

Prior-guided Source-free Domain Adaptation for Human Pose Estimation

Dripta S. Raychaudhuri^{1,2*} Calvin-Khang Ta¹ Arindam Dutta¹ Rohit Lal¹ Amit K. Roy-Chowdhury¹

¹University of California, Riverside ²AWS AI Labs

{drayc001@, cta003@, adutt020@, rlal011@, amitrc@ece.}ucr.edu

Abstract

Domain adaptation methods for 2D human pose estimation typically require continuous access to the source data during adaptation, which can be challenging due to privacy, memory, or computational constraints. To address this limitation, we focus on the task of source-free domain adaptation for pose estimation, where a source model must adapt to a new target domain using only unlabeled target data. Although recent advances have introduced source-free methods for classification tasks, extending them to the regression task of pose estimation is non-trivial. In this paper, we present *Prior-guided Self-training (POST)*, a pseudo-labeling approach that builds on the popular Mean Teacher framework to compensate for the distribution shift. *POST* leverages prediction-level and feature-level consistency between a student and teacher model against certain image transformations. In the absence of source data, *POST* utilizes a human pose prior that regularizes the adaptation process by directing the model to generate more accurate and anatomically plausible pose pseudo-labels. Despite being simple and intuitive, our framework can deliver significant performance gains compared to applying the source model directly to the target data, as demonstrated in our extensive experiments and ablation studies. In fact, our approach achieves comparable performance to recent state-of-the-art methods that use source data for adaptation.

1. Introduction

Human pose estimation is a fundamental task in computer vision that involves determining the precise locations of keypoints, such as joints, on a human body in an image or video [1]. The growing need for pose estimation in various applications such as action recognition [2], human-computer interaction [3], and video surveillance [4] has driven the rapid development of highly accurate deep learning tech-

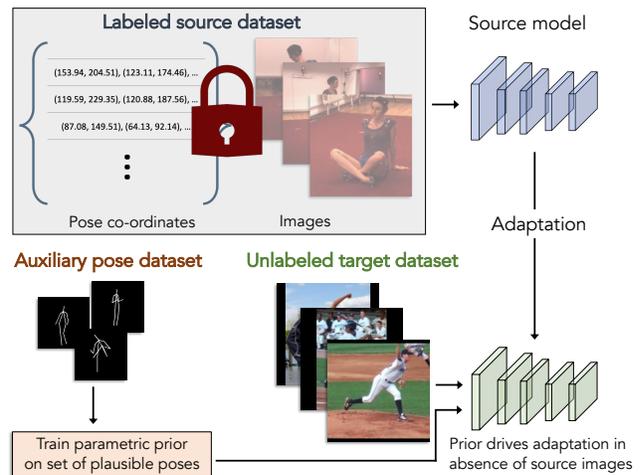


Figure 1. **Problem setup.** Existing UDA methods for pose estimation rely on a labeled source dataset while adapting to an unlabeled target dataset. However, privacy concerns surrounding the use of personally identifiable information in these labeled datasets, as well as the significant storage and computational requirements, can limit access to such data. Hence, our work focuses on source-free UDA of pose estimation models.

niques. However, the challenge of obtaining large annotated datasets for training, compounded with the susceptibility to a performance decline in the face of data distribution shifts still poses limitations for current pose estimation models.

To overcome these limitations, recent studies have focused on unsupervised domain adaptation (UDA) of pose estimators [5, 6]. UDA allows for transferring a pose estimation model trained on a source domain, where labeled data is available, to a target domain where labeled data is unavailable. Despite improved and robust pose estimation, the requirement of simultaneous access to both source and target domains during adaptation hinders real-world implementation. For instance, the labeled source data may not be accessible post-deployment due to privacy or proprietary issues. This is particularly relevant for human pose datasets, which contain *personally identifiable information (PII)* [7]. Furthermore, adaptation using the entire source data might

* Currently at AWS AI Labs. Work done while the author was at UCR.

be infeasible due to both memory and computational constraints. In light of these issues, we focus on *source-free* UDA of human pose estimation models.

Concretely, our objective is to adapt a 2D human pose estimation model to a new target domain utilizing only its trained parameters and unlabeled target data. This presents a major challenge as the absence of source data for regularization can cause catastrophic forgetting. While recent advances have introduced methods to tackle this issue in classification tasks [8, 9, 10], extending them to the regression task of pose estimation is non-trivial. To address this challenge, we introduce Prior-guided Self-training (POST), a self-training regression framework that employs a human pose prior to effectively guide the adaptation process in the absence of source data. An overview of our problem setup is shown in Figure 1.

Our approach builds on the Mean Teacher [11] framework, which uses consistency in the prediction space of a student and teacher model to produce trustworthy pseudo-labels and learn from the unlabeled target domain. To achieve this, we create two augmented views of each target image, varying in scale, spatial context, and color statistics. Aligned pose predictions from both models in both views are then obtained, and consistency between the predictions across the different views is encouraged to facilitate *prediction space adaptation*. However, our empirical results show that relying solely on consistency in the output space is insufficient when supervision from the source data is lacking. To address this, we also introduce *feature space adaptation*, which aims to encourage consistency across features extracted from the two separate views. We adopt the Barlow Twins [12] approach to accomplish this. Specifically, we seek to make the cross-correlation matrix calculated from a pair of feature embeddings as close to the identity matrix as possible.

In addition to the adaptation across both outputs and features, we employ a human pose prior that models the full manifold of plausible poses in some high-dimensional pose space to refine possible noisy pseudo-labels that may arise during self-training. The plausible poses are represented as points on the manifold surface, with zero distance from it, while non-plausible poses are located outside the surface, with a non-zero distance from it. This manifold is learned using a high-dimensional neural field, similar to Pose-NDF [13]. The pose prior acts as a strong regularizer, directing the model to generate more accurate pose pseudo-labels on the target data and leading to improved adaptation. The learning of this prior requires an auxiliary dataset of plausible human poses, but this does not compromise the privacy aspect of our framework as the prior does not make use of RGB images. In addition, it is worth noting that the prior can be trained offline, separately from the adaptation process. This not only saves computational resources but also reduces the amount of storage required. Compared to storing entire

images, it is much more efficient to store pose coordinates, which requires approximately $3000\times$ less memory.

