

FuRPE: Learning Full-body Reconstruction from Part Experts

Zhaoxin Fan¹, Yuqing Pan¹, Hao Xu², Zhenbo Song³
Zhicheng Wang⁴, Kejian Wu⁴, Hongyan Liu⁵, Jun He^{1*}
¹Renmin University of China, ²Psyche AI Inc,
³Nanjing University of Science and Technology, ⁴Xreal,
⁵Tsinghua University

Abstract

In the field of full-body reconstruction, the scarcity of annotated data often impedes the efficacy of prevailing methods. To address this issue, we introduce FuRPE, a novel framework that employs part-experts and an ingenious pseudo ground-truth selection scheme to derive high-quality pseudo labels. These labels, central to our approach, equip our network with the capability to efficiently learn from the available data. Integral to FuRPE is a unique exponential moving average training strategy and expert-derived feature distillation strategy. These novel elements of FuRPE not only serve to further refine the model but also to reduce potential biases that may arise from inaccuracies in pseudo labels, thereby optimizing the network’s training process and enhancing the robustness of the model. We apply FuRPE to train both two-stage and fully convolutional single-stage full-body reconstruction networks. Our exhaustive experiments on numerous benchmark datasets illustrate a substantial performance boost over existing methods, underscoring FuRPE’s potential to reshape the state-of-the-art in full-body reconstruction.

Introduction

Interpreting human behavior and appearance from real-world imagery and video data is a crucial prerequisite for a multitude of applications, with robotics (Dupont et al. 2021; Anwar et al. 2019; Ortenzi et al. 2021) and augmented reality (Fan et al. 2021; Xiong et al. 2021; Siriwardhana et al. 2021) standing out as prime examples. The task of human body reconstruction, a critical yet challenging aspect within computer vision, substantially contributes to this understanding. It involves estimating a mesh from a given image or video that aptly represents the target human’s pose and appearance, serving as a cornerstone in applications that hinge on digital human representations.

The landscape of human body reconstruction is enriched with diverse methodologies (Kolotouros et al. 2019; Kocabas, Athanasiou, and Black 2020; Choi et al. 2021; Rong, Liu, and Loy 2022; Zhang et al. 2021). These can be broadly bracketed into two categories: methods that predict the mesh topology directly from input images, and those that estimate the parameters of the parametric human body model SMPL

*Corresponding author: hejun@ruc.edu.cn

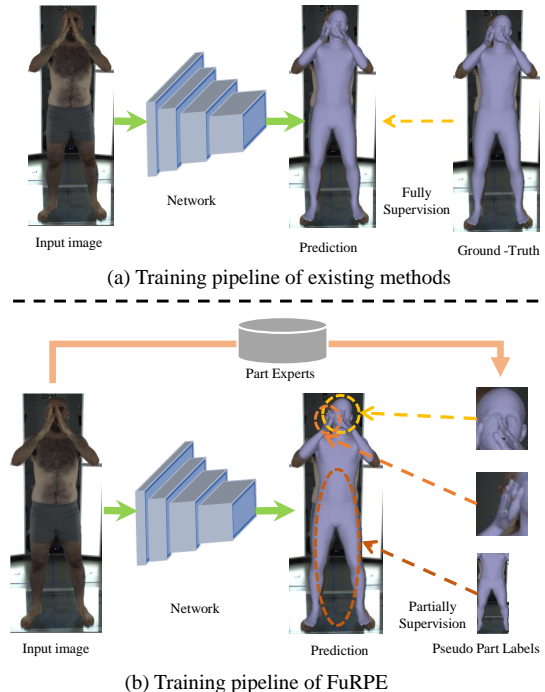


Figure 1: (a) Traditional methods train on costly, scarce annotated data. (b) Our method utilizes affordable, high-quality pseudo labels from part-experts.

(Loper et al. 2015). Despite significant strides, these methods have a common limitation - they primarily focus on reconstructing only the body parts, thereby constraining their broader applicability. To address this limitation, several recent works (Choutas et al. 2020; Rong, Shiratori, and Joo 2021; Moon and Lee 2020; Zhou et al. 2021; Lin et al. 2023) have embarked on full-body reconstruction methods. These methodologies harness the power of the SMPL-X model (Pavlakos et al. 2019), which incorporates the head and hands into SMPL’s representation. However, these methods encounter a significant challenge: the scarcity of adequately annotated data for training. This challenge originates from the intricacies associated with annotating full-body parameters from an image. Conventional methods involve using optical or inertial tools for recording the body’s pose. However, capturing the head and hands requires advanced, more

expensive equipment like the Vrtrix, adding a layer of financial burden and complexity to the process. Furthermore, harmonizing the data from these three distinct captures into a cohesive representation is a significant challenge.

