

Camera Distortion-aware 3D Human Pose Estimation in Video with Optimization-based Meta-Learning

Hanbyel Cho

Yooshin Cho

Jaemyung Yu

Junmo Kim

School of Electrical Engineering, KAIST, South Korea

{t1r14658, choys95, jaemyung, junmo.kim}@kaist.ac.kr

Abstract

Existing 3D human pose estimation algorithms trained on distortion-free datasets suffer performance drop when applied to new scenarios with a specific camera distortion. In this paper, we propose a simple yet effective model for 3D human pose estimation in video that can quickly adapt to any distortion environment by utilizing MAML, a representative optimization-based meta-learning algorithm. We consider a sequence of 2D keypoints in a particular distortion as a single task of MAML. However, due to the absence of a large-scale dataset in a distorted environment, we propose an efficient method to generate synthetic distorted data from undistorted 2D keypoints. For the evaluation, we assume two practical testing situations depending on whether a motion capture sensor is available or not. In particular, we propose Inference Stage Optimization using bone-length symmetry and consistency. Extensive evaluation shows that our proposed method successfully adapts to various degrees of distortion in the testing phase and outperforms the existing state-of-the-art approaches. The proposed method is useful in practice because it does not require camera calibration and additional computations in a testing set-up. Code is available at <https://github.com/hanbyel10105/CamDistHumanPose3D>.

1. Introduction

3D human pose estimation is a task that localizes 3D human body joint from an RGB input. As a fundamental task in computer vision, it is applied to many downstream applications, e.g., action recognition [35, 15, 3], human body reconstruction [33, 9], and human-computer interaction [6]. Particularly, 3D pose estimation for monocular video, which predicts 3D joint in inputs from a single-camera, has attracted a lot of academic interest recently [17, 36, 21, 2, 4, 16] because of the simplicity of the hardware setting in use and its advantage of being able to leverage temporal information to resolve inherent depth ambiguity.

Recently, many state-of-the-art studies adopted two-

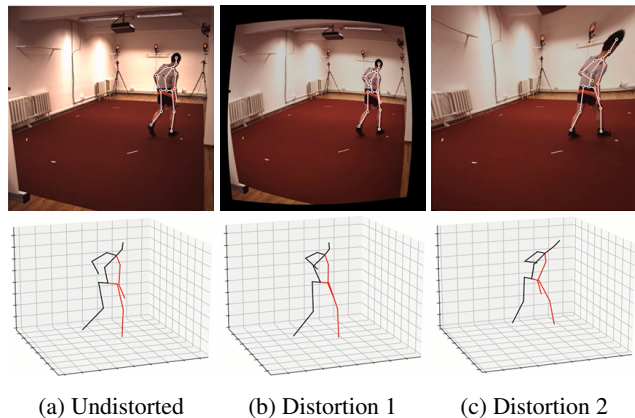


Figure 1: 3D reconstruction for videos with varying degrees of distortion using a network trained with a distortion-free dataset. **Top:** input video frames with 2D pose overlay. **Bottom:** 3D reconstruction. 3D reconstruction of (a) is predicted from undistorted video, and (b) and (c) are predicted from video with different degrees of distortion, respectively.

Condition	MPJPE(↓)	P-MPJPE(↓)	PCKh@0.5(↑)
Undistorted	48.5	37.1	87.1
Distortion 1	94.4(+45.9)	65.6(+28.5)	57.7(-29.4)
Distortion 2	133.8(+85.3)	79.2(+42.1)	38.2(-48.9)

Table 1: Performance drop in environments with distortion of a network trained with a distortion-free dataset. *Distortion 1* and *Distortion 2* are the cases of barrel distortion and pin-cushion distortion, with tangential distortion, respectively.

stage architecture to achieve higher performance. In this architecture, 2D keypoints are first extracted from the off-the-shelf 2D keypoint detector [5, 8, 20, 28], and 3D keypoints are inferred using the predicted 2D keypoints sequence as input. These approaches simplify the 3D pose estimation problem to solve depth ambiguity from 2D joint sequences. This allows the study [32, 4, 36] of algorithms explicitly using information such as skeleton kinematics and motion of human, which showed plausible results.

Despite significant advances in 2D-keypoint-based 3D pose estimation, there still remain certain limitations. That is existing 3D human pose estimation algorithms trained on distortion-free datasets show severe performance drop when applied to new scenarios with a specific camera distortion, as shown in Figure 1 and Table 1. Previously, preprocessed images were used when inferring 3D joints from distorted inputs. However, it is important to make models that can adapt themselves to arbitrary distortion in the testing phase, as algorithms that are needed in preprocessing, such as camera calibration, are sometimes difficult to apply and they also introduce certain errors of their own. This is substantially important issue in the wild use of the algorithms; however, cross-scenario research on camera distortion has been out of scope due to the absence of a dataset with various degrees of camera distortion.

To overcome this limitation, in this work, we propose a simple yet effective model for 3D human pose estimation in video that can quickly adapt to any distortion environment by utilizing model-agnostic meta-learning (MAML) [7], a representative optimization-based meta-learning algorithm. We focus on training a distorted-2D-keypoints-conditioned 3D pose estimator to be able to quickly adapt to camera distortion, because we found that 2D keypoint detector is good at finding distorted 2D keypoints consistent with distorted images. Therefore, we consider a sequence of 2D keypoints in a particular distortion as a single task of MAML. However, due to the absence of a large-scale dataset with a distorted environment, we propose an efficient method to generate synthetic distorted data from undistorted 2D keypoints. Note that, the goal of training phase is not to just increase the performance at a particular distortion, but to train a network sensitive to distortion, allowing the network to adapt quickly to arbitrary distortion in the testing phase. For the testing phase, the trained network is first adapted to a specific camera distortion environment by fine-tuning or Inference Stage Optimization, which as proposed in recent work [34] in the following two scenarios.