Main contributions. To summarize, our primary contributions are as follows:

- We address the problem of adapting a human pose estimation model to a target domain consisting of unlabeled data, without access to the original source dataset. This ameliorates the privacy concern associated with the current domain adaptive pose estimation methods.
- We introduce Prior-guided Self-training (POST), a simple source-free unsupervised adaptation algorithm. POST leverages both prediction-level and feature-level consistency, in addition to a human pose prior, to drive self-training for improved adaptation to the target domain.
- We evaluate our method qualitatively and quantitatively on three challenging domain adaptive scenarios, demonstrating comparable performance to existing UDA methods that have access to the source data.

2. Related Works

Pose Estimation. 2D human pose estimation aims to locate human anatomical keypoints, such as the elbow and knee. Prior works can be categorized into two primary frameworks: the top-down framework and the bottom-up framework. Top-down methods [14, 15, 16, 17, 18, 19, 1] first detect each person from the image and then perform single-person pose estimation on each bounding box independently. On the other hand, bottom-up methods [20, 21, 22, 23, 24, 25, 26] predict keypoints of each person directly in an end-to-end manner. Typical bottom-up methods consist of two steps: predicting keypoint heatmaps and grouping the detected keypoints into separate poses. In this work, we focus on the bottom-up framework for efficiency purposes and adopt the Simple Baseline [18] architecture following [6] to ensure fair comparisons with prior domain adaptation algorithms.

Unsupervised Domain Adaptation. UDA methods have been extensively applied to a broad range of computer vision tasks, including image classification [27], semantic segmentation [28], object detection [29], and reinforcement learning [30] to tackle the issue of data distribution shift. Most approaches aim to align the source and target data distributions through techniques such as maximum mean discrepancy [31] and adversarial learning [32, 27]. Another line of research utilizes image translation methods to perform adaptation by transforming the source images into the target domain [33, 34]. More recently, there has been a surge of interest in adaptation using only a pre-trained source model due to privacy and memory storage concerns related to the source data. These include techniques such as information maximization [8, 9], pseudo-labeling [35, 36] and self-supervision [37]. Compared to other tasks, domain adaptation for regression tasks, such as pose estimation, remains

relatively unexplored.

Domain Adaptive Pose Estimation. UDA methods for pose estimation have explored various techniques for overcoming the domain gap, including adversarial feature alignment and pseudo-labeling. RegDA [5] estimates the domain discrepancy by evaluating false predictions on the target domain and minimizes it. Mu *et al.* [38] proposed consistency regularization with respect to transformations and temporal consistency learning within a video. Li *et al.* [39] proposed a refinement module and a self-feedback loop to obtain reliable pseudo-labels. Recently, Kim *et al.* [6] introduced a unified framework for both human and animal keypoint detection, which aligns representations using input-level and output-level cues. Typically, these methods require access to the source data, which may raise data privacy, memory and computation concerns. In contrast, our method addresses the domain adaptation problem in a source-free setting.

3. Prior-guided Self-training

Our work investigates source-free domain adaptation for 2D human pose estimation. In the pose estimation task, given an input image $\mathbf{x} \in \mathbb{R}^{H \times W \times 3}$, the goal is to predict the corresponding $y \in \mathbb{R}^{K \times 2}$ representing the 2D coordinates of K keypoints using a pose regression model f . In this paper, we assume access to a pre-trained model, denoted by f_S , as well as N unlabeled images $\mathcal{D} = \{\mathbf{x}_i\}_{i=1}^N$ from a target domain \mathcal{T} . Our goal is to adapt the source model to the target such that it performs better on images drawn from the target distribution than when directly using the source model on the target images.

Overview. In the absence of source data, we propose to use self-training to adapt the source pose estimation model to the target domain. However, self-training methods are prone to error accumulation, particularly when labeled data is absent to act as regularization. Hence, we introduce POST, an enhanced self-training strategy that employs three essential ideas to prevent such errors:

1. A weight-averaged teacher model is used to generate the pseudo-labels for self-training. This ensures better retention of the source knowledge within the teacher model by reducing the effect of updating the weights via noisy pseudo-labels (Section 3.1.1).
2. In addition to adaptation over the output space via pseudo-labels, the model is adapted in the feature space as well. For each target image, two aligned predictive views are generated via data augmentation, and consistency across features extracted from the two separate views is encouraged via a contrastive learning strategy (Section 3.1.2).
3. A human pose prior is used to regularize the adaptation by directing the model to generate more accurate and anatomically plausible pose pseudo-labels (Section 3.2). An overview of our framework is presented in Figure 2.

3.1. Self-training via Mean Teacher

Motivated by research suggesting that weight-averaged models over training steps tend to perform better than the final model [6, 39], we utilize the Mean Teacher framework [11] to generate pseudo-labels for self-training. The framework involves creating two identical models, a teacher model f_{tea} and a student model f_{stu} , both of which are initialized with the pre-trained network f_S at time step $t = 0$. At each subsequent time step t , the student model parameters θ are updated by backpropagating the supervisory signals provided by the teacher model. The parameters of the teacher model $\tilde{\theta}$ are updated via an exponential moving average (EMA) of the student model parameters:

$$\tilde{\theta}_t = \alpha \tilde{\theta}_{t-1} + (1 - \alpha) \theta_t, \quad (1)$$

where α denotes the smoothing coefficient which is set to 0.999 by default. The EMA update prevents the teacher model from overfitting to noisy pseudo-labels during the initial rounds of self-training, thereby, preserving the source knowledge. This is especially advantageous in our scenario where source data is unavailable to regularize the adaptation.

In the following sections, we demonstrate how to adapt f_{stu} using supervisory signals from f_{tea} on both the feature space and the pose prediction space.

3.1.1 Prediction Space Adaptation

At each time step t , we apply two different data augmentations, \mathcal{A}_1 and \mathcal{A}_2 , to a target image \mathbf{x} to generate two views. We then obtain the keypoint heatmap corresponding to the first transformed image $\tilde{h}^t = f_{tea}^t(\mathcal{A}_1(\mathbf{x}))$. Here $\tilde{h}^t \in \mathbb{R}^{K \times H' \times W'}$ denotes the spatial likelihood of the K different keypoints on each channel. The pseudo-label from the teacher model is generated by obtaining the coordinates which produce the maximum activations $\hat{y}^t = \operatorname{argmax}_u \tilde{h}_{[:,u]}^t$.