In this paper, we present FuRPE, a novel method designed to tackle the complexities of annotating full-body parameters from an image. FuRPE harnesses the potential of the SMPL-X model, which can be decomposed into three sub-parametric models: SMPL (Loper et al. 2015), FLAME (Li et al. 2017), and MANO (Romero, Tzionas, and Black 2022). Several part experts (Kolotouros et al. 2019; Kocabas, Athanasiou, and Black 2020; Choi et al. 2021; Ge et al. 2019; Zhou et al. 2020; Feng et al. 2021b; Daněček, Black, and Bolkart 2022) have shown promising results for these individual components. As illustrated in Fig. 1, adopting the concept from (Weinzaepfel et al. 2020), FuRPE utilizes these experts to generate pseudo full-body ground-truth labels, enabling us to exploit large-scale datasets for training a full-body reconstruction model. While (Weinzaepfel et al. 2020) predicts only the 3D keypoints of the body, FuRPE extends this by predicting SMPL-X parameters with improved pseudo labels and carefully crafted modules. Although this makes our task more challenging, it is also more practical, considering the full-body shape and appearance, and enabling a broader range of applications.

Specifically, FuRPE ingeniously leverages the pseudo labels produced by part experts. This approach significantly augments the valuable training data available, leading to an enhancement in the overall performance of the system. To ensure the quality of these pseudo labels, a simple yet highly effective pseudo ground-truth selection scheme is employed. This scheme plays a pivotal role in refining the training process and mitigating potential bias that could be introduced by inaccurate pseudo labels. Alongside the pseudo ground-truth selection scheme, FuRPE adopts an Exponential Moving Average (EMA) training strategy (Klinker 2011) and an expert-derived feature distillation strategy. These strategies are designed to refine the training outcomes and further curtail the bias induced by inaccurate pseudo labels. The EMA training strategy is especially noteworthy as it promotes a self-supervised joint training process between student-teacher networks, which significantly bolsters the training performance. Crucially, at the inference stage, FuRPE only requires the student model to predict the full-body SMPL-X parameter.

We apply FuRPE to train both two-stage and fully convolutional single-stage full-body reconstruction networks, using a variety of well-established human body reconstruction datasets. This application showcases the method’s adaptability and robustness across different network structures. The empirical results from these comprehensive experiments significantly underscore FuRPE’s superiority over the baseline models, thereby establishing a new benchmark in the field of full-body reconstruction performance. A consistent improvement in the performance of our model was observed as the size of the training dataset increased. This trend effectively demonstrates FuRPE’s capability to utilize existing publicly available datasets in a highly efficient manner.

Our contributions can be summarized as follows: 1) We

present FuRPE, an innovative method the full-body reconstruction task, which notably amplifies the performance of both two-stage and one-stage strong baselines. 2) We propose to harness knowledge from part-experts for full-body reconstruction, complemented by a simple yet effective pseudo ground-truth selection scheme. In addition, we incorporate an Exponential Moving Average training strategy and introduce an expert-derived feature distillation strategy to reduce bias exist in pseudo labels. 3) We conduct extensive experiments on several publicly available datasets, demonstrating the effectiveness of using pseudo labels and expert-derived features for full-body reconstruction.

Related Works

3D Human Pose Estimation

3D human pose estimation is closely related to body mesh reconstruction, as both tasks involve predicting the 3D keypoints of a person from an image. There are two main categories of 3D human pose estimation methods: direct prediction methods and 2D-3D lifting methods. Direct prediction methods (Moon, Chang, and Lee 2019; Mehta et al. 2018; Pavlakos et al. 2017; Rogez, Weinzaepfel, and Schmid 2019) directly predict the 3D keypoints from the input image. For example, LCR-Net++ (Rogez, Weinzaepfel, and Schmid 2019) predicts 2D and 3D poses of multiple people simultaneously to increase the robustness of 3D keypoints estimation from a fully-connected prediction branch. On the other hand, 2D-3D lifting methods (Martinez et al. 2017; Pavllo et al. 2019; Zhan et al. 2022; Li et al. 2022a) first estimate the 2D keypoints of the human and then use a lifting network to lift these keypoints to 3D keypoints. For instance, (Li et al. 2022b) proposes to predict several hypotheses from 2D keypoints for 3D keypoints prediction. Despite significant progress, 3D human pose estimation methods only estimate the keypoints of a human and neglect the prediction of shape and appearance. In contrast, our proposed method reconstructs the parameters of the SMPL-X model, which represent a human’s pose and shape parameters. This allows for a more comprehensive representation of the human body, enabling a wider range of applications.