For evaluation, we assume two practical situations in which the proposed method will be used and confirm that our algorithm is useful for each case. First, *Scenario 1* is a situation in which a user can collect data using motion capture sensors in front of a testing environment, as shown in Figure 3 (b). In this case, data with the same distortion as the testing environment can be obtained, but it would be in relatively small amounts compared with a large-scale dataset (e.g., Human3.6M [10]) collected in the laboratory environment. Therefore, it is important to transfer knowledge trained with a large-scale dataset as much as possible. To validate the usefulness of the proposed method, we construct a small-scale dataset with the same distortion as the test environment, and evaluate whether the network can adapt well through naive fine-tuning.

Second, *Scenario 2* is when the user is unable to obtain data in a testing environment, as shown in Figure 3 (c). In this case, the network should be adapted to specific distortions using only test videos. In a recent study [34], the authors proposed the concept named Inference Stage Optimization (ISO) to adapt network using only test data before testing. We also utilize ISO in this case. To this end, we propose a novel ISO method based on skeleton symmetry and consistency. This might be a weak constraint, but we confirm that our network is fully adaptable even with these constraints because it has been sensitively trained on distortion.

In summary, our overall contribution is four-fold:

- To the best of our knowledge, our method, which utilized optimization-based meta-learning, is the first algorithm that can adapt to arbitrary camera distortion at the testing phase.
- We propose an efficient method to generate synthetic distorted data from undistorted 2D keypoints, enabling cross-scenario research on camera distortion, which has been out-of-scope due to the absence of datasets with distortion.
- We validate the effectiveness of the proposed method for each case, assuming two practical testing environments. In particular, we propose the ISO method using bone-length symmetry and consistency.
- Our proposed method is useful in practical applications because it does not require calibration for the testing camera and additional computational complexity.

2. Related Work

2.1. 3D Human Pose Estimation

Since the success of 2D human pose estimation, 3D human pose estimation has been widely studied. Martinez *et al.* [17] successfully predicted 3D poses from 2D joint locations using simple and lightweight networks. It showed better results than previous studies that involved training with raw image pixels. To make better use of 2D keypoints, GCN and attention mechanism were applied to learn the global relationship between joints [36, 16]. In contrast, Pavlo *et al.* [21] predicted 3D poses using video to overcome inherent ambiguity that multiple 3D poses can be mapped to the same 2D pose. Furthermore, prior knowledge about human body structure was explicitly utilized to give constraints [2, 4]. Despite substantial progress in this field, the performance severely drops when camera distortion occurs due to the changes in camera parameters in test environment.

2.2. Cross-scenario Pose Estimation

Deep learning models have substantially improved in recent years. However, due to the limitation of supervised learning on datasets that lack diversity, even state-of-the-art

algorithms show poor results in-the-wild. To be robust on a domain gap between training and inference, many studies have been conducted [25, 27, 34]. A recent study [34] proposed the domain (*e.g.*, varying poses, camera viewpoints, body size, and appearances) robust 3D pose estimation algorithm that adapts to the target domain using self-supervised learning schemes named Inference Stage Optimization (ISO) using cycle consistency among 2D and 3D spaces. In this paper, we focus on the domain gap of camera distortion caused by the different camera settings at the testing phase, which has been out-of-scope.

2.3. Optimization-based Meta-Learning

There are three common categories in meta-learning. The first category is the metric-based approach [12, 30, 26, 29], which learns a good metric that expresses the relationship between inputs in task space and applies it well to new samples. The second category is a model-based approach [24, 19, 18] that controls the structure of a target model through another model called meta-learner. The last category is an optimization-based approach [1, 23, 7, 22] that looks for sensitive initial parameters for tasks and quickly adapts to new tasks with only a few samples. In this work, we utilize MAML [7], which belongs to the optimization-based approach so that the network can adapt quickly to arbitrary camera distortion in the testing phase.

3. Preliminary

In this section, we introduce background knowledge on camera distortion and framework of the MAML algorithm.

Camera Distortion. There are two kinds of camera distortion. The first is radial distortion, which is caused by the refractive index of the convex lens. Radial distortion is determined by the distance from the center of the image, and it is usually expressed in parameters k_1 , k_2 , and k_3 . The value of k_1 determines the main form of the distortion. A negative k_1 and a positive k_1 result in barrel distortion and pincushion distortion respectively, as shown in Figure 2 (b) and (c). The second is tangential distortion, which is caused by the misalignment of the camera lens and the image sensor (*e.g.*, CCD and CMOS) during manufacturing of the camera. This can be approximated by parameters p_1 and p_2 . The p_1 and p_2 cause ladder shape distortion mainly in x -axis and y -axis respectively, as shown in Figure 2 (d) and (e). Both distortions are common in commercial cameras, and radial distortion is particularly severe in wide-angle cameras.

Model-Agnostic Meta-Learning. The stage of meta-learning consists of meta-training and meta-testing. We consider a model represented by a function g_θ with parameters θ , that outputs \mathbf{y} with input as \mathbf{x} . The objective of meta-training is to find initial transferable weights that can be adapt to new tasks. For meta-training, a batch of

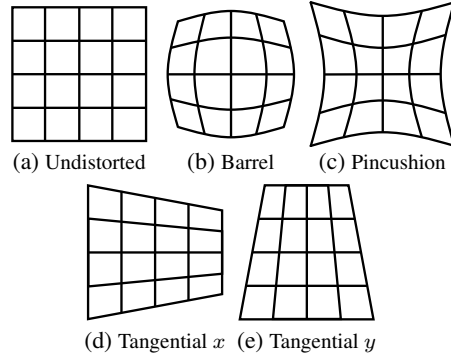


Figure 2: Types of camera distortion. (b) and (c) represent radial distortion, and (d) and (e) represent tangential distortion. Radial distortion and tangential distortion can occur simultaneously.

tasks \mathcal{T}_i is sampled from task distribution $p(\mathcal{T})$. The model is first optimized through the task-specific loss $\mathcal{L}_{\mathcal{T}_i}$ using training samples within a task (*task-level training*), and meta-optimization is performed using test samples (*task-level testing*). In meta-testing, the model adapts to a new task \mathcal{T}_{new} using only a few samples. In this study, we use MAML [7], in which input \mathbf{x} and output \mathbf{y} are distorted 2D keypoint trajectory and 3D joints respectively. Various distortion parameters construct task distribution, and each task corresponds to a 3D pose estimation from the distorted 2D keypoint trajectory with a particular distortion parameter.

