This WACV paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

POISE: Pose Guided Human Silhouette Extraction under Occlusions

Arindam Dutta^{1,*} Rohit Lal^{1,*} Dripta S. Raychaudhuri^{1,2,†} Calvin-Khang Ta¹ Amit K. Roy-Chowdhury¹ ¹University of California, Riverside ²AWS AI Labs

{adutt0200, rlal0110, drayc0010, cta0030, amitrc0ece.}ucr.edu

Abstract

Human silhouette extraction is a fundamental task in computer vision with applications in various downstream tasks. However, occlusions pose a significant challenge, leading to incomplete and distorted silhouettes. To address this challenge, we introduce POISE : Pose Guided Human Silhouette Extraction under Occlusions, a novel self-supervised fusion framework that enhances accuracy and robustness in human silhouette prediction. By combining initial silhouette estimates from a segmentation model with human joint predictions from a 2D pose estimation model, POISE leverages the complementary strengths of both approaches, effectively integrating precise body shape information and spatial information to tackle occlusions. Furthermore, the self-supervised nature of POISE eliminates the need for costly annotations, making it scalable and practical. Extensive experimental results demonstrate its superiority in improving silhouette extraction under occlusions, with promising results in downstream tasks such as gait recognition. The code for our method is available https://github.com/take2rohit/poise.

1. Introduction

Human silhouette extraction is a fundamental task in computer vision with wide-ranging applications in human motion analysis [30], surveillance [36], augmented reality [1], and human-computer interaction [2]. The precise extraction of human silhouettes from images and videos enables the accomplishment of higher-level tasks, including gait recognition [9] and activity recognition [18]. However, the presence of occlusions [3] poses substantial challenges to the extraction process, resulting in incomplete or distorted silhouettes that impede the performance of downstream tasks.

Occlusions are common in real-world scenarios, where human subjects navigate complex and cluttered environments. They occur when certain areas of the human body are



Figure 1. **Problem overview.** Existing works on segmentation struggle to effectively handle occlusion, leading to fragmented silhouettes (*Top row*). One approach to remedy this is by using 2D pose estimates, however, this fails to preserve the body shape (*Middle row*). In this paper, we propose POISE, a novel self-supervised framework that integrates both segmentation and 2D pose predictions to enable accurate prediction of human silhouettes under occlusion while preserving the body shape (*Bottom row*).

concealed or overlapped by objects or other individuals in the scene. Examples include objects obstructing body parts and self-occlusions caused by body parts coming into contact. State-of-the-art techniques for extracting human silhouettes, which involve pretrained semantic segmentation models like DeepLabv3 [5,6], struggle in such scenarios. Despite being trained on large-scale natural datasets, these models are not explicitly trained to handle occlusions. Consequently, they tend to misidentify occluded regions as background, resulting in fragmented and inaccurate silhouette predictions, as shown in Figure 1. Additionally, it is infeasible to acquire

^{*} Equal contribution.

[†] Currently at AWS AI Labs. Work done while the author was at UCR.

datasets with complete ground-truth human segmentation masks under occlusions due to visual uncertainty.

To overcome these challenges, we present a novel fusion framework, Pose Guided Silhouette Extraction under Occlusions (POISE), that significantly enhances the accuracy and robustness of human silhouette prediction. Our approach combines the predictions from a segmentation model, which initially estimates the silhouettes, with the human joint predictions from a 2D human pose estimation model. The pose estimation model provides crucial spatial information by predicting the keypoints and inferring the body structure, aiding in handling occlusions. However, it lacks detailed information about the body shape. This is addressed via the initial silhouettes predicted by the segmentation model for the unoccluded regions, capturing precise body shape information. By fusing these two complementary sources of information, we refine the silhouette predictions, resulting in silhouettes that are unfragmented and preserve body shape, as illustrated in Figure 1.

Existing approaches for jointly learning poses and silhouettes primarily rely on supervised learning [28, 50] which necessitates large annotated datasets for training. In this work, we relax this assumption and instead learn the fusion framework in an *self-supervised fashion*. Specifically, given pretrained segmentation and pose estimation models, we train the fusion model to mimic the predictions given by the pair via a pseudo-labeling approach. In order to utilize the pose predictions for silhouette estimation, we design a novel auxiliary transformation function to convert the sparse keypoint estimates into dense human silhouettes. Our self-supervised approach alleviates the need for costly annotations and makes the method more scalable and applicable to real-world scenarios.

Main contributions. Our primary contributions are summarized as follows:

- 1. We address the problem of self-supervised human silhouette extraction under occlusion.
- 2. We present an effective fusion framework that leverages pretrained 2D pose estimation and silhouette extraction models to produce accurate and robust human silhouettes.
- Our framework is based on self-supervised learning, which eliminates the need for costly pixel-level annotations or pose annotations, thereby enhancing scalability.
- 4. Beyond showcasing excellent results on silhouette extraction benchmarks, we also demonstrate the utility of the predicted silhouettes on downstream tasks such as gait recognition.

The rest of the paper is organized as follows: in Section 2, we review prior works in human silhouette extraction,

pose estimation, and gait recognition. Section 3 describes the proposed framework in detail. Experimental results and analysis are presented in Section 4, followed by conclusions and future directions in Section 5.