3D Human Body Reconstruction

3D human body reconstruction aims to predict the mesh of a target human from a single image or video. HMR (Kanazawa et al. 2018) recovers SMPL parameters from a single image using a model trained on pseudo labels generated from SMPLify (Loper et al. 2015). SPIN (Kolotouros et al. 2019) proposes to reconstruct 3D human pose and shape via model-fitting in the loop, making the training pipeline self-supervised. Both TCMR (Choi et al. 2021) and VIBE (Kocabas, Athanasiou, and Black 2020) introduce the use of temporal information for 3D human body reconstruction to improve stability and smoothness of the prediction results. Many subsequent works (Lin, Wang, and Liu 2021; Rempe et al. 2021; Kocabas et al. 2021) have been proposed to improve the performance of SMPL parameter estimation. Some methods (Alldieck, Xu, and Sminchisescu 2021) also estimate the parameters of other parametric models such as

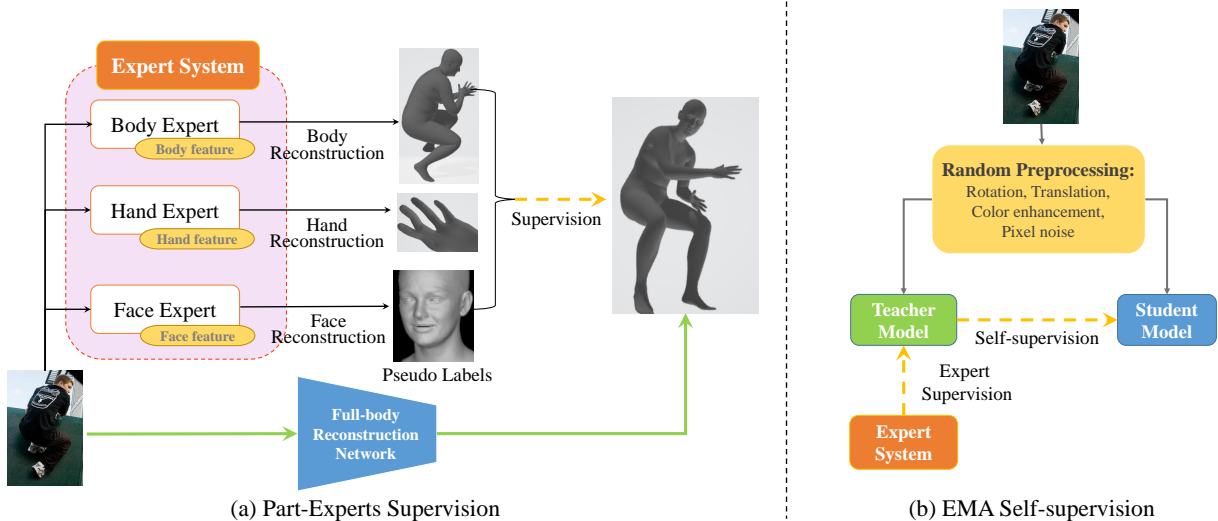


Figure 2: Pipeline of our work. (a) The training pipeline of using part-experts to generate supervision signals. (b) The training pipeline of Exponential Moving Average self-supervision.

GHUM (Xu et al. 2020). However, these methods only predict the mesh of body parts, limiting their broader application. In contrast, our paper focuses on the more general task of full-body reconstruction.

Full-body Reconstruction

Building upon SMPL-X’s introduction (Pavlakos et al. 2019), several strategies offer full-body reconstruction. SMPLify-X, an optimization-based method, suffers from impractical speed. ExPose (Choutas et al. 2020) utilizes a neural network for direct full-body SMPL-X parameters prediction from a single image and introduces a body-attention mechanism for interrelations among head, body, and hands. FrankMocap (Rong, Shiratori, and Joo 2021) suggests a strategy that stitches together different body parts to form whole-body parameters. Concurrently, PIXIE (Feng et al. 2021a) merges features from different part modules in the network through a moderator, allowing all parts to contribute to the whole-body reconstruction. Despite acceptable performance, these methods face limitations due to the scarcity of well-annotated training data. OSX (Lin et al. 2023) introduces an optimization-based method for whole-body parameters annotation, but it remains computationally expensive. Contrarily, our method efficiently generates pseudo labels and features from part experts and leverages large-scale public datasets for training, significantly enhancing full-body reconstruction models’ performance. Our work resembles (Weinzaepfel et al. 2020), which also use pseudo labels generation to predict 3D human body parameters. However, their approach predicts only the full-body 3D keypoints, while ours estimates both human pose and shape and employs an Exponential Moving Average strategy for further performance improvement.