4. Method

The overall framework of our method is shown in Figure 3. In this section, we first propose a method for *generating synthetic distorted tasks*. Then, we describe two phases that constitute our method: *training phase* and *adaptation before testing*.

4.1. Synthetic Distorted Task Generation

We found that 2D keypoint detector is good at finding distorted 2D keypoints consistent with distorted images, as shown in the top row of Figure 1 because it is based on the texture of images. Thus, in the training phase, our goal is to train a 3D pose estimator conditioned on distorted 2D detection to be able to quickly adapt to various distortions by applying meta-learning. Meta-learning in our case requires tasks under varying degrees of distortion. In this section, we describe how to efficiently generate distorted tasks from undistorted videos.

Given a video clip with frame length of T , first 2D keypoints are obtained by a pretrained 2D keypoint detector (*e.g.*, Mask R-CNN [8]). Let $\tilde{\mathbf{p}}_t \in \mathbb{R}^{J \times 2}$ denotes predicted 2D coordinates of J keypoints of the human in the frame and $\tilde{\mathbf{P}} = \{\tilde{\mathbf{p}}_t\}_{t=1}^T$ denotes the set of joints for a video clip. Specifically, $\tilde{\mathbf{p}}_t = \{[\tilde{a}_{t,j}, \tilde{b}_{t,j}]\}_{j=1}^J$ where $\tilde{a}_{t,j}$ and $\tilde{b}_{t,j}$ denote x and y coordinates of j th joint at frame t , respectively.

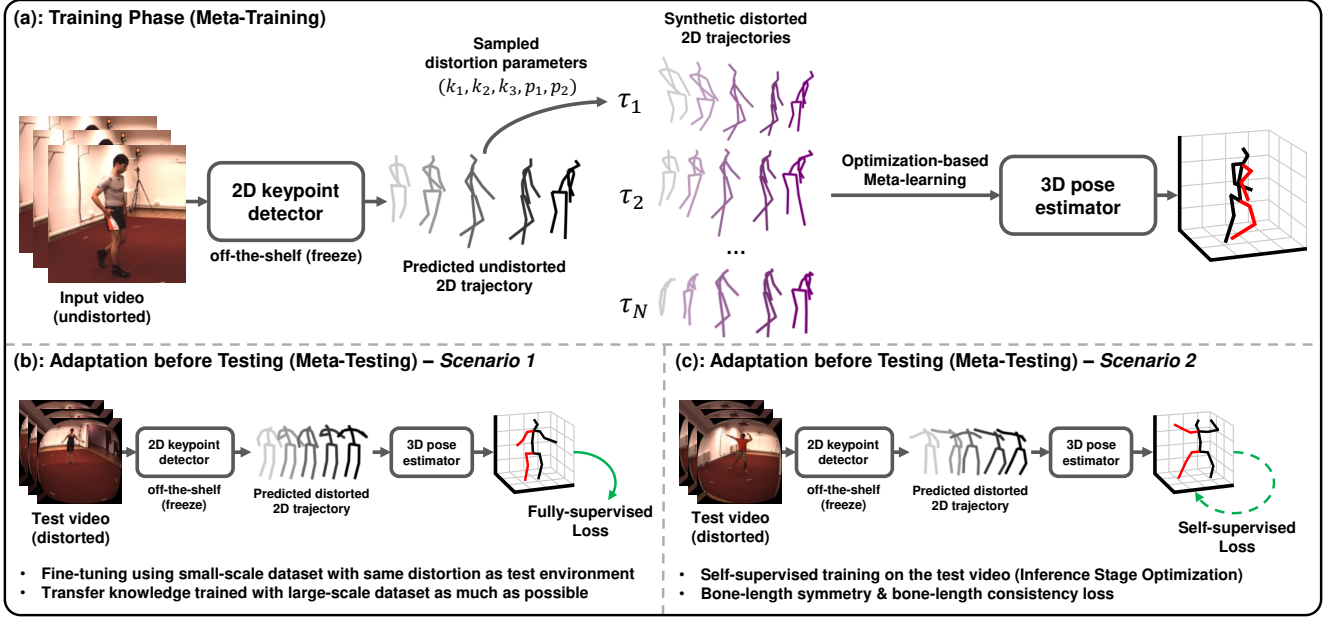


Figure 3: Overall framework of our methods. (a) We train a 2D-keypoint-conditioned 3D pose estimator that can quickly adapt to any distortions using only an undistorted large-scale dataset. Before the trained network can be used in practice, it must be adapted to a certain distortion. (b) and (c) represent adaptation method for *Scenario 1* and *Scenario 2*, respectively.

To generate synthetic distorted tasks, we apply the camera distortion model [31] directly to predicted 2D keypoints. We omit subscript t and j for simplicity. As shown in Figure 4 (a), the process of generating the task with particular distortion parameters (i.e., k_1, k_2, k_3, p_1, p_2) is divided into three steps. The first is obtaining normalized 2D keypoints (denoted as $[\tilde{a}_n, \tilde{b}_n]$) and distance between the point and image center (denoted as r). As camera distortion models should be applied on a normalized image plane, we first normalize 2D keypoints with focal length (denoted as $\mathbf{f} = [f_x, f_y]$) and optical center (denoted as $\mathbf{c} = [c_x, c_y]$) using the following equations:

$$\tilde{a}_n = \frac{\tilde{a} - c_x}{f_x}, \tilde{b}_n = \frac{\tilde{b} - c_y}{f_y}, r = \sqrt{\tilde{a}_n^2 + \tilde{b}_n^2}. \quad (1)$$

Then, we apply distortion to the normalized 2D keypoints using the following equations:

$$\tilde{a}_{n,d} = \tilde{a}_n(d_r + d_t) + p_1 r^2, \tilde{b}_{n,d} = \tilde{b}_n(d_r + d_t) + p_2 r^2, \quad (2)$$

where intermediate variable d_r and d_t are obtained by $d_r = 1 + k_1 r^2 + k_2 r^4 + k_3 r^6$ and $d_t = 2p_1 \tilde{a}_n + 2p_2 \tilde{b}_n$ respectively. Finally, distorted 2D keypoints (denoted as $[\tilde{a}_d, \tilde{b}_d]$) are obtained by unnormalization using following equations:

$$\tilde{a}_d = \tilde{a}_{n,d} f_x + c_x, \tilde{b}_d = \tilde{b}_{n,d} f_y + c_y. \quad (3)$$

