the teacher network adopts pose-dictionary-based modeling for regularization to estimate a physically plausible 3D pose. To handle the decomposition ambiguity in the teacher network, we propose a cycle-consistent architecture promoting a 3D rotation-invariant property to train the teacher network. To further improve the estimation accuracy, the student network adopts a novel graph convolution network for flexibility to directly estimate the 3D coordinates. Another cycle-consistent architecture promoting 3D rotation-equivariant property is adopted to exploit geometry consistency, together with knowledge distillation from the teacher network to improve the pose estimation performance. We conduct extensive experiments on Human3.6M and MPI-INF-3DHP. Our method reduces the 3D joint prediction error by 11.4% compared to state-of-the-art unsupervised methods and also outperforms many weakly-supervised methods that use side information on Human3.6M. Code will be available at <https://github.com/sjtxcx/ITES>.

1 Introduction

3D human pose estimation from 2D landmarks receives substantial attention due to its broad applications (Li et al. 2019, 2020a). Many fully-supervised or weakly-supervised estimating algorithms have been proposed, which require either adequate 3D ground-truth annotations or side information such as unpaired 3D poses (Tung et al. 2017) and multi-view images (Kocabas, Karagoz, and Akbas 2019) as supervision. However, obtaining such 3D annotation or side information is time-consuming and prohibitive, which is not feasible for many applicable scenarios. In this paper, we consider the unsupervised-learning setting, which estimates a 3D pose from a given 2D pose without any additional information in either the training or the inference phase.

Estimating a 3D pose solely from a given 2D pose is an ill-posed problem because there exists an infinite number of 3D skeletons corresponding to the same 2D pose with the perspective projection ambiguity. Some simple self-supervision

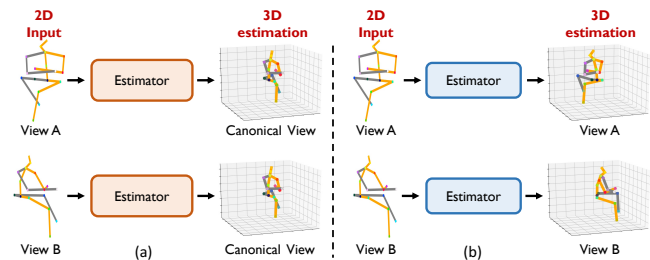


Figure 1: (a) 3D rotation-invariant property. Different input views result in the same view-invariant 3D estimation. (b) 3D rotation-equivariant property. The 3D estimation has the same view as the input views.

methods through straightforward reprojection might be under-constrained, resulting in overfitting. Extra regularization or priors of a human body is required to eliminate solutions that are not physically plausible. However, too much regularization tends to limit the flexibility of the model and leads to underfitting. An appropriate trade-off between regularization and flexibility is usually hard to reach, which leads to the regularization-flexibility dilemma.

To resolve the above dilemma, we propose a teacher-student framework to balance the regularization and flexibility, which achieves a coarse-to-fine estimation. The proposed pose dictionary-based teacher network emphasizes the regularization, ensuring the feasibility of the pose estimation. The proposed graph-convolution-based student network emphasizes the flexibility, further refining the estimation. To enable unsupervised estimation, both teacher and student networks are trained based on respective cycle-consistent architectures.

Following (Novotny et al. 2019), the teacher network adopts a pose-dictionary-based modeling, where the 3D pose is decomposed into the camera viewpoint and a linear combination of pose atoms in a pose dictionary. However, this pose modeling suffers from decomposition ambiguity (Dai, Li, and He 2014), with an infinite number of valid camera-pose pairs. To resolve decomposition ambiguity, we propose a novel cycle-consistent architecture promoting the 3D rotation-invariant property to train the teacher network. The 3D rotation-invariant property means that 2D poses from different views of a same 3D skeleton should lead to the same

*Corresponding author

3D pose and different camera viewpoints, thus making the decomposition unique; see Figure 1(a). In the cycle-consistent architecture, the estimated 3D pose is randomly rotated and reprojected to obtain a new 2D pose, whose estimation is forced to be the same as the original 2D input through a cycle-consistent loss.

Our student network leverages a trainable skeleton graph structure capturing joints correlation to directly estimate the 3D pose in the input view. The student network receives the knowledge from the teacher and further improves the estimation performance without rigid regularization. Since the student network provides the final estimated 3D pose, which should be adaptive to the input camera view, we propose another cycle-consistent architecture to promote the 3D rotation-equivariant property for the student network. The 3D rotation-equivariance means that the 3D estimation has the same view as the input view; see Figure 1(b). The cycle-consistent architecture provides a self-supervision way by exploiting geometric consistency and enhances the training of the student network.

Overall, the proposed method is characterized by a rotation-Invariant Teacher and a rotation-Equivariant Student (named ITES thereafter). Different from many traditional teacher-student frameworks (Zhou et al. 2017a), our student network refines the coarse teacher estimation result and outperforms the teacher network. It is still worth noting that the student network can not be trained alone without the teacher network because no extra constraints or priors are added into the network to eliminate solutions not physically plausible. The effectiveness of the proposed ITES is validated comprehensively on standard 3D benchmarks. Experimental results show that our approach outperforms the state-of-the-art unsupervised methods on Human3.6M (Ionescu et al. 2013) and MPI-INF-3DHP (Mehta et al. 2017). We furthermore visualize the estimation provided by the proposed ITES, which demonstrates the effectiveness of our approach qualitatively.

The main contributions of this paper are as follow:

- We propose an unsupervised teacher-student learning framework, called invariant teacher and equivariant student (ITES), to estimate the 3D human pose from a given 2D pose, without requiring additional annotation or side information.
- We propose a cycle-consistent architecture to train a pose-dictionary-based teacher network. This architecture promotes the 3D rotation-invariant property for the estimation.
- We propose a similar, yet different cycle-consistent architecture exploiting geometry consistency to enhance the training of the graph-convolution-based student network. This architecture promotes the 3D rotation-equivariant property.
- We conduct extensive experiments on Human3.6M and MPI-INF-3DHP datasets. We show that our ITES outperforms the state-of-the-art unsupervised method by 11.4% (51.4mm vs 58mm) measured by P-MPJPE on Human3.6M.