2. Related Works

Human Parsing and Silhouette Extraction. Extensive research has been conducted on human silhouette extraction, with earlier works such as [52] utilizing Hidden Markov models for modeling human silhouettes. More recent approaches leverage deep learning by primarily addressing the problem as semantic segmentation [4-6, 13, 26, 29, 34, 51]. These methods have achieved high accuracy, but require large labeled datasets of human images with corresponding segmentation masks. Pretrained models like DeepLabV3 [6] perform well in diverse settings due to their training on extensive datasets. Despite the success of these methods, occlusion remains a challenge. Zhou et. al. [55] proposed a multistage architecture for de-occluding humans, while Li et. al. [22] introduced a state-of-the-art human parsing framework that can also be utilized for silhouette extraction by mapping all parts of the human body to the foreground. However, these algorithms require labeled training data, which is a limiting factor.

Human Pose Estimation. Human pose estimation involves localizing keypoints on the human body, such as the head, elbows and knees, in 2D or 3D space. Deep learningbased pose estimation methods such as [27, 37, 44, 46], have achieved remarkable success on challenging academic datasets. However, these models are typically trained in supervised settings and often exhibit limited generalization capabilities when applied to unseen images. To overcome this, Zhang et. al. [48] proposed a novel domain adaptive 3D pose estimation algorithm. Jiang et. al. [16] proposed RegDA, a domain adaptive 2D pose estimation algorithm, which was further improved upon by Kim et. al. [17] in their work on UDAPE. Further, a source-free approach was recently proposed in [33]. These unsupervised methods have made significant contributions to enhancing pose estimation in settings where labeled data is unavailable.

Human Pose Estimation under Occlusions: Modern human pose estimation algorithms often fail to localize human keypoints under occlusions. To counter the same, Zhou *et. al.* [54] introduced a novel algorithm based on siamese networks and feature matching to improve 2D human pose estimation performance under occlusions. Cheng *et. al.* [8] exploits spatio-temporal continuity for handing occlusions thus leading to improved pose estimation. Qiu *et. al.* [31] used a novel graph formulation for improving human pose estimation under multi-person occlusion scenario. Liu *et. al.* [25] introduced a novel multi-stage framework for obtaining human keypoints under occlusions. Note that, these algorithms are entirely supervised learning algorithms necessitating costly annotations. To mitigate the same, Wang *et. al.* [41] introduced a novel contrastive learning based occlusion-aware algorithm for predicting 3D human pose from given 2D keypoints. However, this algorithm still needs access to ground-truth 2D keypoints.

Multi-Task Learning for Pose and Parsing. Multitask learning for human pose and body parsing aims to leverage the complementary information from both tasks to enhance their individual performances. Nie *et. al.* [28] introduced an adaptive convolutional architecture that facilitates joint human pose estimation and body parsing. Their approach was trained in a supervised manner, minimizing a linear combination of losses for each task. Liang *et. al.* [23] proposed a novel architecture that utilizes the semantic correlation between pose estimation and body parsing tasks. By exploiting this correlation, they were able to improve the accuracy of both tasks. More recently, Zhang *et. al.* [50] explored the explicit cross-task consistency between pose estimation and parsing. However, all of these methods are fully supervised and thus, require labeled data for both tasks.

Gait Recognition. Gait recognition, a fundamental task in computer vision, involves identifying individuals based on their unique walking patterns. Silhouettes have been extensively studied and utilized for gait recognition for nearly two decades. Pioneering works by Collins *et. al.* [9] and Wang *et. al.* [42] laid the groundwork for employing silhouettes in this context. Since then, numerous studies [40] have explored the use of silhouettes for gait recognition. As gait recognition under occlusions poses a challenge, it is crucial to preprocess silhouettes to handle occlusions before incorporating them into gait recognition systems. This preprocessing step ensures the robustness and reliability of the gait recognition process, even in the presence of occlusions.

3. Method

3.1. Problem Formulation

Given a dataset of images $\mathcal{T} = \{x_i\}_{i=1}^{N}$, where each $x_i \in \mathbb{R}^{H \times W \times 3}$, our goal is to train a model $\mathcal{M} : \mathbb{R}^{H \times W \times 3} \rightarrow \mathbb{R}^{H \times W}$ that accurately predicts the silhouette *y* corresponding to an input image *x*. To achieve this, we leverage the predictions of pretrained segmentation (\mathcal{S}) and 2D pose estimation (\mathcal{P}) models. These models provide initial predictions for the silhouette and sparse keypoints, respectively.

In order to align the sparse keypoints from \mathcal{P} with dense human silhouettes, we utilize a pose-to-silhouette transformation function $g : \mathbb{R}^{K \times 2} \to \mathbb{R}^{H \times W}$, where K represents the number of keypoints. We term this function **Pose2Sil**. By carefully fusing the predictions from S and the transformed keypoints using g, POISE aims to produce accurate and robust silhouettes.

3.2. Pose Guided Silhouette Extraction

For each image $x \in \mathcal{T}$, we generate two silhouettes, \mathcal{I}_{S} and $\mathcal{I}_{\mathcal{P}}$, using pretrained models S and \mathcal{P} , respectively. Specifically,

$$\mathcal{I}_{\mathcal{S}} = \mathcal{S}(x) \tag{1a}$$

$$\mathcal{I}_{\mathcal{P}} = g \circ \mathcal{P}(x) . \tag{1b}$$

Here, \mathcal{I}_{S} represents the silhouette obtained from the pretrained segmentation model, while \mathcal{I}_{P} is derived from the sparse pose predictions transformed by g.