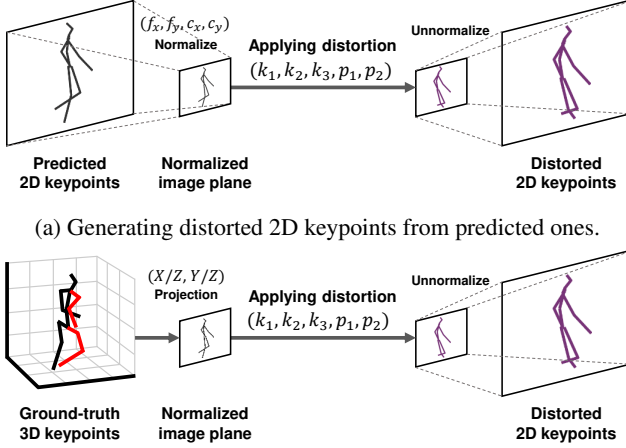
We apply this process to all joints J and frames T to obtain a distorted 2D trajectory (denoted as $\tilde{\mathbf{P}}_{dist}$) reflecting a specific distortion. Then, we consider a pair of the distorted 2D trajectory and ground-truth 3D joints (denoted as $\mathbf{s} = [\mathbf{x}, \mathbf{y}, \mathbf{z}] \in \mathbb{R}^{J \times 3}$) as a single task \mathcal{T} of MAML.

This method is highly efficient because it does not apply distortion in image domain, and thus, we can generate numerous distortions in the training phase, as shown in Figure 3 (a). Furthermore, it can reflect the jittered outputs of the 2D keypoint detector caused by inherent ambiguity (e.g., occlusion), because it generates distorted joints from predicted 2D keypoints. Synthetic tasks can also be generated from ground-truth 3D joints, as shown in Figure 4 (b). In this case, normalized 2D keypoints are obtained from the ground-truth 3D joints through projection. However, as shown in Table 4, this method is less effective because it could not reflect the noisy output of the 2D keypoint detector, resulting in a domain gap during training and testing.


4.2. Training Phase

In the training phase, we will perform meta-learning using synthetic distorted tasks from undistorted videos. Our goal in the training phase is to find sensitive initial transferable weights to camera distortion by utilizing optimization-based meta-learning. Our algorithm mostly follows the framework of MAML, but to achieve better performance, there are two modifications: *stratified sampling* and *random distortion pretraining*.

As shown in Figure 3 (a), given predicted undistorted 2D trajectory, we generate a batch of distorted 2D trajectories (denoted as $\{\tilde{\mathbf{P}}_{dist,i}\}_{i=1}^N$, where N represents the number of tasks in meta-batch) with sampled distortion parameters. Then, we construct each task \mathcal{T}_i by pairing a distorted 2D trajectory $\tilde{\mathbf{P}}_{dist,i}$ and ground-truth 3D joints \mathbf{s} .



(a) Generating distorted 2D keypoints from predicted ones.



(b) Generating distorted 2D keypoints from 3D ground-truth.

Figure 4: Methods to generate distorted 2D keypoints.

Specifically, parameters related to radial distortion are sampled by $k_1, k_2, k_3 \sim \mathcal{U}[-\lambda_1, \lambda_1]$ and tangential distortion parameters are sampled by $p_1, p_2 \sim \mathcal{U}[-\lambda_2, \lambda_2]$ where λ_1 and λ_2 denote the maximum value of each distribution. We basically use the sampling method to both task-level training and task-level testing. However, for the task-level training, we adopt *stratified sampling* for sampling parameter k_1 , which determines the main form of distortion. In this case, a k_1 of i th sample in the meta-batch is sampled as follows:

$$k_{1,i} \sim -\lambda_1 + 2 \cdot \lambda_1 \cdot \mathcal{U} \left[\frac{i-1}{N}, \frac{i}{N} \right]. \quad (4)$$

By sampling the distortion parameter k_1 from evenly spaced bins, the meta-batch can consist of tasks with varying degrees of distortion. This enhances the adaptability of our network as shown in Table 3. We denote the distribution of tasks generated using stratified sampling as $p_{strat}(\mathcal{T})$, and using only uniform distribution as $p_{rand}(\mathcal{T})$.

Finally, we consider a 3D pose estimator model represented by a parameterized function g_θ with parameters θ . We perform only one gradient descent update when the parameters θ is adapted to a new task \mathcal{T}_i . Thus, the newly adapted parameters θ'_i are obtained by

$$\theta'_i = \theta - \alpha \nabla_\theta \mathcal{L}_{\mathcal{T}_i}(g_\theta), \quad (5)$$

where α is the learning rate for task-level training.

The parameters θ of model are optimized by maximizing the performance of $g_{\theta'_i}$ with respect to θ across tasks sampled for task-level testing. Specifically, the meta-objective is expressed as follows:

$$\begin{aligned} & \arg \min_{\theta} \sum_{\mathcal{T}_i \sim p(\mathcal{T})} \mathcal{L}_{\mathcal{T}_i}(g_{\theta'_i}) \\ & = \arg \min_{\theta} \sum_{\mathcal{T}_i \sim p(\mathcal{T})} \mathcal{L}_{\mathcal{T}_i}(g_{\theta - \alpha \nabla_\theta \mathcal{L}_{\mathcal{T}_i}(g_\theta)}). \end{aligned} \quad (6)$$

Finally, we perform meta-optimization by using the Eq. 6. For the stochastic gradient descent, model parameters θ are updated as follows:

$$\theta \leftarrow \theta - \beta \nabla_\theta \sum_{\mathcal{T}_i \sim p(\mathcal{T})} \mathcal{L}_{\mathcal{T}_i}(g_{\theta'_i}), \quad (7)$$

where β is the learning rate for meta-optimization. We use a loss function, MPJPE, which is the L2 distance between ground-truth 3D joints and predicted ones as a task-level objective in the entire process of meta-optimization.