Method

Overview

In this section, we delve into the comprehensive description of our proposed methodology, FuRPE. The process flow of our method is illustrated in Fig. 2. As depicted in Fig. 2 (a), the process commences with the generation of pseudo labels and pseudo features. We denote the pseudo labels and features for the face, hand, and body as (L_{pre}^f, F^f) , (L_{pre}^h, F^h) , and (L_{pre}^b, F^b) respectively before the selection process.

To ensure high-quality training data, we implement an elaborate pseudo ground-truth selection scheme, which refines these initial pseudo labels to (L^f, F^f) , (L^h, F^h) , and (L^b, F^b) , where the subscript indicates the corresponding body part (face: ‘f’, hand: ‘h’, body: ‘b’).

As shown in Fig. 2 (b), we introduce an Exponential Moving Average (EMA) self-supervision pipeline to further enhance performance. At the onset of training, we simultaneously instantiate a student network and a teacher network. We denote an input image as I , which is randomly augmented into two images, I_a and I_b . I_a is processed by the teacher network, and I_b is processed by the student network. The outputs and weights of the two networks should align according to the data augmentation method. To maintain this consistency, we adopt an EMA training strategy to update the weights of the networks, enabling them to be trained jointly. During inference, only the student network is utilized to predict the full-body parameters.

Part Experts and Pseudo labels

The task of integrating part-specific annotations from diverse devices into a unified full-body annotation is addressed by leveraging the separability of the SMPL-X model. It is segregated into the FLAME model for the face, SMPL model for the body, and the MANO model for the hands. We utilize the predictions of part-expert models, trained on part-specific annotated data, as pseudo part labels. Accordingly, we introduce three part-experts in our work, each ded-

icated to a specific part: face, body, and hands. These experts - SPIN, DECA, and FrankMocap - are well-trained deep learning models that perform robust part-specific predictions from respective cropped images.

Body Expert: We utilize the SMPL model as our body expert. The pose parameters $\theta_{body} \in R^{72}$ and shape parameters $\beta_{body} \in R^{10}$ of the body are estimated from the cropped body image I_{body} with the SPIN model, which has been well-trained on a large dataset comprising diverse body shapes and poses.

$$\theta_{body}, \beta_{body} = f_{SPIN}(I_{body}) \quad (1)$$

Face Expert: The FLAME model is employed for face reconstruction. The parameters for facial identity $\beta_{face} \in R^{200}$, pose $\theta_{face} \in R^{3k+3}$ (where $k = 4$ joints: neck, jaw, and eyeballs), and expression $\psi_{face} \in R^{100}$, are estimated from the cropped face image I_{face} with the DECA model, which has been well-trained on a large dataset comprising diverse facial expressions and identities.

$$\beta_{face}, \theta_{face}, \psi_{face} = f_{DECA}(I_{face}) \quad (2)$$

Hand Expert: The MANO model is used for hand reconstruction. The pose parameters $\theta_{hand} \in R^{21 \times 3}$ and shape parameters $\beta_{hand} \in R^{10}$ of the hands are estimated from the cropped hand image I_{hand} with the FrankMocap model, which has been well-trained on a large dataset comprising diverse hand poses and shapes.

$$\beta_{hand}, \theta_{hand} = f_{FrankMocap}(I_{hand}) \quad (3)$$

Pseudo Ground-truth Selection Scheme

Though the three part-experts can generate expressive pseudo labels and pseudo features, we can't guarantee that every pseudo ground-truth is in high quality. To filter our low quality pseudo ground-truths, we propose a three-step pseudo label select scheme.