These pair of silhouettes provide complementary information. While \mathcal{I}_{S} retains image-specific body shape details, it may become fragmented in the presence of occlusions. On the other hand, $\mathcal{I}_{\mathcal{P}}$ produces a continuous silhouette but may lose accuracy in capturing the body shape. Thus, our problem of obtaining robust silhouettes under occlusions boils down to learning a mixed representation of the two individually noisy silhouettes. To learn a robust silhouette representation, we train \mathcal{M} to effectively combine the information from these two noisy silhouettes under occlusions. This is accomplished by minimizing the pixel-wise binary cross-entropy loss between the predicted silhouette and those obtained from the pretrained models,

$$\mathcal{L}_{\mathcal{S}} = \mathcal{L}_{ce}\left(\mathcal{M}(x), \mathcal{I}_{\mathcal{S}}\right) \tag{2a}$$

$$\mathcal{L}_{\mathcal{P}} = \mathcal{L}_{ce} \left(\mathcal{M}(x), \mathcal{I}_{\mathcal{P}} \right) .$$
 (2b)

Notably, we train our model to focus exclusively on the foreground information from \mathcal{I}_S and ignore the background. This selective attention allows \mathcal{M} to concentrate solely on the identity-specific features derived from \mathcal{I}_S . This is particularly crucial since the segmentation model may misidentify body parts as background due to occlusion, rather than the other way around. In practice, this is carried out via a simple masking operation, which ignores all background pixels in \mathcal{I}_S .

We also use the pseudo-labels obtained from \mathcal{M} itself to regularize the training. Given an image x, we first obtain the pseudo-label $\mathcal{I}_{pl} = \mathcal{M}(x)$. These pseudo-labels are subsequently used to train the model using the pixel-wise binary cross-entropy loss,

$$\mathcal{L}_{pl} = \mathcal{L}_{ce} \left(\mathcal{M}(x), \mathcal{I}_{pl} \right) .$$
(3)

In order to prevent noisy predictions from impeding the learning process, we use a confidence threshold τ to mask out possible incorrect pseudo-labels. This ensures only the high-quality predictions are reinforced by the model.

The overall training objective is given by

$$\min_{\mathcal{M}} \lambda_1 \mathcal{L}_{\mathcal{S}} + \lambda_2 \mathcal{L}_{\mathcal{P}} + \lambda_3 \mathcal{L}_{pl} , \qquad (4)$$

where each λ_i , for $i \in [1,3]$, controls the influence of individual loss terms in generating the final silhouette. An overview of our framework can be found in Figure 2.



Figure 2. Framework overview. The key idea in POISE is to use two noisy silhouettes: \mathcal{I}_S and \mathcal{I}_P , obtained from S and $g \circ \mathcal{P}$ respectively, to train our model \mathcal{M} to predict robust silhouettes under occlusion. Note that, \mathcal{I}_S and \mathcal{I}_P provide complimentary information - \mathcal{I}_S retains identity specific shape features while \mathcal{I}_P provides a generic human silhouette for the pose keypoints. POISE learns a self-supervised fusion of \mathcal{I}_S and \mathcal{I}_P .



Figure 3. Learning Pose2Sil. The model g takes in the posekeypoints: $p \in \mathbb{R}^{K \times 2}$ as inputs and predicts corresponding silhouettes: $s \in \mathbb{R}^{H \times W}$. The model g is trained on a synthetic dataset \mathcal{D} and is trained by minimizing the binary cross-entropy loss between the predicted silhouettes and ground-truth silhouettes.

3.3. Obtaining Silhouette from Pose Keypoints

To leverage the valuable information derived from the pose keypoints for training \mathcal{M} , we incorporate a dedicated module, **Pose2Sil**, which facilitates the transformation from pose to silhouette. In this section, we outline the training methodology employed for this module.

We assume access to an auxiliary synthetic dataset $\mathcal{D} = \{(x_i, s_i, p_i)\}_{i=1}^{M}$ containing images of humans with corresponding annotations for pose keypoints p and segmentation masks s. Due to its synthetic nature, it is simple to obtain annotations [39], unlike the real-world images in \mathcal{T} . We

train g by minimizing the binary cross-entropy loss between the silhouette predicted from the pose and the corresponding ground-truth silhouette, $\mathcal{L}_{ce}(g(p), s)$. We illustrate this training process in Figure 3. Since the model g does not rely on RGB images during training or inference, it remains unaffected by domain changes. Thus, it is domain-agnostic and can be applied seamlessly across different domains without retraining.

3.4. Estimating Pose Keypoints under Occlusion

Due to the limited generalization capabilities of pretrained pose estimation models [16, 17], directly applying them to the images in \mathcal{T} may result in subpar keypoint localization, as illustrated in figure 4. Consequently, self-supervised domain adaptation of these pretrained models becomes imperative to ensure their effectiveness on our specific domain.

In this work, we build upon the state-of-the-art domain adaptive pose estimation algorithm, UDAPE [17]. We leverage the availability of a synthetic dataset \mathcal{D} as our source dataset, and \mathcal{T} serves as our target dataset. The images in \mathcal{T} may contain occlusions, which can significantly hinder the accuracy of pose estimation, as highlighted in previous research [54]. To enhance pose estimation under occlusions on the target dataset, we introduce occlusions into the source dataset itself [20]. By doing so, we force the model to learn robust representations that facilitate improved pose estimation under occlusions in the target dataset. Please note that we do not have access to the ground-truth labels of \mathcal{T} .