Additionally, we pretrain the network before training meta-learner through *random distortion pretraining* that regresses 3D joints from randomly distorted 2D keypoint trajectories. This allows the network to learn feature representation under various distortions and consequently enables stable MAML training. However, while random distortion pretraining helps in the stability of MAML, the pretraining *without* meta-learning shows poor results, as shown in Figure 7, when the network is adapted to the specific distortion, because it is not a transferable initial weights.

4.3. Adaptation before Testing

When the trained model that can quickly adapt to arbitrary distortions is used, it must, first, be adapted to the specific distortion of the testing environment. This is similar to meta-testing in the MAML framework. We assume two practical situations, *Scenario 1* and *Scenario 2*, and propose an adaptation method for each case.

Scenario 1 is a situation in which a user can collect data using motion capture sensors in front of a testing environment. In this case, data with the same distortion as the testing environment can be obtained. Thus, we adopt naive fine-tuning using the MPJPE loss function to adapt the network to the specific distortion, as shown in Figure 3 (b). The collected data would be in relatively small amounts than the large-scale dataset (e.g., Human3.6M). Therefore, we will check whether the network can adapt well with small amounts of collected data. Detailed settings are provided in Section 5.1.

Scenario 2 is a situation when the user is unable to obtain data in a testing environment. In this case, the network should be adapted to the specific distortion using only test videos. As shown in Figure 3 (c), we adopt Inference Stage Optimization (ISO) [34], which performs self-supervised training using the test data before testing. Usually, the inferred 3D joints are orthogonally projected to 2D plane and compared with the predicted 2D keypoints to perform ISO. However, if there is distortion in the video, this method cannot be used. Therefore, we propose the novel ISO method which utilizes *bone-length symmetry* and *bone-length consistency* that allows self-supervision within the inferred 3D joint itself (details in Appendix A.2). The former constrains the length of a person's left and right bones to be equal,

whereas the latter constrains each bone to be equal in length between consecutive frames within a video. The constraints based on bone-length have been used for regularization in fully-supervised training, but have never been used for ISO. Also, these methods might be a weak constraint, but our network is fully adaptable even with these constraints because it has been sensitively trained on distortion via MAML.

4.4. Algorithm

Algorithm 1 shows the entire process of Section 4.2. As shown in lines 2-8, random distortion pretraining is performed before meta-learning. Subsequently, meta-learning is performed, as shown in lines 9-17. Lines 11-14 and lines 15-16 present task-level training and meta-optimization with task-level testing, respectively.

5. Experiments

5.1. Datasets and Evaluation

Human3.6M [10] is a large-scale dataset containing 3.6 million video frames and corresponding 2D and 3D human keypoint labels. We construct a cross-scenario on distortion to validate the effectiveness of the proposed method. For training, we use five subjects (S1, S5, S6, S7, S8) with undistorted videos, as in previous works [21, 4], since our method can generate synthetic distorted tasks from undistorted 2D keypoints. For testing, only one subject (S11) is used. Due to the absence of test videos with distortion, we generate four different kinds of distorted videos (denoted as d_1 , d_2 , d_3 , and d_4 , details in Appendix A.1) from undistorted videos of S11 by using Blender¹ software, as shown in Figure 5. We evaluate the proposed method in each kind of distortion. For *Scenario 1* in adaptation before the testing phase, collected small-scale dataset with the same distortion as the test data is required. Therefore, we adopt only 1% of S9 and apply the same distortion as S11 to it.

Evaluation Metrics. We use three evaluation protocols following previous works [17, 36, 21, 4, 16, 34]. The first is mean per joint position error (MPJPE) in millimeters, the L2 distance between the predicted 3D joints and ground-truth joints. The second is P-MPJPE. This is similar to MPJPE, but calculates the error between the joints after alignment using Procrustes Analysis. The last one is percentage of correct 3D joints with a threshold as 50% of the head segment length (PCKh@0.5).


5.2. Implementation Details

For 3D pose estimator, the proposed method is not about network architecture but training methods. Therefore, we adopt the state-of-the-art model for 3D human pose estimation in video proposed in the previous work [21] as our base model. It is fully convolutional and based on dilated

Algorithm 1: Training Phase

Input: \mathcal{D} : a large-scale 3D human pose dataset
Input: α, β : learning rate hyperparameters
Output: Model parameters θ

- 1 Randomly initialize θ
- 2 **while not done do**
- 3 Sample batch of tasks $\mathcal{T}_{rand,i} \sim p_{rand}(\mathcal{T})$
- 4 **for all** $\mathcal{T}_{rand,i}$ **do**
- 5 Calculate loss by MPJPE: $\mathcal{L}_{\mathcal{T}_{rand,i}}(g_\theta)$
- 6 Compute updated parameters:
 $\theta = \theta - \beta \nabla_\theta \mathcal{L}_{\mathcal{T}_{rand,i}}(g_\theta)$
- 7 **end**
- 8 **end**
- 9 **while not done do**
- 10 Sample batch of tasks $\mathcal{T}_{strat,i} \sim p_{strat}(\mathcal{T})$
- 11 **for all** $\mathcal{T}_{strat,i}$ **do**
- 12 Calculate loss by MPJPE: $\mathcal{L}_{\mathcal{T}_{strat,i}}(g_\theta)$
- 13 Compute updated parameters:
 $\theta'_i = \theta - \alpha \nabla_\theta \mathcal{L}_{\mathcal{T}_{strat,i}}(g_\theta)$
- 14 **end**
- 15 Update θ with respect to average test loss:
- 16 $\theta \leftarrow \theta - \beta \nabla_\theta \sum_{\mathcal{T}_{rand,i} \sim p_{rand}(\mathcal{T})} \mathcal{L}_{\mathcal{T}_{rand,i}}(g_{\theta'_i})$
- 17 **end**



(a) B+T (d_1) (b) P+T (d_2) (c) B+T (d_3) (d) P+T (d_4)

Figure 5: Videos rendered with different kinds of distortion. B, P, and T represent barrel, pincushion, and tangential distortion respectively. For (a) and (b) heavy distortion is applied, and moderate distortion is applied to (c) and (d).