To reduce the influence of erroneous pseudo-labels in our training process, we utilize a confidence threshold to discard potentially unreliable labels. Specifically, we only retain the keypoint activations with the top $p\%$ maximum values among all activations and discard the rest. We set the threshold τ accordingly to reflect this, thereby ensuring that only the most confident labels are used for training.

Following prior work on supervised pose estimation, we first convert the pseudo-labels to normalized Gaussian heatmaps [40] and then use the mean squared error (MSE) loss to update the student model over batches of target images sampled from \mathcal{D} :

$$\mathcal{L}_{\text{out}} = \frac{1}{|\mathcal{B}|} \sum_{\mathbf{x} \in \mathcal{B}} \sum_{k=1}^K \mathbb{1}(\hat{h}_k^t \geq \tau) \|\tilde{\mathcal{A}}_1(\hat{h}_k^t) - \tilde{\mathcal{A}}_2(h_k^t)\|_2. \quad (2)$$

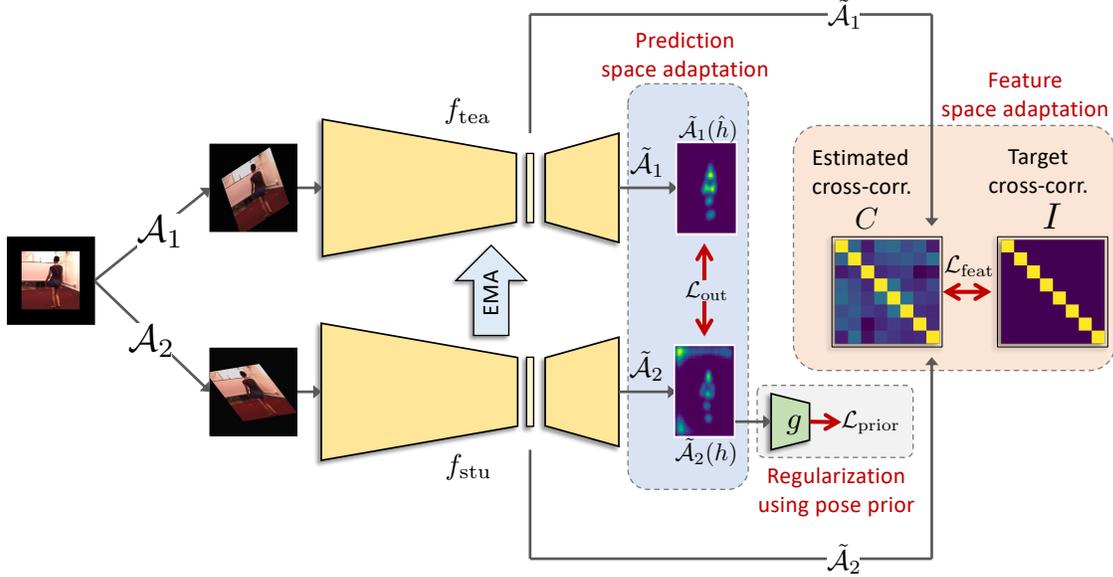


Figure 2. **Framework overview.** Our approach builds on the Mean Teacher framework and performs adaptation both in the pose prediction space using \mathcal{L}_{out} , and the feature space using $\mathcal{L}_{\text{feat}}$. This is supplemented by a human pose prior g that scores the predicted pseudo-labels in terms of plausibility. These scores are used to regularize the adaptation process in the absence of labeled source data via $\mathcal{L}_{\text{prior}}$. The student model f_{stu} is trained by the combination of the three losses, while the teacher model f_{tea} is updated with the exponential moving average (EMA) of the weights of the student model.

We denote the inverse functions of the chosen augmentations \mathcal{A}_1 and \mathcal{A}_2 as $\tilde{\mathcal{A}}_1$ and $\tilde{\mathcal{A}}_2$, respectively. The heatmap predicted by the student model for the k -th keypoint is represented by $h_k^t = f_{\text{stu}}^t(\mathcal{A}_2(x))^k$, while the heatmap generated from the pseudo-labels predicted by the teacher model is represented by $\hat{h}_k^t = L(\hat{y}^t)^k$. Here, $L(\cdot)$ represents the heatmap generating function and \mathcal{B} denotes a batch of target images.

3.1.2 Feature Space Adaptation

Pose estimation models primarily rely on a high-to-low-resolution feature encoder to generate low-resolution representations, which are then used to recover high-resolution pose heatmaps [18]. Since the upsampling process is inherently noisy, providing intermediate supervision to explicitly adapt the features, in addition to adapting the output pose keypoints as shown in the previous section, can be beneficial [17]. While previous work has demonstrated the benefits of joint adaptation across the output and feature space via adversarial learning in tasks such as semantic segmentation [41], this has been exclusively focused on scenarios where source data is available. Here, we propose an alternative way to accomplish feature space adaptation in the absence of source data via contrastive learning.

We begin by creating two different views of each target image x using a pair of sampled augmentations \mathcal{A}_1 and \mathcal{A}_2 , as previously shown. Next, considering the pose estimation model as a composition of a feature encoder and an output regressor, *i.e.*, $f = \text{Dec} \circ \text{Enc}$, we extract

the augmentation reversed feature maps from the teacher model $\tilde{z} = \tilde{\mathcal{A}}_1(\text{Enc}_{\text{tea}}(\mathcal{A}_1(x)))$ and the student model $z = \tilde{\mathcal{A}}_2(\text{Enc}_{\text{stu}}(\mathcal{A}_2(x)))$. We extract pairs of features for every image in a batch, pool them along the spatial dimensions, and then normalize them along the batch dimension to ensure that each covariate has a mean of 0 over the batch. For simplicity, we overload z, \tilde{z} to represent the normalized features, and drop the time index t .