The adaptation algorithm follows the self-training framework inspired by the widely-used Mean-Teacher approach [38]. It involves the utilization of two identical models: a teacher model and a student model. Both models are initially initialized with the same weights at time step t = 0. Subsequently, at each



Figure 4. Limited generalization capacity of pre-trained pose estimation models. *Left:* Direct Inference of pre-trained model trained on dataset \mathcal{D} on BRIAR dataset. *Right:* Inference after adapting the pre-trained model to the same.

time step t, the student model parameters θ are updated by leveraging the supervisory signals provided by the teacher model, as well as the annotated data from the synthetic dataset \mathcal{D} . The parameters of the teacher model, denoted as $\tilde{\theta}$, are updated using an exponential moving average (EMA) of the student model parameters, ensuring a smoother and more stable learning process.



Figure 5. Adaptation of the pose estimation model. The algorithm is based on [17] with explicit occlusion handling by occlusion augmentation on the source dataset. A_1 and A_2 are augmentation operators such as rotation, scaling, etc. Similarly, A_1^{-1} and A_2^{-1} are inverse augmentation operators which essentially undo respective augmentations.

To update the student model, a combination of supervised and self-supervised losses is employed. The supervised loss \mathcal{L}_{sup} is computed using a mean square criterion on the source dataset, leveraging the annotated ground-truth labels. In addition to the supervised loss, a self-supervised consistency criterion \mathcal{L}_{unsup} is introduced to encourage consistency in the pose predictions of two different augmentations of an image. Figure 5 shows a schematic diagram of the adaptation algorithm for learning \mathcal{P} . Once the adaptation is done, we use the updated teacher model as \mathcal{P} to obtain the pose estimates on the images in \mathcal{T} . Figure 6 shows the importance of introducing occlusions in the source data which leads to improved pose estimation under occlusions.



Figure 6. Domain Adaptive Pose Estimation under Occlusions. *Left:* Inference of adapted model <u>without</u> occluding the source data (\mathcal{D}) . *Right:* Inference of adapted model <u>with</u> occlusions on the source data (\mathcal{D}) .

4. Experiments and Results

In this section, we provide a thorough assessment of POISE, highlighting its exceptional ability to accurately extract human silhouettes even when they are partially obscured. We evaluate POISE on five datasets, assessing its performance not only in human silhouette extraction but also in gait recognition. Our method outperforms existing off-the-shelf solutions and requires no extra annotations.

4.1. Datasets

We use the following datasets in our experiments.

- Humans3.6M [15] is large scale real-world video dataset with over 3 million frames featuring from 11 professional actors performing different actions such as walking, eating, etc.. Following standard protocol, we use subjects 'S1', 'S5', 'S6', 'S7' and 'S8' for training and subjects 'S9' and 'S11' for testing. Similar to [17], we use $\approx 20,000$ frames for training and another ≈ 3000 for evaluation. We use the mean-IoU (mIoU) metric [35] to report segmentation performance on this dataset.
- UP-S31 [21] is an image-based human part segmentation dataset with over 8000 images with 31 corresponding part annotations. However, several images containing more than one person were removed from our experiments leading to a total of 5226 images being used for our experiments. This was dis-jointly split into 4227 images for training and 1059 images for testing. We use the mean-IoU

(mIoU) metric [35] to report segmentation performance on this dataset.

- CASIA-B [45] is a large-scale multi-view indoor gait recognition dataset with 124 subjects across 11 views. Each subject has a total 10 sequences - six sequences under Normal (NM) conditions and two each under Carrying Bag (BG) and Wearing different Clothing (CL) conditions. We adhere to existing works such as [43] for training/testing and gallery/probe partitions. We report Rank-1 accuracy for gait recognition experiments using GaitBase [11] on this dataset.
- **BRIAR** [10] is a recent large-scale real-world biometric dataset with images of individuals under challenging conditions such as atmospheric turbulence and natural occlusions. For this work, we prepare the frames by tracking the person with Byte-Track [49]. We then select about 20,000 frames at different distances, allocating roughly 75% for training and the remainder for testing. Since we lack ground-truth silhouettes, we present qualitative silhouette extraction results. Additionally, we evaluate gait recognition using GaitBase [11] on a subset of 87 subjects divided into 60 for training and 27 for testing.
- **3DOH50K** [47] is large-scale real-world occlusion dataset with over 50000 frames. As the ground-truth segmentation masks are incomplete and images do not have ids associated with them, we report qualitative results for silhouette extraction on this dataset.

Generating occluded images. We use the datasets mentioned above to generate corresponding occluded RGB images for evaluation purposes. In particular, we consider two kinds of occlusions - Random Erase (RE) Occlusions [53] and Common Objects in Context (COCO) Occlusions [24]. For both Random Erase and COCO Occlusions, we first estimate keypoints on clean images using the adapted pose estimation model described in Section 3.4 and then add occlusions on randomly selected keypoints. This ensures that the occlusion always covers a certain part of the human body. Additional details on the same is presented in the section 1 of the supplementary.

4.2. Implementation details

In our experiments, we employ the DeepLabv3 architecture [7] with the ResNet-101 [14] feature extractor as the backbone for our model \mathcal{M} . For most of our experiments, we leverage the ResNet-101 pre-trained on the COCO dataset [24] as the segmentation network \mathcal{S} . However, for the specific scenario of COCO occlusions on Huamns3.6M and UP-S31 dataset, we adopt the SCHP architecture [22] pre-trained on the LIP dataset [12] as \mathcal{S} . For experiments with Random Erase Occlusions, \mathcal{M} is trained with five different severities of occlusion (12, 16, 20, 24, and 28), and the same model is used for inference on the five different severities of occlusion. Data augmentation strategies such

as random rotation, translation, and shear are also used to regularize the training process. Additional implementation details are provided in the section 2 of supplementary.