temporal convolutions with residual blocks. For 2D keypoint detector, we use *Mask R-CNN* [8] with a ResNet-101-FPN [13] backbone as off-the-shelf 2D keypoint detector. We fine-tune the COCO [14] pretrained model on 2D keypoints of Human3.6M. Similar to previous work [21], the 2D keypoint format of COCO differs from Human3.6M, and hence, we reinitialize the last layer of the keypoint network of the detector and carry out fine-tuning, after which the 2D keypoint detector is frozen in the entire training process for the 3D pose estimator because it is robust to camera distortion. We use Adam [11] optimizer, with batch size 1024. During the training phase, we set the required λ_1 and λ_2 to 5 and 0.5 respectively, to sample the distortion parameters. The learning rate α for task-level training is set to 0.1 and the β for meta-optimization is set to 0.001. We use 5 for the number of samples in meta-batch. The learning rate decay is set to 0.95 and network is trained with 60 epochs.

¹<https://www.blender.org/>

Method	Scenario 1			Scenario 2		
	MPJPE(↓)	P-MPJPE(↓)	PCKh@0.5(↑)	MPJPE(↓)	P-MPJPE(↓)	PCKh@0.5(↑)
Martinez <i>et al.</i> [17] ICCV'17	78.3 / 63.1	58.1 / 48.7	66.6 / 76.5	128.0 / 68.3	86.8 / 49.1	47.3 / 74.1
Zhao <i>et al.</i> [36] CVPR'19	86.3 / 64.0	64.2 / 47.4	63.2 / 76.9	119.7 / 71.4	85.5 / 51.9	45.0 / 72.2
Pavlo <i>et al.</i> [21] CVPR'19	79.9 / 65.0	59.4 / 48.3	67.3 / 76.7	114.1 / 64.5	72.4 / 45.7	47.9 / 76.6
Chen <i>et al.</i> [4] TCSVT'21	89.4 / 62.7	61.9 / 46.3	59.2 / 77.8	107.3 / 65.1	71.0 / 46.3	49.0 / 77.3
Liu <i>et al.</i> [16] CVPR'20	81.5 / 68.8	60.9 / 51.0	66.4 / 74.7	110.7 / 64.0	77.5 / 46.5	49.5 / 76.8
Ours	62.0 / 53.6	46.4 / 40.6	78.4 / 83.3	66.1 / 51.6	47.8 / 39.2	76.3 / 85.7

Table 2: Comparison of average performance on (heavy) / (moderate) with other state-of-the-art models. The top two rows [17, 36] are based on a single-frame and others [21, 4, 16], including our method, are based on a video with a frame length of 27. Best in bold, second-best underlined. More results can be seen in the supplementary material (Appendix A.3).

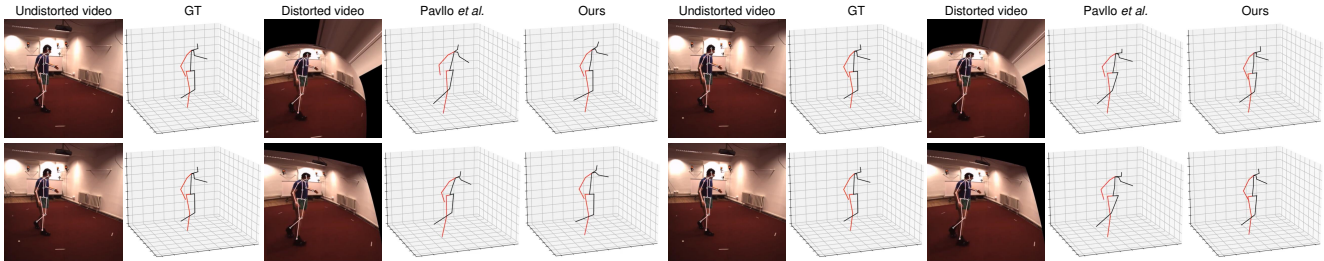


Figure 6: Qualitative results on heavily distorted videos of Human3.6M. The five columns from the leftmost are the result under the Scenario 1 setting, while the rest columns are the result under the Scenario 2 setting. **Top row:** 3D reconstruction results on d_1 . **Bottom row:** 3D reconstruction results on d_2 . More results can be seen in Appendix A.4.

In both Scenario 1 and Scenario 2 of adaptation before testing, learning rate is set to 0.6 and epochs for adaptation is set to 100. Note that, during the adaptation process, we train the model for 100 epochs, however since it is done on a small-scale dataset, the overall training time required for the adaptation is within a few minutes.

5.3. Experiment Results

In this section, we validate effectiveness of the proposed method. We evaluate performance on four different kinds of distortion in all experiments and report the average performance of d_1 and d_2 which applied heavy distortion and the average performance of d_3 and d_4 which applied moderate distortion. All reported values are performance after adaptation to the specific distortion. For Scenario 1, the trained network is adapted by fine-tuning on 1% of S9, which went through the same distortion as S11. For Scenario 2, the network is adapted by ISO on 0.1% of test videos (S11).

Comparison with State-of-the-Art. Table 2 shows the performance of existing 3D pose estimation algorithms and our method. The baseline models do not take cross-scenario about distortion into account. However, for fair evaluation, we evaluate the performance in Scenario 1 after fine-tuning on the small-scale dataset, just like our method. For Scenario 2, we did not apply ISO to baseline models because they show poor performance when ISO is applied.

The proposed method outperforms other methods re-

gardless of the kinds of distortions and scenarios. Specifically, compared to our base model [21], the proposed method shows -14.64mm, -10.35mm, and +17.7% average performance improvement in Scenario 1 for each metric (*i.e.*, MPJPE, P-MPJPE, and PCKh@0.5) and -30.45mm, -15.55mm, and +18.75% average performance improvement in Scenario 2. Especially in Scenario 2, our method rather shows better performance than Scenario 1 under moderate distortion. This demonstrates that our bone-length-based ISO method is effective and also that the trained model has transferable initial weights.