2 Related Works

2.1 Non-Rigid Structure from Motion

Non-rigid structure from motion (NRSfM) considers the problem of reconstructing the 3D shapes of non-rigid objects and the viewpoint from 2D points as those objects are moving

along with the corresponding shape deforming. NRSfM becomes an ill-posed problem since the number of unknown variables that need to be estimated is more than the number of known equations. To alleviate this problem, extra constraints need to be added. Bregler et al. (Bregler, Hertzmann, and Biermann 2000) proposed a constraint to model the 3D shapes of objects in a low-rank subspace. Based on this modeling, various constraints about the shapes and the viewpoints are exploited, including the low-rank subspaces in the spatial domain (Dai, Li, and He 2014; Fragkiadaki et al. 2014), fixed articulation (Ramakrishna, Kanade, and Sheikh 2012), low-rank subspaces in the temporal domain (Akhter et al. 2009, 2010), the union-of-subspaces (Zhu et al. 2014; Agudo, Pijoan, and Moreno-Noguer 2018) and block-sparsity (Kong and Lucey 2016). Many of these previous approaches consider arbitrary non-rigid objects and deal with small-scale datasets with thousands of images, while we focus on the human body and the proposed method can be applied to handle large-scale datasets with millions of images.

2.2 Weakly-Supervised 3D Pose Estimation

Weakly-supervised 3D pose estimation does not directly use the corresponding 3D ground truth pose of the input 2D pose or image. Instead, they use a collection of unpaired 3D pose (Tung et al. 2017; Wandt and Rosenhahn 2019; Zhou et al. 2016), multi-view supervision (Kocabas, Karagoz, and Akbas 2019; Rhodin et al. 2018; Li et al. 2020b) or ordinal depth supervision (Pavlakos, Zhou, and Daniilidis 2018) to alleviate the need for paired 2D-3D annotations. However, those auxiliary supervisions are still expensive to access. In this paper, we propose an unsupervised method that only uses a single 2D pose without auxiliary information.

2.3 Unsupervised 3D Pose Estimation

Due to the lack of 3D ground truth annotations, extra constraints are needed to solve the unsupervised 3D pose estimation problem. (Drover et al. 2018; Kudo et al. 2018) introduce implicit constraints based on adversarial training. Following this branch, (Chen et al. 2019) proposes a domain adaptation method with a geometric self-supervision loss. Instead of using adversarial training, (Novotny et al. 2019) considers the constraints based on low-rank matrix factorization and the canonicalization of 3D shapes. Recently, (Wang, Kong, and Lucey 2019) proposes a teacher-student framework, where the student is guided by a teacher to alleviate the projection ambiguity. In this work, we also consider a teacher-student framework. Note that (Wang, Kong, and Lucey 2019) adds a block sparsity constraint into the teacher network structure, while our framework promotes the 3D rotation-invariant property for the teacher network. We also propose a novel graph-convolution-based student network promoting the rotation-equivariance comparing to the convolutional-based student network in (Wang, Kong, and Lucey 2019).

3 Invariant Teacher and Equivariant Student

In this section, we formulate the task of 3D pose estimation and propose our estimation framework, called invariant teacher and equivariant student (ITES). Let $\mathbf{X} =$

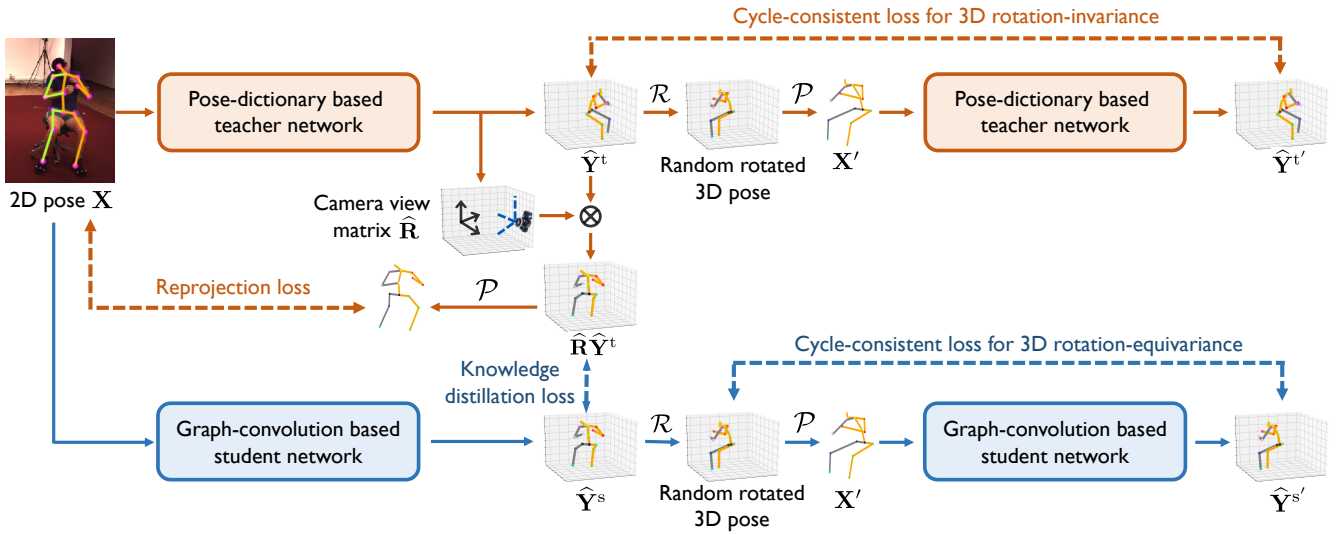


Figure 2: An overview of ITES. To train the teacher network (orange branch), a reprojection loss and a cycle-consistent loss for 3D rotation-invariance are applied. The student network (blue branch) is trained by knowledge distillation and a cycle-consistent loss for the 3D rotation-equivariance. \mathcal{R} and \mathcal{P} represent the random rotation and perspective projection operation.