We utilize feature-level consistency between the different views in order to accomplish feature space adaptation. This is achieved via a contrastive learning strategy which encourages the cross-correlation matrix between the outputs of the two networks to be as close to the identity matrix as possible [12],

$$\mathcal{L}_{\text{feat}} = \sum_i (1 - C_{ii})^2 + \gamma \sum_i \sum_{j \neq i} C_{ij}^2. \quad (3)$$

We define γ as a positive constant that balances the importance of the first and second terms of the loss function. C represents the cross-correlation matrix computed between the outputs of the student and teacher networks along the batch dimension:

$$C_{ij} = \frac{\sum_{b=1}^{|\mathcal{B}|} \tilde{z}_{b,i} z_{b,j}}{\sqrt{\sum_{b=1}^{|\mathcal{B}|} (\tilde{z}_{b,i})^2} \sqrt{\sum_{b=1}^{|\mathcal{B}|} (z_{b,j})^2}}. \quad (4)$$

Here, b represents the batch index, while i and j index the feature dimensions of the networks' outputs. We set $\gamma = 5e - 3$ following [12].

The first term of $\mathcal{L}_{\text{feat}}$ encourages the consistency of pose features within the same image by equating the diagonal elements of the cross-correlation matrix to 1, effectively reversing the effects of augmentations. In contrast, the second term aims to decorrelate the different feature dimensions of the embedding by forcing the off-diagonal elements of the cross-correlation matrix to 0.

3.2. Regularization via Pose Prior

Enhancing the performance of the pose estimation model on the target domain through joint adaptation over the output and feature space is undoubtedly valuable. However, this approach has its limitations, as it relies solely on general domain adaptation principles and overlooks the rich structural priors associated with human poses. To address this issue, we propose incorporating a parametric human pose prior to better adapt the pose estimation model to the target domain.

3.2.1 2D Human Pose Prior

Building upon the work of [13], we propose a human pose prior modeled as a manifold consisting of plausible 2D poses. To represent the 2D poses while ignoring aspects such as size and scale, we use a set of 2D orientation vectors that connect pairs of joints in the human skeleton, denoted by $\mathcal{G} = \{\theta = (\theta_1, \dots, \theta_L) \mid \theta_l \in \mathbb{R}^2, \|\theta_l\|_2 = 1 \forall l \in [L]\}$. We assume that plausible 2D human poses lie on a manifold embedded in this pose space \mathcal{G} . We use a function $g : \mathcal{G} \rightarrow \mathbb{R}^+$, which maps a pose to a non-negative scalar, to represent the manifold of plausible poses as the zero-level set:

$$\mathcal{P} = \{\theta \in \mathcal{G} \mid g(\theta) = 0\}, \quad (5)$$

where g represents the unsigned distance to the manifold. We construct this distance function by first encoding the pose using a hierarchical network g_{enc} that encodes the human pose based on its anatomical structure [42], and subsequently, use g_{dec} to predict the distance based on the pose representation.

Specifically, for a given pose θ , we encode it as follows,

$$\mathbf{v}_1 = g_{\text{enc}}^1(\theta_1), \quad \mathbf{v}_l = g_{\text{enc}}^l(\theta_l, \mathbf{v}_{\Omega(l)}), l \in \{2, \dots, L\}. \quad (6)$$

Here, $\Omega(l)$ is a function that maps the index of each orientation vector to its parent orientation vector in the kinematic chain of the human skeleton. We obtain the overall pose encoding as $\mathbf{p} = [\mathbf{v}_1 \mid \dots \mid \mathbf{v}_L]$ by concatenating all the individual orientation encodings. This pose encoding is then processed by $g_{\text{dec}} : \mathbb{R}^{d \cdot L} \rightarrow \mathbb{R}^+$, which predicts the unsigned distance for the given pose representation \mathbf{p} . A lower distance value for a pose implies that the configuration of joints is more likely to be a plausible human pose.

3.2.2 Pose Prior Training

To train the parametric prior g , we rely on an auxiliary dataset of M human poses $\mathcal{D}_A = \{\theta^i\}_{i=1}^{N_A}$, where $\theta^i =$

$(\theta_1^i, \dots, \theta_L^i)$. Importantly, these poses are not associated with their corresponding RGB images, which preserves the privacy aspect of the method. Additionally, storing only the pose coordinates instead of entire images makes data storage much more efficient and feasible.

We adopt a supervised approach to train g to predict the $L2$ distance to the plausible pose manifold for a given pose. To achieve this, we construct a dataset $\tilde{\mathcal{D}} = \{(\theta_i, d_i)\}_{i=1}^M$, consisting of pose and distance pairs, from \mathcal{D}_A . As the poses from \mathcal{D}_A lie on the desired manifold, we assign $d = 0$ to all poses in the dataset. To diversify our training samples, we randomly generate negative samples with distance $d > 0$ by perturbing the poses from \mathcal{D}_A with noise.

We train the network with the standard $L1$ loss,

$$\mathcal{L}_{\text{dist}} = \sum_{(\theta, d) \in \tilde{\mathcal{D}}} \|g(\theta) - d\|_1. \quad (7)$$

More details on the training process of the prior are provided in Section 4.2.

3.2.3 Adaptation using Pose Prior

We leverage the trained pose prior g to regularize the adaptation process by incentivizing the pose estimator to generate pseudo-labels that resemble plausible human poses.

Given the heatmaps $\{h_k^t\}_{k=1}^K$ generated for each keypoint by the student model f_{stu}^t for a target image \mathbf{x} , we calculate the corresponding orientation vectors in a differentiable manner to evaluate the plausibility of the predicted pose using the prior. First, we renormalize each heatmap to a probability distribution via spatial softmax and condense it to a point by computing the spatial expected value of the latter. For computational efficiency, we carry this out in a separable manner along the two spatial dimensions. Namely, assuming $u = (u_1, u_2)$ to be the two components of each pixel coordinate, we set

$$u_i^k = \frac{\sum_{u_i} u_i e^{h_k^t(u_i)}}{\sum_{u_i} e^{h_k^t(u_i)}}, \quad h_k^t(u_i) = \sum_{u_j} h_k^t(u_1, u_2), \quad (8)$$

where $i = 1, 2$ and $j = 2, 1$ respectively. Next, we use the pose coordinates to determine the orientation vectors between pairs of connected keypoints. Specifically, for every pair (a, b) of connected keypoints in the human skeleton (denoted by the set \mathcal{E}), we calculate the unit vector θ in the direction from u^a to u^b , where u^a and u^b are the estimated softmax coordinates of keypoints a and b , respectively, i.e.,

$$\theta = \frac{u^a - u^b}{\|u^a - u^b\|_2}, \quad \forall (a, b) \in \mathcal{E}. \quad (9)$$

Finally, we use the prior as a regularization term to minimize the distance of the current pose from our learned manifold,

$$\mathcal{L}_{\text{prior}} = \frac{1}{|\mathcal{B}|} \sum_{\mathbf{x} \in \mathcal{B}} g(T(f_{\text{stu}}^t(\mathbf{x}))), \quad (10)$$

where $T(\cdot)$ converts the predicted heatmaps to the orientation vector format required as input by the prior.