Figure 6 shows qualitative results of the estimated 3D pose from distorted videos. Unlike others, our method successfully adapts to the distortion that the test video has, and consequently, we can see that estimated 3D joints from the distorted video by using the proposed method is almost the same with ground-truth joints. Results predicted from videos with more diverse poses and distortion can be seen in the supplementary material (Appendix A.4).

Ablation Studies. We first look at the contribution of each of the proposed method. We evaluate the performance changes, adding each proposed method with Pavlo *et al.* [21] as our base model. As shown in Table 3, we can notice that each method provides a positive contribution under all metrics. In particular, the significant improvement comes from utilizing MAML using synthetic distorted tasks and learning rich feature representation on distortion through

	MPJPE(↓)	P-MPJPE(↓)	PCKh@0.5(↑)
<i>base model</i> [21]	84.2 / 79.6	62.8 / 59.7	64.8 / 66.9
+ MAML (with synthetic tasks)	73.5 / 67.5	55.6 / 51.7	72.0 / 74.5
+ stratified sampling	71.7 / 66.2	54.3 / 50.4	72.8 / 75.2
+ random distortion pretraining	67.2 / 61.9	51.0 / 47.0	75.7 / 78.2

Table 3: Effectiveness of each proposed method based on input frame length of 9 under *Scenario 1* setting. Each value denotes performance on (distortion d_1) / (distortion d_2).

Method	MPJPE(↓)	P-MPJPE(↓)	PCKh@0.5(↑)
Predicted 2D keypoints	62.0 / 53.6	46.4 / 40.6	78.4 / 83.3
Ground-truth 3D joints	64.7 / 56.1	48.2 / 42.0	77.0 / 82.0
Predicted 2D keypoints	66.1 / 51.6	47.8 / 39.2	76.3 / 85.7
Ground-truth 3D joints	71.3 / 55.6	51.9 / 42.6	72.8 / 83.5

Table 4: Comparison of average performance on (heavy) / (moderate) between the methods generating synthetic 2D keypoints. **Top rows:** *Scenario 1*. **Bottom rows:** *Scenario 2*.

random distortion pretraining.

Table 4 shows the performance when applying each of the two methods that generate synthetic distorted 2D keypoints based on frame length of 27. Predicted 2D keypoints denotes the case in which synthetic distorted 2D keypoints are generated using noisy results inferred from the 2D keypoint detector, as shown in Figure 4 (a) and Ground-truth 3D keypoints denotes the case in which synthetic keypoints are generated using the ground-truth 3D keypoints as shown in Figure 4 (b). We can notice that the former method shows better performance under all metrics and scenarios since there is less domain gap between training and testing.

Table 5 reports the performance and complexity of the model (*i.e.*, parameters and FLOPs) with respect to different input frame lengths. Our method uses the same model structure as Pavllo *et al.* [21] because our work is about the learning method rather than the structure of the model. When the input frame length is 3, the proposed method shows comparable performance to the base model, even though the capacity is one-fifty of the base model. Furthermore, our method with an input frame length of 27 outperforms the base model of the same size significantly. Moreover, our method has the same floating-point operations (FLOPs) for inference as the base model [21], thus no additional computational cost is required compared to the base model when testing after adaptation to the test environment.

Performance Changes during Adaptation. We also validate the ability of the model, trained in the training phase, to adapt well to the specific camera distortion. In this experiment, we observe the performance changes of the model with and without MAML during the adaptation process based on frame length 27. As shown in Figure 7, in the case of a model using MAML, we can notice that it adapts well regardless of the degrees of distortion and scenarios.

Model	Parameters	\approx FLOPs	MPJPE	P-MPJPE	PCKh@0.5
Pavllo <i>et al.</i> [21] 27f	8.56M	17.11M	72.4	53.8	72.0
Ours 3f	0.16M	0.32M	75.0	56.1	69.6
Ours 9f	4.36M	8.71M	59.8	45.4	79.5
Ours 27f	8.56M	17.11M	57.6	43.4	80.9

Table 5: Performance and computational complexity of various models under *Scenario 1*. The reported performance is the average value for all kinds of distortions.

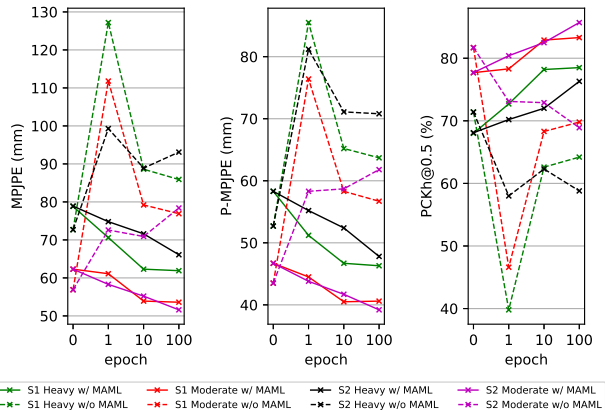


Figure 7: Performance changes during adaptation to the specific distortion. S1 and S2 denote *Scenario 1* and *Scenario 2*, respectively. A solid line w/ MAML denotes our final model trained using all the elements proposed in Section 4.2, and a dashed line w/o MAML denotes a model trained only with *random distortion pretraining*.

Also, the mean and standard deviation of MPJPE are 6.5mm (10%) and 2.3mm (25%) lower than those of w/o MAML (at epoch 0), respectively. In contrast, in the case of the model not using MAML, the model is not stably adapted, and its performance is rather significantly degraded. Specifically, it performs well at epoch 0 with the effect of the proposed random distortion pretraining, however since it is not a transferable initial weight, it is highly degraded when the adaptation process starts. This demonstrates the superior potential of MAML to adapt to various distortion environments. Note that, as mentioned in Section 5.2, training time required for the adaptation process is within a few minutes.

6. Conclusion

We have introduced a model for 3D human pose estimation that can adapt quickly to arbitrary camera distortion. Our model finds initial transferable weights that are sensitive to distortion through meta-learning. For this, we overcome the limitations of the absence of publically available distorted data by generating synthetic distorted tasks from undistorted data. Furthermore, we propose a novel ISO method based on bone-length that can adapt the model to the test environment without 3D joint labels. Our method is expected to be very useful in practice because once trained, it can adapt to any distortion without camera calibration.