$[\mathbf{x}_1 \cdots \mathbf{x}_N] \in \mathbb{R}^{2 \times N}$ be a 2D pose matrix, where N is the number of body joints and $\mathbf{x}_i \in \mathbb{R}^2$ is the 2D coordinate of the i th body joint. The root joint (pelvis) is set to the origin. Given this 2D pose matrix \mathbf{X} , a 3D pose estimator $\mathcal{E}(\cdot)$ is designed to estimate the corresponding 3D pose matrix $\hat{\mathbf{Y}} \in \mathbb{R}^{3 \times N}$; that is, $\hat{\mathbf{Y}} = \mathcal{E}(\mathbf{X}) \in \mathbb{R}^{3 \times N}$. We apply a perspective camera model as (Chen et al. 2019) by assuming a virtual camera with an intrinsic matrix of \mathbf{I}_3 centered at the world origin and fix the distance from skeleton to the camera in the 3D space with a constant t unit.

3.1 Framework Overview

We consider a teacher-student framework to address the trade-off between regularization and flexibility. The teacher network considers a pose-dictionary-based model limiting the solution space spanned by the pose atoms for regularization. To train this teacher network, we propose a cycle-consistent loss for 3D rotation-invariance, see details in Section 3.2. The student network captures the correlations between body joints and reflects the physical constraints of the human body by leveraging a trainable graph structure. To train this graph-convolution-based student network, we consider knowledge distilling from the teacher network and a cycle-consistent loss for 3D rotation-equivariance; see details in Section 3.3.

3.2 Pose-Dictionary-based Teacher Network

The specific aim of the teacher network is to recover the corresponding 3D pose and camera-view matrix from a given 2D pose. To achieve this, we consider two features for our design: the pose modeling based on a pose dictionary, which regularizes the solution space of 3D poses, as well as a cycle-consistent estimation architecture, which promotes the rotation-invariant property of 3D pose estimation.

Pose modeling. Following the common representation of

3D poses in NRSfM methods (Bregler, Hertzmann, and Biermann 2000) we model a 3D pose in our teacher network \mathbf{Y}^t as a linear combination of K pose atoms $\mathbf{B}_k \in \mathbb{R}^{3 \times N}$, $k = 1, \dots, K$; that is $\mathbf{Y}^t = \sum_{k=1}^K c_k \mathbf{B}_k$, where c_k is the coefficient of the k -th pose atom. We can rewrite the formula in a matrix representation, $\mathbf{Y}^t = (\mathbf{c} \otimes \mathbf{I}_3) \mathbf{B} \in \mathbb{R}^{3 \times N}$ where $\mathbf{B} = [\mathbf{B}_1^\top \cdots \mathbf{B}_K^\top]^\top \in \mathbb{R}^{3K \times N}$ is a pose dictionary, consisting of K pose atoms; $\mathbf{c} = [c_1 \cdots c_K]^\top \in \mathbb{R}^K$ is the coefficient vector and \otimes denotes the Kronecker product. The 2D pose can be projected from \mathbf{Y}^t as

$$\mathbf{X} = \mathcal{P}(\mathbf{R}\mathbf{Y}^t) = \mathcal{P}(\mathbf{R}(\mathbf{c}^\top \otimes \mathbf{I}_3)\mathbf{B}) \in \mathbb{R}^{2 \times N}, \quad (1)$$

where \mathcal{P} is the perspective projection function, which maps the 3D pose on a 2D plane, and $\mathbf{R} \in SO(3)$ is the camera view matrix.

Network architecture. Motivated by Eq. (1), we design a pose-dictionary-based teacher network to solve an inverse problem; that is, given the 2D pose \mathbf{X} , we aim to estimate the pose dictionary \mathbf{B} , the coefficient \mathbf{c} and the camera view matrix \mathbf{R} through a neural network.

The network consists of three parts: a backbone network that extracts features from the input 2D pose, two output heads that produce the coefficients $\hat{\mathbf{c}}$ and the camera view matrix $\hat{\mathbf{R}}$, respectively, and a pose-generation module that produces the estimated 3D pose. The backbone and two output heads are constructed by fully-connected layers. In the pose-generation module, we multiply a trainable matrix of the pose-dictionary \mathbf{B} with the pose coefficients $\hat{\mathbf{c}}$ to generate a 3D pose; that is, $\hat{\mathbf{Y}}^t = (\hat{\mathbf{c}}^\top \otimes \mathbf{I}_3)\mathbf{B} \in \mathbb{R}^{3 \times N}$; see Figure 3. Note that \mathbf{B} works as trainable parameters shared across different inputs, while $\hat{\mathbf{c}}$ and $\hat{\mathbf{R}}$ are the corresponding outputs of the teacher network.

To train this network, we consider the following two losses: a self-supervised 2D reprojection loss and a cycle-consistent




Figure 3: The pose-dictionary-based teacher network.

loss for 3D rotation-invariance.

Self-supervised 2D reprojection loss. To promote the equality in Eq.(1), the teacher network is learned in a self-supervised way by minimizing the 2D reprojection loss

$$\mathcal{L}_{\text{REP}} = \frac{1}{N} \left\| \mathcal{P}(\widehat{\mathbf{R}}\widehat{\mathbf{Y}}^t) - \mathbf{X} \right\|_{\text{F}}^2, \quad (2)$$

where $\|\cdot\|_{\text{F}}$ indicates the Frobenius norm, $\widehat{\mathbf{Y}}^t$ is the estimation of the teacher network, the camera view matrix $\widehat{\mathbf{R}}$ rotates the estimation 3D pose into the input view, and \mathcal{P} is the perspective projection function that projects a 3D pose to a 2D pose. The 2D reprojection loss constrains the 2D projection of the estimated 3D pose close to input 2D pose in a self-supervised manner.

Cycle-consistent loss for 3D rotation-invariance. According to Eq. (1), one 2D pose can actually be decomposed into multiple pairs of 3D poses and camera views. To be specific, given a 2D pose \mathbf{X} , any rotation matrix $\mathbf{G} \in SO(3)$ could introduce a new pair of the rotation matrix $\mathbf{R}\mathbf{G}$ and the 3D pose $\mathbf{G}^{-1}\mathbf{Y}^t$, leading to a new valid decomposition, $\mathbf{X} = \mathcal{P}(\mathbf{R}\mathbf{Y}^t) = \mathcal{P}((\mathbf{R}\mathbf{G})(\mathbf{G}^{-1}\mathbf{Y}^t))$. This decomposition ambiguity causes a problem that there are innumerable solutions to generate 3D human poses from a single 2D pose.