Synthetic Dataset. We use SURREAL [39] as our synthetic dataset \mathcal{D} . SURREAL is a large-scale dataset containing over 6 million frames of synthetically generated human images against an indoor background.

Pose Estimation Model (\mathcal{P}). Adhering to [17], we use the Simple Baseline decoder [44] with Resnet-101 backbone [14] the architectures for our student and teacher networks. The networks are trained for a total of 80 epochs with the first 40 epochs being used for supervised training of the student model and the other 40 epochs for adaptation to the target domain. We use a batch size 32 while optimizing using the Adam optimizer [19] with an initial learning rate of 1e - 4, decaying by a factor of 0.1 after 45^{th} and 60^{th} epochs.

Pose2Sil Model (g). The architecture of g is the same as that of the DCGAN [32]. The model is trained for a total of 200 epochs using the Adam optimizer [19] with a learning rate of 1e - 4 and a batch size of 32.

Table 1. Quantitative Results using mIoU metric for POISE against \mathcal{I}_{S} and \mathcal{I}_{P} on the Humans3.6M dataset with COCO occlusions.

Method	mIoU
$\mathcal{I}_{\mathcal{S}}$	80.14
$\mathcal{I}_{\mathcal{P}}$	75.97
POISE	87.22
POISE + Weak Sup.	89.67
Full Sup.	92.75

4.3. Results

4.3.1 Silhouette Extraction

We evaluate the effectiveness of POISE for human silhouette extraction under occlusion in terms of segmentation accuracy. Table 1 shows the efficacy of POISE is learning strong feature representations to perform human silhouette extraction under COCO occlusions on Humans3.6M dataset. While our results are inferior to a fully-supervised baseline by ≈ 5.5 %, we show that we can bridge this gap to ≈ 3 % by fine-tuning the network with limited supervision, i.e. considering 5 % of the training dataset as annotated. This shows that POISE learns generalized feature representations that can be effectively used to obtain optimal human silhouettes under occlusions. We report the mIoU on UP-S31 dataset at six different occlusion severities in Table 2. We note that for Random Erase occlusions, POISE provides an improvement of $\approx 8\%$ against $\mathcal{I}_{\mathcal{S}}$ and $\approx 3\%$ against $\mathcal{I}_{\mathcal{P}}$. For COCO occlusions, POISE provides an improvement of $\approx 2\%$ against $\mathcal{I}_{\mathcal{S}}$ and $\approx 3\%$ against $\mathcal{I}_{\mathcal{P}}$. This shows the deficiencies of using state-of-the-art human segmentation methods in sce-

Occlusion Severity	$\mathcal{I}_{\mathcal{S}}$	$\mathcal{I}_{\mathcal{P}}$	POISE
RE - 12	79.67	81.98	85.58
RE - 16	78.25	81.72	85.35
RE - 20	76.71	81.43	85.11
RE - 24	74.88	81.34	84.55
RE - 28	73.28	80.84	83.66
COCO	78.68	77.34	80.86

Table 2. Quantitative Results using mIoU metric for POISE against \mathcal{I}_{S} and \mathcal{I}_{P} on the UP-S31 dataset. RE - Random Erase Occlusion, COCO - COCO occlusion.

narios involving occlusion. We present additional qualitative results on the same in the section 3 of the supplementary.

We also assess the performance of POISE on the BRIAR dataset under natural occlusion scenarios. In Figure 7, we visually demonstrate how POISE excels at extracting silhouettes, closely resembling the original body shape from \mathcal{I}_S . In occluded regions, POISE relies on \mathcal{I}_P for guidance. Further, as shown in Figure 7, when dealing with human subjects of diverse body shapes (compared to the synthetic dataset \mathcal{D}), \mathcal{I}_P fails to retain body-specific information, while \mathcal{I}_S retains most of it. POISE combines information from both \mathcal{I}_S and \mathcal{I}_P to produce optimal silhouettes.

Figure 8 shows the efficacy of POISE in handling natural occlusions on the 3DOH50K dataset. As POISE uses complementary self-supervisory signals (from \mathcal{I}_S and \mathcal{I}_P) during training, it is able to handle occlusions reasonably well. The results are particularly interesting as it shows silhouettes from an self-supervised method i.e. POISE is better off as compared against ground-truth silhouettes. Additional results using recent segmentation methods are provided in section 4 of the supplementary.

4.3.2 Gait Recognition

In this section, we present the efficacy of POISE on the downstream task of gait recognition on two datasets: CASIA-B and BRIAR.

Results on CASIA-B: We present gait recognition results on the CASIA-B dataset for three different walking conditions: Normal (NM), Carrying Bag (BG) and Clothing (CL). We evaluate POISE under the following two different settings. **Duration of Occlusion.** In this setting the severity of the occlusion remains the same across frames but the number of occluded frames changes. As an example, 20% occlusion means that for every video in the dataset, a temporally continuous set of frames of length 20% of the total number of frames are occluded. Table 3 presents quantitative results on the performance of POISE against \mathcal{I}_S and \mathcal{I}_P at



Figure 7. Qualitative results on BRIAR. Silhouettes extracted using POISE compared against \mathcal{I}_{S} and \mathcal{I}_{P} on natural occlusions encountered in the BRIAR dataset. Note that while we use COCO occlusions for training, the model \mathcal{M} generalizes seamlessly to natural occlusions.