References

- [1] Marcin Andrychowicz, Misha Denil, Sergio Gómez Colmenarejo, Matthew W. Hoffman, David Pfau, Tom Schaul, Brendan Shillingford, and Nando de Freitas. Learning to learn by gradient descent by gradient descent. In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS'16*, page 3988–3996, Red Hook, NY, USA, 2016. Curran Associates Inc. 3
- [2] Yujun Cai, Lihao Ge, Jun Liu, Jianfei Cai, Tat-Jen Cham, Junsong Yuan, and Nadia Magnenat Thalmann. Exploiting spatial-temporal relationships for 3d pose estimation via graph convolutional networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2272–2281, 2019. 1, 2
- [3] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. 1
- [4] Tianlang Chen, Chen Fang, Xiaohui Shen, Yiheng Zhu, Zhili Chen, and Jiebo Luo. Anatomy-aware 3d human pose estimation with bone-based pose decomposition. *IEEE Transactions on Circuits and Systems for Video Technology*, PP:1–1, 02 2021. 1, 2, 6, 7, 12
- [5] Yilun Chen, Zhicheng Wang, Yuxiang Peng, Zhiqiang Zhang, Gang Yu, and Jian Sun. Cascaded pyramid network for multi-person pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 1
- [6] Irene Connolly, Marion Palmer, Hannah Barton, and Grinne Kirwan. *An Introduction to Cyberpsychology*. Routledge, USA, 2016. 1
- [7] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1126–1135, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR. 2, 3
- [8] Kaiming He, Georgia Gkioxari, Piotr Dollar, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Oct 2017. 1, 3, 6
- [9] Zhongyue Huang, Jingwei Xu, and Bingbing Ni. Human motion generation via cross-space constrained sampling. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 757–763. International Joint Conferences on Artificial Intelligence Organization, 7 2018. 1
- [10] C. Ionescu, D. Papava, V. Olaru, and C. Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(7):1325–1339, 2014. 2, 6
- [11] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. 6
- [12] Gregory Koch, Richard Zemel, Ruslan Salakhutdinov, et al. Siamese neural networks for one-shot image recognition. In *ICML deep learning workshop*, volume 2. Lille, 2015. 3
- [13] Tsung-Yi Lin, Piotr Dollar, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. 6
- [14] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft coco: Common objects in context. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision – ECCV 2014*, pages 740–755, Cham, 2014. Springer International Publishing. 6
- [15] Jun Liu, Gang Wang, Ping Hu, Ling-Yu Duan, and Alex C. Kot. Global context-aware attention lstm networks for 3d action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. 1
- [16] Ruixu Liu, Ju Shen, He Wang, Chen Chen, Sen-ching Cheung, and Vijayan Asari. Attention mechanism exploits temporal contexts: Real-time 3d human pose reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5064–5073, 2020. 1, 2, 6, 7, 12
- [17] Julieta Martinez, Rayat Hossain, Javier Romero, and James J Little. A simple yet effective baseline for 3d human pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2640–2649, 2017. 1, 2, 6, 7, 12
- [18] Nikhil Mishra, Mostafa Rohaninejad, Xi Chen, and Pieter Abbeel. A simple neural attentive meta-learner. In *International Conference on Learning Representations*, 2018. 3
- [19] Tsendsuren Munkhdalai and Hong Yu. Meta networks. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 2554–2563, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR. 3
- [20] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *European conference on computer vision*, pages 483–499. Springer, 2016. 1
- [21] Dario Pavlo, Christoph Feichtenhofer, David Grangier, and Michael Auli. 3d human pose estimation in video with temporal convolutions and semi-supervised training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7753–7762, 2019. 1, 2, 6, 7, 8, 11, 12
- [22] Aravind Rajeswaran, Chelsea Finn, Sham M Kakade, and Sergey Levine. Meta-learning with implicit gradients. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. 3
- [23] S. Ravi and H. Larochelle. Optimization as a model for few-shot learning. In *ICLR*, 2017. 3

- [24] Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. Meta-learning with memory-augmented neural networks. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1842–1850, New York, New York, USA, 20–22 Jun 2016. PMLR. [3](#)
- [25] Assaf Shocher, Nadav Cohen, and Michal Irani. “zero-shot” super-resolution using deep internal learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3118–3126, 2018. [3](#)
- [26] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. [3](#)
- [27] Jae Woong Soh, Sunwoo Cho, and Nam Ik Cho. Meta-transfer learning for zero-shot super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3516–3525, 2020. [3](#)
- [28] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. [1](#)
- [29] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip H.S. Torr, and Timothy M. Hospedales. Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. [3](#)
- [30] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, koray kavukcuoglu, and Daan Wierstra. Matching networks for one shot learning. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016. [3](#)
- [31] Jianhua Wang, Fanhuai Shi, Jing Zhang, and Yuncai Liu. A new calibration model of camera lens distortion. *Pattern Recognition*, 41(2):607–615, 2008. [4](#)
- [32] Jingwei Xu, Zhenbo Yu, Bingbing Ni, Jiancheng Yang, Xiaokang Yang, and Wenjun Zhang. Deep kinematics analysis for monocular 3d human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. [1](#)
- [33] Y. Yan, B. Ni, W. Zhang, J. Xu, and X. Yang. Structure-constrained motion sequence generation. *IEEE Transactions on Multimedia*, 21(7):1799–1812, 2019. [1](#)
- [34] Jianfeng Zhang, Xuecheng Nie, and Jiashi Feng. Inference stage optimization for cross-scenario 3d human pose estimation. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. [2](#), [3](#), [5](#), [6](#)
- [35] Pengfei Zhang, Cuiling Lan, Wenjun Zeng, Junliang Xing, Jianru Xue, and Nanning Zheng. Semantics-guided neural networks for efficient skeleton-based human action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. [1](#)
- [36] Long Zhao, Xi Peng, Yu Tian, Mubbasir Kapadia, and Dimitris N Metaxas. Semantic graph convolutional networks for 3d human pose regression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3425–3435, 2019. [1](#), [2](#), [6](#), [7](#), [12](#)