To address this issue, we propose a cycle-consistent loss promoting 3D rotation-invariance to enable the potentially finite solutions of the teacher network. The intuition is that when we project a 3D pose to various 2D poses along multiple views, the teacher network should estimate the same 3D pose from those projected 2D poses. To achieve this, we rotate the estimated 3D pose $\widehat{\mathbf{Y}}^t$ by a random rotation matrix $\mathbf{R}_{\text{rand}} \in SO(3)$, and then project the 3D pose on a 2D plane to obtain a new 2D pose $\mathbf{X}' \in \mathbb{R}^{2 \times N}$. We next input this new 2D pose \mathbf{X}' to the same teacher network, producing the 3D pose estimation $\widehat{\mathbf{Y}}^{t'}$ once again; see Figure 2. We then minimize the cycle-consistent loss \mathcal{L}_{RIC} to restrict $\widehat{\mathbf{Y}}^{t'}$ to be close to $\widehat{\mathbf{Y}}^t$; that is,

$$\mathcal{L}_{\text{RIC}} = \frac{1}{N} \left\| \mathcal{F}_t(\mathcal{P}(\mathbf{R}_{\text{rand}}\widehat{\mathbf{Y}}^t)) - \widehat{\mathbf{Y}}^t \right\|_{\text{F}}^2, \quad (3)$$

where $\mathcal{F}_t(\cdot)$ is the teacher network. This consistent 3D pose estimation architecture constrains the teacher network to estimate the canonical 3D pose from various 2D poses.

Some previous works also tried to promote the rotation-invariant property. For example, C3DPO (Novotny et al. 2019) uses an extra canonicalization neural network to recover a random rotation and estimate a canonical Euclidean




Figure 4: Adaptive graph convolution block.

form of a 3D pose. Compared to those previous methods, the proposed consistent teacher network does not introduce any additional network to handle the decomposition ambiguity, which not only reduces the number of training parameters but also improves empirical performances.

3.3 Graph-Convolution-based Student Network

The teacher network with its training architecture makes the estimation problem feasible; however, the learning ability of the teacher network is also limited by linear approximation. To improve the flexibility, we propose a novel graph convolutional network as a student network, which is pose-dictionary free and leverages graphs to model the correlations between body joints. The teacher network estimates the 3D coordinates of the pose because the pose is represented by a linear combination of pose atoms, providing strong regularization. In the student network, we emphasize the flexibility, thus we can simplify the problem to the depth estimation to ensure the reprojection loss is always zero. Inspired by (Chen et al. 2019), the student network estimates the depth offset of every joint to the root joint. Given an input 2D pose \mathbf{X} , the proposed student network outputs a N -dimensional row vector $\widehat{\mathbf{d}} = [\widehat{d}_1, \widehat{d}_2, \dots, \widehat{d}_N]$, whose i th element \widehat{d}_i represents the depth offset from the i th body joint to the root joint. Suppose the i th input 2D coordinate is (u_i, v_i) , utilizing the perspective camera model, the 3D coordinates of the i th body-joint $(\widehat{\mathbf{Y}}^s)_i$ is then $(u_i z_i, v_i z_i, \widehat{d}_i)$, where $z_i = \max(1, t + \widehat{d}_i)$ and t is the constant distance between the camera and root joint. To achieve the depth estimation, we consider an adaptive graph convolutional network.

Network structure. The key component of the network is an adaptive graph convolution layer that consists of a physical-graph convolution operation and a nonphysical-graph convolution operation. Let \mathbf{A} be the adjacency matrix that defines the physical connections among body joints and $\mathbf{H} \in \mathbb{R}^{D \times N}$, $\mathbf{H}' \in \mathbb{R}^{D' \times N}$ be the features of body joints before and after the physical-adaptive graph convolution layer.

For the physical-graph convolution, we leverage the physical constraints of a human body to extract features. Different from the common graph convolution (Kipf and Welling 2016), the proposed physical-graph convolution operation introduces a set of trainable edge-weight matrix $\mathcal{M} = \{\mathbf{M}_d \mid \mathbf{M}_d \in \mathbb{R}^{N \times N}, d = 1, 2, \dots, D'\}$ to adaptively model joints' correlations (Zhao et al. 2019) and work as

$$\mathbf{H}_{\text{P}} = \bigcup_{d=1}^{D'} \{\sigma(\mathbf{w}_d \mathbf{H} \rho(\mathbf{M}_d \odot \mathbf{A}))\} \in \mathbb{R}^{D' \times N}, \quad (4)$$

where \cup is the concatenate operation along the first dimension, $\sigma(\cdot)$ is the nonlinear activation function, such as the ReLU function, $\mathbf{w}_d \in \mathbb{R}^{1 \times D}$ is the trainable weight vector, $\rho(\cdot)$ is the column-wise softmax operation, which normalizes the edge weights and enhances the highest relationships and \odot is an element-wise multiplier. Each edge-weight matrix \mathbf{M}_d and its associated weight vector \mathbf{w}_d correspond to one feature channel for the output features of body joints \mathbf{H}_P . They are trainable and can adaptively adjust the edge weight between each pair of body joints, providing more flexibility to capture the correlations of the body joints.

The physical-graph convolution operation relies on the physical connections provided by the human skeleton. To model the nonphysical connections, we further propose a nonphysical-graph convolution operation to capture the correlations between an arbitrary pair of body joints. The proposed nonphysical-graph convolution operation works as

$$\mathbf{H}_{NP} = f(g(\mathbf{H})\rho(\mathbf{H}^\top \mathbf{W}_1^\top \mathbf{W}_2 \mathbf{H})) \in \mathbb{R}^{D' \times N}, \quad (5)$$

where $f(\cdot)$ and $g(\cdot)$ are two multi-layer perceptrons that operate along the feature dimension and are shared by all the body joints, $\mathbf{W}_1 \in \mathbb{R}^{D \times D}$ and $\mathbf{W}_2 \in \mathbb{R}^{D \times D}$ are two trainable weight matrices and $\rho(\cdot)$ is a softmax operation on each column. The N -by- N matrix $\rho(\mathbf{H}^\top \mathbf{W}_1^\top \mathbf{W}_2 \mathbf{H})$ is data-adaptive and models the affinities between all pairs of body points, capturing the nonlocal and nonphysical relationships. Its matrix multiplication with $g(\mathbf{H})$ thus aggregates the features of all the body joints with nonuniform attentions.