3.3. Overall Adaptation

The final training objective for the student model f_{stu} can be expressed as:

$$\min_{f_{\text{stu}}} \mathcal{L}_{\text{out}} + \lambda_1 \mathcal{L}_{\text{feat}} + \lambda_2 \mathcal{L}_{\text{prior}}. \quad (11)$$

Here, λ_1 and λ_2 are hyper-parameters that control the influence of feature space adaptation and prior regularization, respectively. The teacher model f_{tea} is updated asynchronously by computing an exponential moving average of the student model weights as shown in Equation 1.

4. Experiments

In this section, we demonstrate POST’s ability to adapt a 2D human pose estimation model to a target domain using only unlabeled data from that domain. We conduct experiments on three domain adaptive scenarios and compare with state-of-the-art domain adaptation baselines that utilize source data during adaptation. We also conduct extensive analysis to analyze the contribution and interaction between each component in our framework.

4.1. Datasets

SURREAL: SURREAL [43] is a large-scale dataset of synthetically generated images of people rendered from 3D sequences of human motion capture data against indoor backgrounds. It contains over 6 million frames, making it one of the largest and most diverse datasets of its kind.

Human3.6M: Human3.6M [44] is a real-world video dataset captured in indoor environments, comprising 3.6 million frames. The dataset features human subjects performing various actions. In order to reduce redundancy and computational complexity, we down-sampled the videos from 50fps to 10fps as per the approach proposed in [5]. For training, we follow the standard protocol proposed in [6] and use 5 subjects (S1, S5, S6, S7, S8), while the remaining 2 subjects (S9, S11) are reserved for testing.

LSP: Leeds Sports Pose (LSP) [45] is a real-world dataset that contains 2,000 images with annotated human body joint locations collected from sports activities. The images in LSP are captured in the wild, featuring a wide variety of human poses that are often challenging to detect.

BRIAR: BRIAR [46] is a cutting-edge biometric dataset featuring a large-scale collection of videos of human subjects captured in extremely challenging conditions. The videos are recorded at varying distances, *i.e.*, close range, 100m, 200m, 400m, 500m, and unmanned aerial vehicles (UAV), with each video lasting around 90 seconds. We randomly

sample 20 frames from each sequence for each of the 158 subjects for our experiments.

4.2. Experiment Protocols

Pose estimation model. We adopt the Simple Baseline [18] as our pose estimation model, with the ResNet-101 [47] as the backbone. To train the model, we use the Adam optimizer [48] with a base learning rate of $1e-4$, scheduled to decrease to $1e-5$ after 5 epochs and to $1e-6$ after 20 epochs. The model is trained for 30 epochs. We use a batch size of 32 and run 500 iterations per epoch. To threshold the model predictions, we set the confidence thresholding ratio p to 0.5. To augment the images during training, we follow [6] and use rotation, translation, shear, and Gaussian blur. The hyper-parameters λ_1 and λ_2 are set to $1e-3$ and $1e-6$, respectively.

Pose prior model. The training of the parametric prior follows a multi-stage approach that involves using different types of training samples. Initially, we use a combination of manifold poses θ_m and non-manifold poses θ_{nm} with a considerable distance from the desired manifold. Over the course of training, we gradually increase the number of non-manifold poses θ_{nm} with a small distance from the manifold. This enables our model to first learn a smooth surface and then gradually incorporate finer details as training progresses. We create these non-manifold poses θ_{nm} by injecting noise into the manifold poses θ_m obtained from the auxiliary dataset. Specifically, we sample directional noise from the Von-Mises distribution [49] and add it to the manifold poses in order to obtain the implausible poses.

The architecture for the encoder g_{enc} consists of a 2-layer MLP with an output feature size of $d = 6$ for each orientation vector, similar to [42]. The distance field network g_{dec} is implemented as a 5-layer MLP. Given its large size and diverse poses, we train the prior for our primary experiments using the SURREAL dataset.

4.3. Comparison with Baselines

Baselines. We evaluate the performance of our proposed method against several state-of-the-art domain adaptive frameworks. This includes adversarial learning-based feature alignment methods, such as DAN[31], DD[50], and RegDA[5]. Additionally, we consider approaches based on pseudo-labeling, namely CCSSL[38], and UDAPE [6]. It is worth noting that all these methods employ the source data during the adaptation process. To establish a comprehensive performance baseline, we report the results of two additional baselines: *Oracle* and *Source only*. The *Oracle* baseline represents the upper bound of the model’s performance, achieved by training the model jointly with target 2D annotations. On the other hand, *Source only* represents the model’s performance when it is directly applied to the target domain without any adaptation.

Table 1. **PCK@0.05 on SURREAL \rightarrow LSP.** (Best value is in red color, while the second best value is in blue color.)

Method	SF	Sld.	Elb.	Wrist	Hip	Knee	Ankle	Avg.
Source only	-	51.5	65.0	62.9	68.0	68.7	67.4	63.9
Oracle	-	-	-	-	-	-	-	-
DAN	✗	52.2	62.9	58.9	71.0	68.1	65.1	63.0
DD	✗	28.4	65.9	56.8	75.0	74.3	73.9	62.4
RegDA	✗	62.7	76.7	71.1	81.0	80.3	75.3	74.6
CCSSL	✗	36.8	66.3	63.9	59.6	67.3	70.4	60.7
UDAPE	✗	69.2	84.9	83.3	85.5	84.7	84.3	82.0
POST	✓	66.5	83.9	81.0	84.6	83.1	82.6	80.3

Table 2. **PCK@0.05 on SURREAL \rightarrow Human3.6M.**

Method	SF	Sld.	Elb.	Wrist	Hip	Knee	Ankle	Avg.
Source only	-	69.4	75.4	66.4	37.9	77.3	77.7	67.3
Oracle	-	95.3	91.8	86.9	95.6	94.1	93.6	92.9
DAN	✗	68.1	77.5	62.3	30.4	78.4	79.4	66.0
DD	✗	71.6	83.3	75.1	42.1	76.2	76.1	70.7
RegDA	✗	73.3	86.4	72.8	54.8	82.0	84.4	75.6
CCSSL	✗	44.3	68.5	55.2	22.2	62.3	57.8	51.7
UDAPE	✗	78.1	89.6	81.1	52.6	85.3	87.1	79.0
POST	✓	81.3	88.5	77.4	46.1	83.4	83.4	76.7

Metrics. We adopt the evaluation metric of Percentage of Correct Keypoint (PCK) for all experiments and report PCK@0.05 that measures the ratio of correct prediction within a range of 5% with respect to the image size.