Figure 8. Qualitative results on 3DOH50K dataset [47] Silhouettes obtained by POISE compared against the original groundtruth silhouettes provided. Clearly, silhouettes from POISE are much more complete as compared against the ground-truth silhouettes.

fixed occlusion severity of 12. It is interesting to note that POISE outperforms both \mathcal{I}_{S} and \mathcal{I}_{P} by significant margins under scenarios of temporally long occlusions.

Severity of Occlusion. Similar to our experimental settings for section 4.3.1, we study the effect of varying occlusion severity on gait recognition, keeping the same duration of occlusion. Table 4 presents quantitative results on the performance of POISE against \mathcal{I}_S and \mathcal{I}_P at an occlusion duration of 50%. We note that, although with increasing severity of occlusion gait recognition performance falls for all the three listed methods, POISE still outperforms \mathcal{I}_S and \mathcal{I}_P by significant margins.

Results on BRIAR: Table 5 shows the efficacy of POISE in obtaining robust silhouettes that help with gait recognition

Table 3. Average Rank-1 gait Recognition Accuracy using GaitBase [11] across 11 different camera positions (0°, 18°, ..., 180°) for POISE against \mathcal{I}_{S} and \mathcal{I}_{P} on the CASIA-B dataset as a function of varying occlusion duration.

Method	20%		33%		50%		67%			80%					
	NM	BG	CL												
$\mathcal{I}_{\mathcal{S}}$	84.44	65.58	48.68	82.35	64.42	47.51	79.75	61.19	44.58	75.45	57.38	40.99	71.34	54.31	39.97
$\mathcal{I}_{\mathcal{P}}$	54.95	33.22	25.75	52.08	31.64	24.47	51.72	30.68	23.36	49.44	29.88	24.40	48.44	29.66	23.50
POISE	83.38	68.38	52.63	83.01	68.15	52.49	81.85	66.87	51.01	81.05	65.15	49.82	78.96	63.54	49.17

Table 4. Average Rank-1 gait Recognition Accuracy using GaitBase [11] across 11 different camera positions (0° , 18° , ..., 180°) for POISE against \mathcal{I}_S and \mathcal{I}_P on the CASIA-B dataset as a function of varying **occlusion severity**: 12, 16, ... 28.

Method	12		16		20		24			28					
	NM	BG	CL												
$\mathcal{I}_{\mathcal{S}}$	79.75	61.19	44.58	79.57	61.99	45.03	80.26	61.03	43.32	79.54	59.99	42.29	79.62	59.15	41.96
$\mathcal{I}_{\mathcal{P}}$	51.72	30.68	23.36	48.78	30.21	22.56	46.75	28.31	22.03	44.72	27.12	21.46	43.23	25.93	20.75
POISE	81.85	66.87	51.01	81.25	65.52	50.08	80.47	66.32	51.11	80.05	64.70	49.57	80.36	65.73	49.02

Table 5. Quantitative Results for top-k gait recognition accuracy for POISE against \mathcal{I}_{S} and \mathcal{I}_{P} on BRIAR dataset using GaitBase [11].

Method	top-1	top-3	top-5
$\overline{\mathcal{I}_{\mathcal{S}}}$	8.96	22.39	31.72
$\mathcal{I}_{\mathcal{P}}$	14.20	28.60	38.68
POISE	16.66	31.69	40.54

Table 6. **Ablation study results.** Pseudo-label loss significantly improves rank-1 gait recognition accuracy on CASIA-B dataset (occlusion duration: 50%, severity: 12). Best result is in bold and second result is underlined.

Method	$\mathcal{L}_{\mathcal{I}_{\mathcal{S}}}$	$\mathcal{L}_{\mathcal{I}_{\mathcal{P}}}$	\mathcal{L}_{pl}	NM	BG	CL	Avg.
$\overline{\mathcal{I}_{S}}$	\checkmark			80.22	62.26	43.31	61.93
$\mathcal{I}_{\mathcal{P}}$		\checkmark		52.23	31.07	24.79	36.03
POISE	\checkmark	\checkmark		81.33	63.89	54.35	<u>66.52</u>
POISE (full)	\checkmark	\checkmark	\checkmark	83.30	67.00	52.11	67.47

on the BRIAR dataset. We obtain improvements of $\approx 7 \%$ and 8 % in terms of top-1 and top-5 gait recognition accuracy over simply using \mathcal{I}_S . These improvements are significant as the videos in the dataset suffer from atmospheric turbulence and have instances of persistent natural occlusions.

A thorough discussion with additional quantitative analysis on gait recognition on the two aforementioned datasets is provided in section 5 of the supplementary.

4.3.3 Importance of Pseudo-Labeling:

We present an ablation study on the importance of using pseudo-labels from the model \mathcal{M} under training. Table 6 shows that there is an improvement of $\approx 1\%$ owing to the

use of the pseudo-label loss as described in Equation 3.

5. Conclusion

We present POISE, an innovative approach for robust human silhouette extraction under occlusions. By leveraging pose estimation and self-supervised learning, POISE effectively addresses the limitations posed by occlusions, ensuring accurate and preserved body shape representations. Unlike traditional methods, POISE eliminates the need for expensive annotations, making it cost-effective and practical. Experimental results demonstrate the superiority of POISE in improving silhouette extraction under occlusions. Furthermore, POISE exhibits promising performance in gait recognition tasks, underscoring its potential impact in various applications. In summary, POISE offers a straightforward yet valuable solution for extracting precise silhouettes in challenging scenarios, contributing to advancements in biometrics and related fields.