To integrally capture the features based on physical and nonphysical connections among body joints, we combine two operations and propose an adaptive graph convolutional layer, which is formulated as $\mathbf{H}' = \mathbf{H}_P + \mathbf{H}_{NP} \in \mathbb{R}^{D' \times N}$.

By using the adaptive graph convolutional layer, we construct the graph-convolution-based student network. Figure 4 shows a graph convolution block in our student network. The graph-convolution-based student network finally produces the depth $\hat{\mathbf{d}}$, which can be further converted into 3D coordinates $\hat{\mathbf{Y}}^s$.

Knowledge distilling loss. A single graph-convolution-based student network is hard to estimate a 3D pose as there is no depth supervision for the network training. We thus design a knowledge distillation loss to supervise the student network by the output of the teacher network,

$$\mathcal{L}_{KD} = \frac{1}{N} \left\| \hat{\mathbf{d}} - (\hat{\mathbf{R}}\hat{\mathbf{Y}}^t)_z \right\|_2^2, \quad (6)$$

where the row vector $\hat{\mathbf{d}} \in \mathbb{R}^N$ is the depth estimated by the student network, $\hat{\mathbf{R}}$ and $\hat{\mathbf{Y}}^t$ are the estimated camera-view matrix and the estimated 3D pose output by the teacher network, respectively, and $(\cdot)_z$ slices the third dimension, which represents the depths of body joints.

Cycle-consistent loss for 3D rotation-equivariance. We use an additional cycle-consistent loss for 3D rotation-equivariance similarly to the teacher network to further improve the performance; see Figure 2. Note that the student network estimates the depth, which relies on the view of input. It thus promotes 3D rotation-equivariance, instead of 3D rotation-invariance as in the teacher network. To be specific,

different 2D poses corresponding to a same 3D pose but from different views should produce rotation-equivariant estimation results. Let $\hat{\mathbf{Y}}^s \in \mathbb{R}^{3 \times N}$ denotes the 3D pose estimated by the student network. We rotate 3D pose $\hat{\mathbf{Y}}^s$ by a random rotation matrix $\mathbf{R}_{rand} \in SO(3)$ and project it to generate a new 2D pose $\mathbf{X}' = \mathcal{P}(\mathbf{R}_{rand}\hat{\mathbf{Y}}^s)$. When we input \mathbf{X}' to the student network, we expect the resulting 3D pose should be close to $\mathbf{R}_{rand}\hat{\mathbf{Y}}^s$. The equation of the cycle-consistent loss for 3D rotation-equivariance is thus

$$\mathcal{L}_{REC} = \frac{1}{N} \left\| \mathcal{F}_s(\mathcal{P}(\mathbf{R}_{rand}\hat{\mathbf{Y}}^s)) - \mathbf{R}_{rand}\hat{\mathbf{Y}}^s \right\|_F^2, \quad (7)$$

where $\mathcal{F}_s(\cdot)$ is the student network. \mathcal{L}_{REC} promotes that the information about the random 3D rotation matrix would be preserved in the 3D estimation of the student network.

Both the cycle-consistent loss in the teacher network and student network provide self-supervisions to train the networks, but they are significantly different. For the teacher network, the cycle-consistent loss for rotation-invariance handles the perspective projection ambiguity. On the other hand, the cycle-consistent loss for rotation-equivariance ensures the student network adaptive to the camera view and enhances the training. We further present the effects of these two losses in our ablation study; see Sec. 4.4.

3.4 Framework Training

We consider a two-stage training procedure:

1. We optimize trainable weights in the pose-dictionary-based teacher network by minimizing $\mathcal{L}^t = \lambda_{REP}\mathcal{L}_{REP} + \lambda_{RIC}\mathcal{L}_{RIC}$.

2. We freeze the teacher network and optimize the graph-convolution-based student network by minimizing $\mathcal{L}^s = \lambda_{KD}\mathcal{L}_{KD} + \lambda_{REC}\mathcal{L}_{REC}$.

λ_{REP} , λ_{RIC} , λ_{KD} and λ_{REC} are weight hyperparameters. Finally, we use the estimated 3D pose from the student network as our final output 3D pose; that is, $\hat{\mathbf{Y}} = \hat{\mathbf{Y}}^s$.

4 Experiments

In this section, we conduct extensive experiments to evaluate our model. We show that our ITES outperforms the state-of-the-art method in terms of both quantitative results and qualitative results. We also conduct ablation studies to discuss the effect of various losses, hyperparameters and sub-networks.

4.1 Datasets and Metrics

Human3.6M. Human3.6M (Ionescu et al. 2013) is a large-scale dataset used widely in 3D pose estimation. There are 3.6 million video frames performed by 11 subjects, of which 7 are annotated with 3D poses. It performs 15 classes of actions captured from 4 camera-views. Following the previous works (Martinez et al. 2017), we consider the poses of the subjects S1, S5, S6, S7, and S8 for training, and use S9 and S11 for testing. We take two types of 2D poses as inputs: the ground truth 2D annotations (GT) and the 2D poses estimated from images (IMG) by CPN (Chen et al. 2018), and we respectively test the model on them. Note that we do not train class-specific model or use the multi-view information.

MPI-INF-3DHP. We also use another large scale 3D human pose dataset, MPI-INF-3DHP (Mehta et al. 2017), which includes poses in both indoor and outdoor scenes. Following the previous works (Kanazawa et al. 2018; Mehta et al. 2017; Zhou et al. 2017b), we train the model on Human3.6M, while we test it on the test set of MPI-INF-3DHP.