4.4. Results

Quantitative results. We evaluate POST in two adaptation scenarios: SURREAL \rightarrow LSP and SURREAL \rightarrow Human3.6M, and report the quantitative results in Table 1 and Table 2, respectively. Specifically, we report the PCK@0.05 on 16 keypoints of the human body, including shoulders (sld.), elbows (elb.), wrists, hips, knees, and ankles. Our method achieves comparable results to many recent approaches that leverage source data for adaptation. Among these methods, UDAPE [6] achieves the highest performance on both cases, with our framework achieving a close second and only 2 percentage points behind on average. Notably, we outperform every other method, including the recently proposed RegDA [5] approach, by a significant margin of up to 5.7 percentage points. It is worth noting that not only do these methods require source data, but also involve additional models such as discriminators or style transfer modules, and unstable adversarial training. In contrast, our framework is lightweight, only involves pseudo-label training, and utilizes a simple prior model that can be trained offline.

Qualitative results. In addition to quantitative results, we also present qualitative results on SURREAL \rightarrow Human3.6M in Figure 3 and on SURREAL \rightarrow LSP in Figure 4. Also,

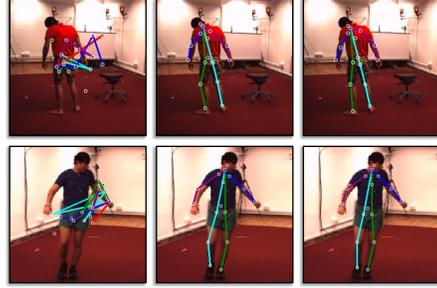


Figure 3. **Qualitative results on SURREAL \rightarrow Human3.6M.** We demonstrate sample results on the Human3.6M dataset. For each row, the leftmost shows the *Source only* prediction, the middle one shows the UDAPE [6] prediction, and the rightmost shows the prediction made POST.

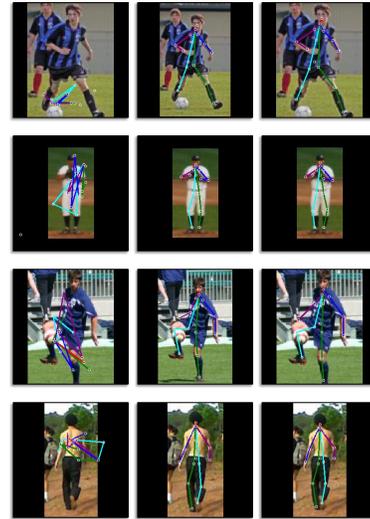


Figure 4. **Qualitative results on SURREAL \rightarrow LSP.** We demonstrate sample results on the LSP dataset. For each row, the leftmost shows the *Source only* prediction, the middle one shows the UDAPE [6] prediction, and the rightmost shows the prediction made POST.

we present only visual results on the SURREAL \rightarrow BRIAR adaptation scenario since pose annotations are absent in the BRIAR dataset. Figure 5 displays the predicted human poses on images taken from six different imaging ranges in the BRIAR dataset. While using the source model directly produces completely inaccurate poses, POST can accurately localize the keypoint locations, even in the presence of occlusions and atmospheric turbulence. Our approach can also accurately reconstruct poses even when the human is imaged from an elevated perspective, resulting in a high camera angle (Figure 5.f). Notably, our approach achieves this without using any source data for adaptation. Our results are comparable to those produced by UDAPE [6], which utilizes source data for adaptation. Additional qualitative results can be found in the supplementary document.



Figure 5. **Qualitative results on SURREAL \rightarrow BRIAR.** We demonstrate sample results on BRIAR at all ranges. For each range, we display three images: the leftmost shows the *Source only* prediction, the middle one shows the *UDAPE* [6] prediction, and the rightmost shows the prediction using *POST*. Although *POST* does not use source data for adaptation, it is able to match the predictions produced by *UDAPE*, which uses source data.

4.5. Ablation Studies

Effect of auxiliary dataset. We conduct experiments to evaluate the impact of the choice of the auxiliary dataset (used to train the prior) on downstream adaptation. The results are presented in Table 3. Our findings indicate that

POST is robust to the choice of the auxiliary dataset, with performance differences of ~ 0.5 percentage points.

Effect of loss terms. We conduct an experiment to evaluate the performance of each component of our framework. The results on the SURREAL \rightarrow Human3.6M adaptation scenario are presented in Table 4. Our findings indicate that in addi-

Table 3. **Effect of auxiliary dataset.** We evaluate the effect of auxiliary dataset on downstream adaptation tasks.

SURREAL→HUMAN3.6M							
Aux. dataset	Slid.	Elb.	Wrist	Hip	Knee	Ankle	Avg.
SURREAL	81.3	88.5	77.4	46.1	83.4	83.4	76.7
Human3.6M	80.9	88.0	77.2	45.0	83.1	82.8	76.2
SURREAL→LSP							
Aux. dataset	Slid.	Elb.	Wrist	Hip	Knee	Ankle	Avg.
SURREAL	66.5	83.9	81.0	84.6	83.1	82.6	80.3
Human3.6M	66.1	83.6	80.7	84.4	83.1	82.5	80.1

Table 4. **Effect of each loss term.** We evaluate the contribution of each loss term on SURREAL→Human3.6M.