Acknowledgements. This research is partially supported by the Office of the Director of National Intelligence (ODNI), specifically through the Intelligence Advanced Research Projects Activity (IARPA), under contract number [2022-21102100007]. The views and conclusions in this research reflect those of the authors and should not be construed as officially representing the policies, whether explicitly or implicitly, of ODNI, IARPA, or the U.S. Government. Nevertheless, the U.S. Government retains the authorization to reproduce and distribute reprints for official government purposes, regardless of any copyright notices included. We further thank our colleague Yash Garg (ygarg002@ucr.edu) for helping us with additional experiments.

References

- Yinoussa Adagolodjo, Raffaella Trivisonne, Nazim Haouchine, Stéphane Cotin, and Hadrien Courtecuisse. Silhouettebased pose estimation for deformable organs application to surgical augmented reality. In 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 539–544. IEEE, 2017. 1
- [2] Biplab Ketan Chakraborty, Debajit Sarma, Manas Kamal Bhuyan, and Karl F MacDorman. Review of constraints on vision-based gesture recognition for human–computer interaction. *IET Computer Vision*, 12(1):3–15, 2018. 1
- [3] Changhong Chen, Jimin Liang, Heng Zhao, Haihong Hu, and Jie Tian. Frame difference energy image for gait recognition with incomplete silhouettes. *Pattern Recognition Letters*, 30(11):977–984, 2009. 1
- [4] Hao Chen, Kunyang Sun, Zhi Tian, Chunhua Shen, Yongming Huang, and Youliang Yan. Blendmask: Top-down meets bottom-up for instance segmentation, 2020. 2
- [5] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs, 2017. 1, 2
- [6] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation, 2017. 1, 2
- [7] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision* (ECCV), pages 801–818, 2018. 6
- [8] Yu Cheng, Bo Yang, Bo Wang, and Robby T Tan. 3d human pose estimation using spatio-temporal networks with explicit occlusion training. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 10631–10638, 2020.
- [9] Robert T Collins, Ralph Gross, and Jianbo Shi. Silhouettebased human identification from body shape and gait. In *Proceedings of fifth IEEE international conference on automatic face gesture recognition*, pages 366–371. IEEE, 2002. 1, 3
- [10] David Cornett, Joel Brogan, Nell Barber, Deniz Aykac, Seth Baird, Nicholas Burchfield, Carl Dukes, Andrew Duncan, Regina Ferrell, Jim Goddard, et al. Expanding accurate person recognition to new altitudes and ranges: The briar dataset. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 593–602, 2023. 6
- [11] Chao Fan, Junhao Liang, Chuanfu Shen, Saihui Hou, Yongzhen Huang, and Shiqi Yu. Opengait: Revisiting gait recognition towards better practicality. In *Proceedings of* the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 9707–9716, 2023. 6, 8
- [12] Ke Gong, Xiaodan Liang, Dongyu Zhang, Xiaohui Shen, and Liang Lin. Look into person: Self-supervised structuresensitive learning and a new benchmark for human parsing. In *Proceedings of the IEEE conference on computer vision* and pattern recognition, pages 932–940, 2017. 6

- [13] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn, 2018. 2
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceed-ings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 6
- [15] 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 Intelli*gence, 36(7):1325–1339, jul 2014. 5
- [16] Junguang Jiang, Yifei Ji, Ximei Wang, Yufeng Liu, Jianmin Wang, and Mingsheng Long. Regressive domain adaptation for unsupervised keypoint detection. In *Proceedings of* the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 6780–6789, 2021. 2, 4
- [17] Donghyun Kim, Kaihong Wang, Kate Saenko, Margrit Betke, and Stan Sclaroff. A unified framework for domain adaptive pose estimation. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXIII*, pages 603–620. Springer, 2022. 2, 4, 5, 6
- [18] Kibum Kim, Ahmad Jalal, and Maria Mahmood. Visionbased human activity recognition system using depth silhouettes: A smart home system for monitoring the residents. *Journal of Electrical Engineering & Technology*, 14:2567– 2573, 2019. 1
- [19] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980, 2014. 6
- [20] Jogendra Nath Kundu, Siddharth Seth, Pradyumna YM, Varun Jampani, Anirban Chakraborty, and R Venkatesh Babu. Uncertainty-aware adaptation for self-supervised 3d human pose estimation. In *Proceedings of the IEEE/CVF conference* on computer vision and pattern recognition, pages 20448– 20459, 2022. 4
- [21] Christoph Lassner, Javier Romero, Martin Kiefel, Federica Bogo, Michael J. Black, and Peter V. Gehler. Unite the people: Closing the loop between 3d and 2d human representations. In *IEEE Conf. on Computer Vision and Pattern Recognition* (CVPR), July 2017. 5
- [22] Peike Li, Yunqiu Xu, Yunchao Wei, and Yi Yang. Selfcorrection for human parsing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(6):3260–3271, 2020. 2, 6
- [23] Xiaodan Liang, Ke Gong, Xiaohui Shen, and Liang Lin. Look into person: Joint body parsing & pose estimation network and a new benchmark. *IEEE transactions on pattern analysis* and machine intelligence, 41(4):871–885, 2018. 3
- [24] 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 *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part* V 13, pages 740–755. Springer, 2014. 6
- [25] Qihao Liu, Yi Zhang, Song Bai, and Alan Yuille. Explicit occlusion reasoning for multi-person 3d human pose estimation.