Metrics. On Human3.6M, we use two quantitative evaluation protocols: 1) the mean per-joint position error (MPJPE), which is the mean Euclidean distance between predicted joints and the ground-truths. For fair comparisons with (Novotny et al. 2019; Kudo et al. 2018), we use the same normalization method to scale the estimated poses before calculating MPJPE; and 2) the Procrustes analysis MPJPE (P-MPJPE), which measures the mean error after alignment to the ground truth in terms of translation, rotation, and scale. On MPI-INF-3DHP, we use the percentage of correct key-points@150mm (PCK) and the corresponding Area Under Curve (AUC) computed for a range of PCK thresholds.

4.2 Experimental Setups

We set the distance between the camera and the root joint (pelvis) t as 5 unit. Similar to the network structure proposed in (Novotny et al. 2019), the feature-extraction module in the teacher network consists of 6 fully-connected residual layers. Each with 1024/256/1024 neurons. In the teacher network, we set the size of the pose dictionary K to 12. In the student network, we use 8 graph convolution blocks and a residual connection is built across consecutive two blocks. We train the entire framework with the SGD (Bottou 2010) optimizer. In the first training stage, we train the teacher network with the learning rate 0.001 for 40 epoches. In the second training stage, we train the student network for 30 epoches with the learning rate 0.001. The weight parameter is set as $\lambda_{REP} = 5$, $\lambda_{RIC} = 1$, $\lambda_{KD} = 5$, $\lambda_{REC} = 1$.

4.3 Comparison with Existing Works

Quantitative results. To evaluate the effectiveness of the proposed ITES, we first compare our model with state-of-the-art methods on Human3.6M, including Chen & Ramanan (Chen and Ramanan 2017), Martinez et al. (Martinez et al. 2017), IGE-Net (Jack et al. 2019), Ci et al. (Ci et al. 2019), Pavlo et al. (Pavlo et al. 2019), Metha et al. (Mehta et al. 2017), HMR (Kanazawa et al. 2018), SPIN (Kolotouros et al. 2019), 3DInterpreter (Wu et al. 2016), AIGN (Tung et al. 2017), RepNet (Wandt and Rosenhahn 2019), Drover et al. (Drover et al. 2018), Li et al. (Li et al. 2020b), Zhou et al. (Zhou et al. 2017b), Pose-GAN (Kudo et al. 2018), C3DPO (Novotny et al. 2019), Nath et al. (Nath Kundu et al. 2020), Wang et al. (Wang, Kong, and Lucey 2019) and Chen et al. (Chen et al. 2019).

Table 1 presents the performance of various models. We report the result of the teacher network after training stage 1 (ITES-T) and the final estimation result of the student network (ITES-TS). We see that i) our method outperforms all the state-of-the-art unsupervised methods and significantly reduces the pose estimation errors by 9.8 mm and 4.6 mm for MPJPE and P-MPJPE in average; ii) we even outperform several weakly-supervised methods that use depth information or multi-view images in training.

Type	Method	MPJPE		P-MPJPE	
		GT	IMG	GT	IMG
Fully	Chen & Ramanan	-	114.2	57.5	82.7
	Martinez et al.	45.5	62.9	37.1	52.1
	IGE-Net	42.6	66.0	37.7	50.5
	Ci et al.	36.3	52.7	27.9	42.2
	Pavlo et al.	37.2	46.8	27.2	36.5
Weak	3DInterpreter	-	-	88.6	98.4
	AIGN	-	-	79.0	97.2
	RepNet	50.9	89.9	38.2	65.1
	Drover et al.	-	-	38.2	64.6
	Li et al.	-	57.0	-	44.1
Unsup	Pose-GAN	130.9	173.2	-	-
	C3DPO	95.6	145.0	-	-
	Nath et al.	-	-	63.8	89.4
	Wang et al.	-	86.4	-	62.8
	Chen et al.	-	-	58.0	-
	ITES-T	85.6	93.7	54.9	62.5
	ITES-TS	77.2	85.3	51.9	59.8

Table 1: Results on Human3.6M. We compare our methods with fully supervised (Fully), Weakly-supervised (Weak) and Unsupervised (Unsup) methods. GT and IMG denote the input data represented by 2D ground truth keypoints and 2D detected keypoints.

Type	Method	Training Set	PCK	AUC
Fully	Metha et al.	H3.6M	64.7	31.7
	HMR	H3.6M+MPI	72.9	36.5
	SPIN	H3.6M+MPI+LSP	76.4	37.1
Weak	Zhou et al.	H3.6M	69.2	32.5
	Li et al.	H3.6M	74.1	41.4
Unsup	Chen et al.	H3.6M	64.3	31.6
	ITES-T	H3.6M	65.8	32.1
	ITES-TS	H3.6M	68.2	35.2

Table 2: Results on MPI-INF-3DHP. The input is the 2D ground truth keypoints following (Chen et al. 2019).

To test the model’s generalization ability, we further train our model on Human3.6M and test it on MPI-INF-3DHP. Table 2 presents the MPJPEs and P-MPJPEs on MPI-INF-3DHP. We see that the proposed ITES method outperforms the state-of-the-art method, Chen et al. (Chen et al. 2019), and is close to some fully-supervised method. This reflects that ITES generalizes well to out-of-domain datasets.

Qualitative results. Figure 5 visualizes the estimation results of C3DPO, one of the state-of-the-art methods, the estimation results of our teacher network (ITES-T) and the final estimation results of our framework (ITES-TS). We see that i) our method produces more precise estimations than C3DPO; and ii) the 3D poses estimated by the teacher-student framework are better than those by the single teacher network.

4.4 Ablation Studies

Effect of the training losses. First we discuss the effect of training losses we used. Table 3 presents the results. We see that i) The lack of 3D rotation-invariant loss in the teacher

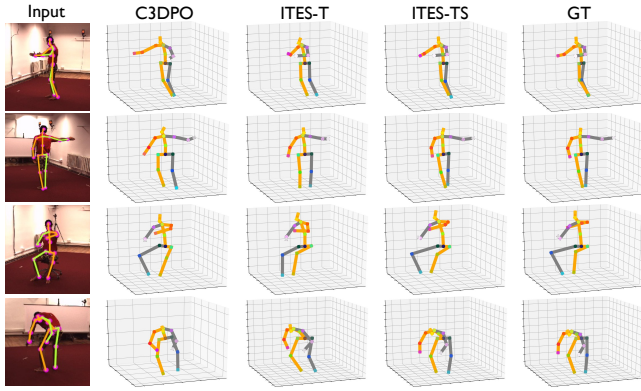


Figure 5: Qualitative results on Human3.6M dataset. We present estimations of state-of-the-art method (C3DPO), our teacher network (ITES-T), our entire framework (ITES-TS).