\mathcal{L}_{out}	\mathcal{L}_{feat}	\mathcal{L}_{prior}	Slid.	Elb.	Wrist	Hip	Knee	Ankle	Avg.
✗	✗	✗	69.4	75.4	66.4	37.9	77.3	77.7	67.3
✓	✗	✗	77.9	86.7	73.7	38.8	83.0	84.3	74.1
✓	✓	✗	81.7	87.1	75.2	44.3	82.3	82.2	75.5
✓	✓	✓	81.3	88.5	77.4	46.1	83.4	83.4	76.7

tion to prediction space adaptation, feature space adaptation also plays a crucial role in enabling effective unsupervised learning from pseudo-labels. Moreover, we observed that the human pose prior brings additional improvements, thus validating our hypothesis that noisy pose pseudo-labels can be refined implicitly by a prior in the absence of source data. Overall, our results demonstrate the effectiveness of our framework and the importance of each of its components in achieving state-of-the-art performance.

Cross-dataset performance of prior. We evaluate the ability of the learned prior to handle distribution shift separately from the adaptation performance. Specifically, we assess the robustness of the prior by computing pose scores across datasets, and the results are presented in Figure 6. The plot demonstrates that our prior is effective in scoring plausible (real) poses with low scores and scoring implausible poses with higher scores. Note that the bell curve shape of the noisy pose scores is due to the Von-Mises noise added to create the noisy poses.

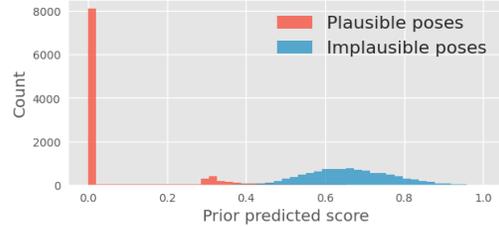
Effect of thresholding. We analyze the impact of the pseudo-label threshold τ on the adaptation performance in Table 5. The results on SURREAL→Human3.6M reveal that as we increase this ratio, the performance gradually decreases. This can be attributed to the fact that higher thresholding ratios tend to include lower confident predictions as pseudo-labels, which can negatively impact adaptation.

5. Conclusion

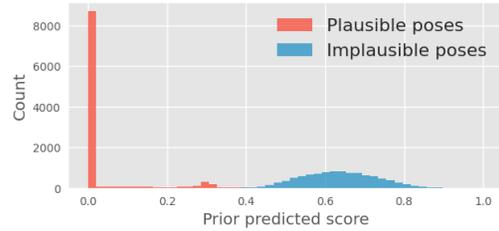
We address the problem of adapting a pre-trained 2D human pose estimator to a new target domain with unlabeled

Table 5. **Effect of τ .** We evaluate adaptation performance on SURREAL→Human3.6M as τ is varied.

$\tau = 0.1$	$\tau = 0.3$	$\tau = 0.5$	$\tau = 0.7$	$\tau = 0.9$
76.4	76.5	76.7	75.3	73.8



(a) Human3.6M → SURREAL



(b) SURREAL → Human3.6M

Figure 6. **Cross-dataset prior transfer.** We plot the histogram of scores predicted by the prior. The prior can clearly demarcate plausible poses from implausible poses across datasets.

target data. To this end, we propose a self-training algorithm, POST, that leverages a Mean Teacher framework to enforce both prediction-level and feature-level consistency between a pair of student and teacher models. Our approach incorporates a human pose prior that captures the manifold of possible poses in a high-dimensional space, which helps to refine noisy pseudo-labels generated during self-training. We evaluate our method on three challenging adaptation scenarios and show that it achieves competitive performance compared to existing UDA methods that have access to source data.

Acknowledgements. This research is based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via [2022-21102100007]. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