In European Conference on Computer Vision, pages 497–517. Springer, 2022. 2

- [26] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. 2
- [27] 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. 2
- [28] Xuecheng Nie, Jiashi Feng, and Shuicheng Yan. Mutual learning to adapt for joint human parsing and pose estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 502–517, 2018. 2, 3
- [29] Hyeonwoo Noh, Seunghoon Hong, and Bohyung Han. Learning deconvolution network for semantic segmentation, 2015.
- [30] Ronald Poppe. Vision-based human motion analysis: An overview. *Computer vision and image understanding*, 108(1-2):4–18, 2007. 1
- [31] Lingteng Qiu, Xuanye Zhang, Yanran Li, Guanbin Li, Xiaojun Wu, Zixiang Xiong, Xiaoguang Han, and Shuguang Cui. Peeking into occluded joints: A novel framework for crowd pose estimation. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIX 16*, pages 488–504. Springer, 2020. 2
- [32] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015. 6
- [33] Dripta S Raychaudhuri, Calvin-Khang Ta, Arindam Dutta, Rohit Lal, and Amit K Roy-Chowdhury. Prior-guided sourcefree domain adaptation for human pose estimation. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 14996–15006, 2023. 2
- [34] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks, 2016. 2
- [35] Hamid Rezatofighi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian Reid, and Silvio Savarese. Generalized intersection over union: A metric and a loss for bounding box regression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 658–666, 2019. 5, 6
- [36] Vinay Sharma and James W Davis. Extraction of person silhouettes from surveillance imagery using mrfs. In 2007 IEEE Workshop on Applications of Computer Vision (WACV'07), pages 33–33. IEEE, 2007. 1
- [37] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep highresolution representation learning for human pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5693–5703, 2019. 2
- [38] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. Advances in neural information processing systems, 30, 2017. 5

- [39] Gul Varol, Javier Romero, Xavier Martin, Naureen Mahmood, Michael J Black, Ivan Laptev, and Cordelia Schmid. Learning from synthetic humans. In *Proceedings of the IEEE* conference on computer vision and pattern recognition, pages 109–117, 2017. 4, 6
- [40] Changsheng Wan, Li Wang, and Vir V Phoha. A survey on gait recognition. ACM Computing Surveys (CSUR), 51(5):1– 35, 2018. 3
- [41] Junjie Wang, Zhenbo Yu, Zhengyan Tong, Hang Wang, Jinxian Liu, Wenjun Zhang, and Xiaoyan Wu. Ocr-pose: Occlusion-aware contrastive representation for unsupervised 3d human pose estimation. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 5477–5485, 2022. 3
- [42] Liang Wang, Tieniu Tan, Huazhong Ning, and Weiming Hu. Silhouette analysis-based gait recognition for human identification. *IEEE transactions on pattern analysis and machine intelligence*, 25(12):1505–1518, 2003. 3
- [43] Zifeng Wu, Yongzhen Huang, Liang Wang, Xiaogang Wang, and Tieniu Tan. A comprehensive study on cross-view gait based human identification with deep cnns. *IEEE transactions* on pattern analysis and machine intelligence, 39(2):209–226, 2016. 6
- [44] Bin Xiao, Haiping Wu, and Yichen Wei. Simple baselines for human pose estimation and tracking. In *Proceedings of the European conference on computer vision (ECCV)*, pages 466–481, 2018. 2, 6
- [45] Shiqi Yu, Daoliang Tan, and Tieniu Tan. A framework for evaluating the effect of view angle, clothing and carrying condition on gait recognition. In 18th international conference on pattern recognition (ICPR'06), volume 4, pages 441–444. IEEE, 2006. 6
- [46] Feng Zhang, Xiatian Zhu, and Mao Ye. Fast human pose estimation. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 3517–3526, 2019. 2
- [47] Tianshu Zhang, Buzhen Huang, and Yangang Wang. Objectoccluded human shape and pose estimation from a single color image. In *IEEE Conference on Computer Vision and Pattern Recognition*, (CVPR), 2020. 6, 7
- [48] Xiheng Zhang, Yongkang Wong, Mohan S Kankanhalli, and Weidong Geng. Unsupervised domain adaptation for 3d human pose estimation. In *Proceedings of the 27th ACM International Conference on Multimedia*, pages 926–934, 2019. 2
- [49] Yifu Zhang, Peize Sun, Yi Jiang, Dongdong Yu, Fucheng Weng, Zehuan Yuan, Ping Luo, Wenyu Liu, and Xinggang Wang. Bytetrack: Multi-object tracking by associating every detection box. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXII*, pages 1–21. Springer, 2022. 6
- [50] Ziwei Zhang, Chi Su, Liang Zheng, and Xiaodong Xie. Correlating edge, pose with parsing. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 8900–8909, 2020. 2, 3
- [51] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network, 2017. 2

- [52] San-Lung Zhao and Hsi-Jian Lee. Human silhouette extraction based on hmm. In 18th International Conference on Pattern Recognition (ICPR'06), volume 2, pages 994–997, 2006. 2
- [53] Zhun Zhong, Liang Zheng, Guoliang Kang, Shaozi Li, and Yi Yang. Random erasing data augmentation. In *Proceedings* of the AAAI conference on artificial intelligence, volume 34, pages 13001–13008, 2020. 6
- [54] Lu Zhou, Yingying Chen, Yunze Gao, Jinqiao Wang, and Hanqing Lu. Occlusion-aware siamese network for human pose estimation. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX 16*, pages 396–412. Springer, 2020. 2, 4
- [55] Qiang Zhou, Shiyin Wang, Yitong Wang, Zilong Huang, and Xinggang Wang. Human de-occlusion: Invisible perception and recovery for humans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3691–3701, 2021. 2