Ablation	\mathcal{L}_{REP}	\mathcal{L}_{RIC}	\mathcal{L}_{KD}	\mathcal{L}_{REC}	MPJPE	P-MPJPE
ITES-T	✓				153.5	120.4
ITES-T	✓	✓			85.6	54.9
ITES-TS				✓	192.3	159.7
ITES-TS	✓	✓	✓		82.4	54.3
ITES-TS	✓	✓	✓	✓	77.2	51.9

Table 3: Effect of various losses component on our framework training on Human3.6M. The input is the 2D ground truth keypoints. Note once we use \mathcal{L}_{KD} , the \mathcal{L}_{RP} and \mathcal{L}_{REC} should also be used to train the teacher network.

network leads to a large estimation error because of the decomposition ambiguity. ii) The student network cannot be trained alone with single cycle-consistent loss for 3D rotation-equivariance because of the perspective projection ambiguity. iii) With the knowledge distillation from the teacher network, adding the 3D rotation-equivariant loss will further improve the performance of the student network.

Effect of the pose dictionary sizes. The second experiment aims to discuss the effect of different pose dictionary sizes in the teacher network. Figure 6 presents the result. We see that either a relatively small or large pose-dictionary size will degrade the performance. A small pose-dictionary size restricts highly on the degree of freedom of the solution space thus cause an inaccurate estimation. A large pose-dictionary size will cause a loose regularization in the teacher network thus prevent the network to learn a set of linearly independent pose atoms. We finally set the pose dictionary size to 12.

Effect of various teachers and students. The third experiment aims to evaluate the effectiveness of the proposed pose-dictionary-based teacher network and graph-convolution-based student network. We consider two substitutions for the teacher network and the student network, respectively. We consider C3DPO (Novotny et al. 2019) as the teacher network and a fully-connected network (ResFC) adopted in (Novotny et al. 2019; Chen et al. 2019), as the student network. Table 4 presents the result. Pose-dictionary and AGCN represent our teacher and student network respectively. We see that i) for

Teacher	Student	MPJPE	P-MPJPE
C3DPO	ResFC	93.2	60.0
C3DPO	AGCN	91.7	59.6
Pose-dictionary	ResFC	82.3	53.6
Pose-dictionary	AGCN	77.2	51.9

Table 4: Effect of the proposed pose-dictionary-based teacher network and graph-convolution-based student network on Human3.6M. We replace the teacher/student network with previous methods. The input is the 2D ground truth keypoints.

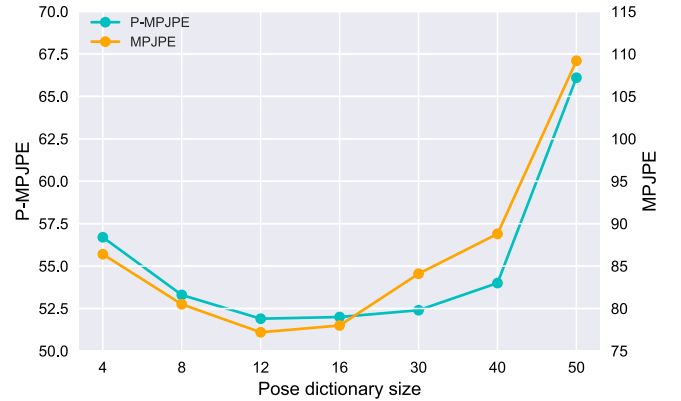


Figure 6: The performance of ITES using different pose dictionary sizes in the teacher network on Human3.6M.

the teacher network, the proposed pose-dictionary-based network works significantly better than C3DPO; and ii) for the student network, the proposed adaptive graph-convolution-based network performs better than a commonly used fully connected residual network.

5 Conclusion

Obtaining 3D human pose data acquires physical sensors, which is time-consuming and inconvenient. To alleviate the data bottleneck, we present a novel unsupervised 3D pose estimation approach. We apply a teacher-student framework in which the teacher emphasizes the regularization and the student emphasizes the flexibility. The teacher aims to make the ill-posed estimation problem feasible and the student aims to further improve the performance. Furthermore, we propose two properties for the two networks training by two cycle-consistent architectures: a 3D rotation-invariant property for the teacher for regularization and 3D rotation-equivariant property for the student for further improvement. We showed that this framework achieves the state-of-the-art performance by evaluating on Human3.6M and MPI-INF-3DHP datasets.

Acknowledgements

This work is supported by the National Key Research and Development Program of China (No. 2019YFB1804304), STCSM (No. 18DZ2270700), and State Key Laboratory of UHD Video and Audio Production and Presentation.