References

- [1] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *CVPR*, 2019. 1, 2

- [2] Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *AAAI*, 2018. 1
- [3] Hai Liu, Tingting Liu, Zhaoli Zhang, Arun Kumar Sangaiah, Bing Yang, and Youfu Li. Arhpe: Asymmetric relation-aware representation learning for head pose estimation in industrial human-computer interaction. *IEEE Transactions on Industrial Informatics*, 18(10):7107–7117, 2022. 1
- [4] Xiang Li, Yasushi Makihara, Chi Xu, Yasushi Yagi, Shiqi Yu, and Mingwu Ren. End-to-end model-based gait recognition. In *ACCV*, 2020. 1
- [5] Junguang Jiang, Yifei Ji, Ximei Wang, Yufeng Liu, Jianmin Wang, and Mingsheng Long. Regressive domain adaptation for unsupervised keypoint detection. In *CVPR*, 2021. 1, 3, 6, 7
- [6] Donghyun Kim, Kaihong Wang, Kate Saenko, Margrit Betke, and Stan Sclaroff. A unified framework for domain adaptive pose estimation. In *ECCV*, 2022. 1, 2, 3, 6, 7, 8
- [7] Paul M Schwartz and Daniel J Solove. The pii problem: Privacy and a new concept of personally identifiable information. *NYUL rev.*, 86:1814, 2011. 1
- [8] Jian Liang, Dapeng Hu, and Jiashi Feng. Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In *ICML*, 2020. 2
- [9] Sk Miraj Ahmed, Dripta S Raychaudhuri, Sujoy Paul, Samet Oymak, and Amit K Roy-Chowdhury. Unsupervised multi-source domain adaptation without access to source data. In *CVPR*, 2021. 2
- [10] Sujoy Paul, Ansh Khurana, and Gaurav Aggarwal. Unsupervised adaptation of semantic segmentation models without source data. *arXiv preprint arXiv:2112.02359*, 2021. 2
- [11] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *NeurIPS*, 2017. 2, 3
- [12] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *ICML*, 2021. 2, 4
- [13] Garvita Tiwari, Dimitrije Antic, Jan Eric Lenssen, Nikolaos Sarafianos, Tony Tung, and Gerard Pons-Moll. Pose-ndf: Modeling human pose manifolds with neural distance fields. In *ECCV*, 2022. 2, 5
- [14] Hao-Shu Fang, Shuqin Xie, Yu-Wing Tai, and Cewu Lu. Rmpe: Regional multi-person pose estimation. In *ICCV*, 2017. 2
- [15] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, 2017. 2
- [16] Shih-En Wei, Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh. Convolutional pose machines. In *CVPR*, 2016. 2
- [17] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hour-glass networks for human pose estimation. In *ECCV*, 2016. 2, 4
- [18] Bin Xiao, Haiping Wu, and Yichen Wei. Simple baselines for human pose estimation and tracking. In *ECCV*, 2018. 2, 4, 6
- [19] Yilun Chen, Zhicheng Wang, Yuxiang Peng, Zhiqiang Zhang, Gang Yu, and Jian Sun. Cascaded pyramid network for multi-person pose estimation. In *CVPR*, 2018. 2
- [20] Leonid Pishchulin, Eldar Insafutdinov, Siyu Tang, Bjoern Andres, Mykhaylo Andriluka, Peter V Gehler, and Bernt Schiele. Deepcut: Joint subset partition and labeling for multi-person pose estimation. In *CVPR*, 2016. 2
- [21] Eldar Insafutdinov, Leonid Pishchulin, Bjoern Andres, Mykhaylo Andriluka, and Bernt Schiele. Deepcut: A deeper, stronger, and faster multi-person pose estimation model. In *ECCV*, 2016. 2
- [22] Zhe Cao, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Openpose: realtime multi-person 2d pose estimation using part affinity fields. *IEEE transactions on Pattern Analysis and Machine Intelligence*, 43(1):172–186, 2019. 2
- [23] Sven Kreiss, Lorenzo Bertoni, and Alexandre Alahi. Pifpaf: Composite fields for human pose estimation. In *CVPR*, 2019. 2
- [24] Bowen Cheng, Bin Xiao, Jingdong Wang, Honghui Shi, Thomas S Huang, and Lei Zhang. Higherhrnet: Scale-aware representation learning for bottom-up human pose estimation. In *CVPR*, 2020. 2
- [25] Zigang Geng, Ke Sun, Bin Xiao, Zhaoxiang Zhang, and Jingdong Wang. Bottom-up human pose estimation via disentangled keypoint regression. *arXiv preprint arXiv:2104.02300*, 2021. 2
- [26] Sheng Jin, Wentao Liu, Enze Xie, Wenhai Wang, Chen Qian, Wanli Ouyang, and Ping Luo. Differentiable hierarchical graph grouping for multi-person pose estimation. In *ECCV*, 2020. 2
- [27] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *CVPR*, 2017. 2
- [28] Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schuster, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. Learning to adapt structured output space for semantic segmentation. In *CVPR*, 2018. 2
- [29] Han-Kai Hsu, Chun-Han Yao, Yi-Hsuan Tsai, Wei-Chih Hung, Hung-Yu Tseng, Maneesh Singh, and Ming-Hsuan Yang. Progressive domain adaptation for object detection. In *WACV*, 2020. 2
- [30] Dripta S Raychaudhuri, Sujoy Paul, Jeroen Vanbaar, and Amit K Roy-Chowdhury. Cross-domain imitation from observations. In *ICML*, 2021. 2
- [31] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In *ICML*, 2015. 2, 6
- [32] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. In *JMLR*, 2016. 2

- [33] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei A. Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *ICML*, 2018. 2
- [34] Luan Tran, Kihyuk Sohn, Xiang Yu, Xiaoming Liu, and Manmohan Chandraker. Gotta adapt 'em all: Joint pixel and feature-level domain adaptation for recognition in the wild. In *CVPR*, 2019. 2
- [35] Hao-Wei Yeh, Baoyao Yang, Pong C Yuen, and Tatsuya Harada. Sofa: Source-data-free feature alignment for unsupervised domain adaptation. In *WACV*, 2021. 2
- [36] Vikash Kumar, Rohit Lal, Himanshu Patil, and Anirban Chakraborty. Conmix for source-free single and multi-target domain adaptation. In *WACV*, 2023. 2
- [37] Haifeng Xia, Handong Zhao, and Zhengming Ding. Adaptive adversarial network for source-free domain adaptation. In *ICCV*, 2021. 2
- [38] Jiteng Mu, Weichao Qiu, Gregory D Hager, and Alan L Yuille. Learning from synthetic animals. In *CVPR*, 2020. 3, 6
- [39] Chen Li and Gim Hee Lee. From synthetic to real: Unsupervised domain adaptation for animal pose estimation. In *CVPR*, 2021. 3
- [40] Jonathan J Tompson, Arjun Jain, Yann LeCun, and Christoph Bregler. Joint training of a convolutional network and a graphical model for human pose estimation. *NeurIPS*, 2014. 3
- [41] Luan Tran, Kihyuk Sohn, Xiang Yu, Xiaoming Liu, and Manmohan Chandraker. Gotta adapt 'em all: Joint pixel and feature-level domain adaptation for recognition in the wild. In *CVPR*, 2019. 4
- [42] Marko Mihajlovic, Yan Zhang, Michael J Black, and Siyu Tang. Leap: Learning articulated occupancy of people. In *CVPR*, 2021. 5, 6
- [43] Gül Varol, Javier Romero, Xavier Martin, Naureen Mahmood, Michael J. Black, Ivan Laptev, and Cordelia Schmid. Learning from synthetic humans. In *CVPR*, 2017. 6
- [44] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(7):1325–1339, 2013. 6
- [45] Sam Johnson and Mark Everingham. Clustered pose and nonlinear appearance models for human pose estimation. In *BMVC*, 2010. 6
- [46] David Cornett, Joel Brogan, Nell Barber, Deniz Aykac, Seth Baird, Nicholas Burchfield, Carl Dukes, Andrew Duncan, Regina Ferrell, Jim Goddard, Gavin Jager, Matthew Larson, Bart Murphy, Christi Johnson, Ian Shelley, Nisha Srinivas, Brandon Stockwell, Leanne Thompson, Matthew Yohe, Robert Zhang, Scott Dolvin, Hector J. Santos-Villalobos, and David S. Bolme. Expanding accurate person recognition to new altitudes and ranges: The briar dataset. In *WACV Workshops*, 2023. 6
- [47] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 6
- [48] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [49] Riccardo Gatto and Sreenivasa Rao Jammalamadaka. The generalized von mises distribution. *Statistical Methodology*, 4(3):341–353, 2007. 6
- [50] Yuchen Zhang, Tianle Liu, Mingsheng Long, and Michael Jordan. Bridging theory and algorithm for domain adaptation. In *ICML*, 2019. 6