References

- Agudo, A.; Pijoan, M.; and Moreno-Noguer, F. 2018. Image collection pop-up: 3d reconstruction and clustering of rigid and non-rigid categories. In *CVPR*.
- Akhter, I.; Sheikh, Y.; Khan, S.; and Kanade, T. 2009. Non-rigid structure from motion in trajectory space. In *NeurIPS*.
- Akhter, I.; Sheikh, Y.; Khan, S.; and Kanade, T. 2010. Trajectory space: A dual representation for nonrigid structure from motion. *IEEE TPAMI* 33(7).
- Bottou, L. 2010. Large-scale machine learning with stochastic gradient descent. In *COMPSTAT*.
- Bregler, C.; Hertzmann, A.; and Biermann, H. 2000. Recovering non-rigid 3D shape from image streams. In *CVPR*, volume 2.
- Chen, C.-H.; and Ramanan, D. 2017. 3d human pose estimation= 2d pose estimation+ matching. In *CVPR*.
- Chen, C.-H.; Tyagi, A.; Agrawal, A.; Drover, D.; Stojanov, S.; and Rehg, J. M. 2019. Unsupervised 3d pose estimation with geometric self-supervision. In *CVPR*.
- Chen, Y.; Wang, Z.; Peng, Y.; Zhang, Z.; Yu, G.; and Sun, J. 2018. Cascaded pyramid network for multi-person pose estimation. In *CVPR*.
- Ci, H.; Wang, C.; Ma, X.; and Wang, Y. 2019. Optimizing network structure for 3d human pose estimation. In *ICCV*.
- Dai, Y.; Li, H.; and He, M. 2014. A simple prior-free method for non-rigid structure-from-motion factorization. *IJCV* 107(2).
- Drover, D.; Chen, C.-H.; Agrawal, A.; Tyagi, A.; and Phuoc Huynh, C. 2018. Can 3d pose be learned from 2d projections alone? In *ECCV*.
- Fragkiadaki, K.; Salas, M.; Arbelaez, P.; and Malik, J. 2014. Grouping-based low-rank trajectory completion and 3D reconstruction. In *NeurIPS*.
- Ionescu, C.; Papava, D.; Olaru, V.; and Sminchisescu, C. 2013. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE TPAMI* 36(7).
- Jack, D.; Maire, F.; Shirazi, S.; and Eriksson, A. 2019. IGE-Net: Inverse graphics energy networks for human pose estimation and single-view reconstruction. In *CVPR*.
- Kanazawa, A.; Black, M. J.; Jacobs, D. W.; and Malik, J. 2018. End-to-end recovery of human shape and pose. In *CVPR*.
- Kipf, T. N.; and Welling, M. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- Kocabas, M.; Karagoz, S.; and Akbas, E. 2019. Self-supervised learning of 3d human pose using multi-view geometry. In *CVPR*.
- Kolotouros, N.; Pavlakos, G.; Black, M. J.; and Daniilidis, K. 2019. Learning to reconstruct 3D human pose and shape via model-fitting in the loop. In *ICCV*.
- Kong, C.; and Lucey, S. 2016. Prior-less compressible structure from motion. In *CVPR*.
- Kudo, Y.; Ogaki, K.; Matsui, Y.; and Odagiri, Y. 2018. Unsupervised adversarial learning of 3d human pose from 2d joint locations. *arXiv preprint arXiv:1803.08244*.
- Li, M.; Chen, S.; Chen, X.; Zhang, Y.; Wang, Y.; and Tian, Q. 2019. Actional-Structural Graph Convolutional Networks for Skeleton-based Action Recognition. In *CVPR*.
- Li, M.; Chen, S.; Zhao, Y.; Zhang, Y.; Wang, Y.; and Tian, Q. 2020a. Dynamic Multiscale Graph Neural Networks for 3D Skeleton Based Human Motion Prediction. In *CVPR*.
- Li, Y.; Li, K.; Jiang, S.; Zhang, Z.; Huang, C.; and Da Xu, R. Y. 2020b. Geometry-Driven Self-Supervised Method for 3D Human Pose Estimation. In *AAAI*.
- Martinez, J.; Hossain, R.; Romero, J.; and Little, J. J. 2017. A simple yet effective baseline for 3d human pose estimation. In *ICCV*.
- Mehta, D.; Rhodin, H.; Casas, D.; Fua, P.; Sotnychenko, O.; Xu, W.; and Theobalt, C. 2017. Monocular 3d human pose estimation in the wild using improved cnn supervision. In *3DV*.
- Nath Kundu, J.; Seth, S.; Rakesh, M.; Venkatesh Babu, R.; Chakraborty, A.; et al. 2020. Kinematic-Structure-Preserved Representation for Unsupervised 3D Human Pose Estimation. *arXiv e-prints arXiv-2006*.
- Novotny, D.; Ravi, N.; Graham, B.; Neverova, N.; and Vedaldi, A. 2019. C3dpo: Canonical 3d pose networks for non-rigid structure from motion. In *ICCV*.
- Pavlakos, G.; Zhou, X.; and Daniilidis, K. 2018. Ordinal depth supervision for 3d human pose estimation. In *CVPR*.
- Pavlo, D.; Feichtenhofer, C.; Grangier, D.; and Auli, M. 2019. 3d human pose estimation in video with temporal convolutions and semi-supervised training. In *CVPR*.
- Ramakrishna, V.; Kanade, T.; and Sheikh, Y. 2012. Reconstructing 3d human pose from 2d image landmarks. In *ECCV*.
- Rhodin, H.; Spörri, J.; Katircioglu, I.; Constantin, V.; Meyer, F.; Müller, E.; Salzmann, M.; and Fua, P. 2018. Learning monocular 3d human pose estimation from multi-view images. In *CVPR*.
- Tung, H.-Y. F.; Harley, A. W.; Seto, W.; and Fragkiadaki, K. 2017. Adversarial inverse graphics networks: Learning 2d-to-3d lifting and image-to-image translation from unpaired supervision. In *ICCV*.
- Wandt, B.; and Rosenhahn, B. 2019. Repnet: Weakly supervised training of an adversarial reprojection network for 3d human pose estimation. In *CVPR*.
- Wang, C.; Kong, C.; and Lucey, S. 2019. Distill knowledge from nrsfm for weakly supervised 3d pose learning. In *ICCV*.
- Wu, J.; Xue, T.; Lim, J. J.; Tian, Y.; Tenenbaum, J. B.; Torralba, A.; and Freeman, W. T. 2016. Single image 3d interpreter network. In *ECCV*.
- Zhao, L.; Peng, X.; Tian, Y.; Kapadia, M.; and Metaxas, D. N. 2019. Semantic graph convolutional networks for 3D human pose regression. In *CVPR*.

Zhou, G.; Fan, Y.; Cui, R.; Bian, W.; Zhu, X.; and Gai, K. 2017a. Rocket launching: A universal and efficient framework for training well-performing light net. *arXiv preprint arXiv:1708.04106* .

Zhou, X.; Huang, Q.; Sun, X.; Xue, X.; and Wei, Y. 2017b. Towards 3d human pose estimation in the wild: a weakly-supervised approach. In *ICCV*.

Zhou, X.; Zhu, M.; Leonardos, S.; Derpanis, K. G.; and Daniilidis, K. 2016. Sparseness meets deepness: 3d human pose estimation from monocular video. In *CVPR*.

Zhu, Y.; Huang, D.; De La Torre, F.; and Lucey, S. 2014. Complex non-rigid motion 3d reconstruction by union of subspaces. In *CVPR*.