Step 1: We leverage Openpose to detect the 2D key points of each person in the image. Then, we count the number of high confident keypoints. When the number of high confident keypoints is smaller than 12, we discard this image. The confidence thresholds are 0.1, 0.2, 0.4 for body, hand, and face respectively. This process can be expressed as below:

$$K_{2D} = \text{OpenPose}(I), n_{conf} = \text{count}(K_{2D}), \quad (4)$$

$$\text{discard } I \text{ if } n_{conf} < 12. \quad (5)$$

Step 2: The left images are input into our part-experts for parameters estimation. We discard these images with invalid output. This process can be expressed as below:

$$\theta, \beta = \text{PartExperts}(I), \quad (6)$$

$$\text{discard } I \text{ if output is invalid.} \quad (7)$$

Step 3: We use the predicted part parameters to drive the SMPL-X model and get the 3D keypoints. These 3D keypoints are projected into 2D keypoints. Then, we calculate the mean square error (MSE) between these projected 2D keypoints and 2D keypoints detected by OpenPose. If the MSE is larger than 1.5cm, we discard this image. This process can be expressed as below:

$$K_{3D} = \text{SMPLX}(\theta, \beta), K_{proj} = \text{Project}(K_{3D}), \quad (8)$$

$$\text{MSE} = \text{mean}((K_{proj} - K_{2D})^2), \quad (9)$$

$$\text{discard } I \text{ if } \text{MSE} > 1.5. \quad (10)$$

The final left images and part-parameters are used for the full-body reconstruction model training.

Training using Pseudo Labels

To extract knowledge from the pseudo labels generated by part experts, we compute three distinct loss types: body loss, head loss, and hand loss to train a full-body reconstruction network.

Body loss is composed of the Mean Absolute Error (MAE) between 2D body joints and the Mean Square Error (MSE) between body poses. The head and hand losses follow a similar structure, but an additional expression loss is included in the head loss.

The overall training loss can be expressed as follows:

$$\begin{aligned} L_{total} &= L_{body} + L_{face} + L_{hand}, \\ L_{body} &= L_{2d-body-joint} + L_{pose}, \\ L_{face} &= L_{2d-face-joint} + L_{expression} + L_{jaw-pose}, \\ L_{hand} &= L_{2d-hand-joint} + L_{hand-pose}. \end{aligned}$$

In the above loss function, the 2D joint loss ($L_{2d-body-joint}$, $L_{2d-face-joint}$, and $L_{2d-hand-joint}$) is calculated as:

$$L_{2d-joint} = \sum_{j=1}^J v_j \|\hat{x}_j - x_j\|_1, \quad (11)$$

where v_j is the binary variable representing the visibility of the j th joint, \hat{x}_j refers to the ground truth value, and x_j is the predicted value.

The pose loss can be calculated as:

$$L_{pose} = \left\| \hat{\theta} - \theta \right\|_2^2, \quad (12)$$

where $\hat{\theta}$ refers to the ground truth value, and θ is the predicted value.

The expression loss follows the same function, with ϕ representing ground truth expression parameters:

$$L_{expression} = \left\| \hat{\phi} - \phi \right\|_2^2. \quad (13)$$

Table 1: Comparison on EFH dataset.

Category	Methods	V2V/Procrustes	PA-V2V Body	PA-V2V L-Hand	PA-V2V R-Hand	PA-V2V Face
Baseline	Expose	55.1	52.9	13.1	12.6	5.7
SOTA Methods	Frankmocap(copy & paste)	58.2	52.7	10.9	11.2	5.7
	Frankmocap(optimization)	54.7	53.3	10.8	11.2	5.4
	Frankmocap(neural network)	57.1	53.2	10.9	11.2	5.7
	PIXLE	55.0	53.0	11.0	11.2	4.6
Ours	FuPRE+Expose	50.6	51.6	12.4	11.7	5.1
	FuPRE+ResNet	54.8	52.9	12.4	12.1	5.6

Expert-Derived Feature Distillation Strategy

Apart from the pseudo labels, we also employ pseudo features generated by part experts for training our full-body reconstruction model, as depicted in Fig. 2. This strategy is not because these expert-derived features are inherently superior, but because they can supplement the information that may be lost in the pseudo labels.

By distilling the knowledge embedded in these expert-derived features, we can effectively guide the training of our model. This approach refines the training outcomes and reduces the bias that inaccurate pseudo labels may induce.

The feature loss for each body part is computed separately. We denote the feature loss for the body part as $L_{b-feature}$. The feature losses for the head and hand parts follow a similar structure and are calculated as follows:

$$L_{feature} = A \cdot KL(\hat{f}, f), \quad (14)$$

where KL represents the Kullback-Leibler divergence, A is an amplification coefficient empirically set as e^5 , \hat{f} is the pseudo feature generated by the part experts, and f is the predicted feature of our model.

Exponential Moving Average Training Strategy

In our training framework, depicted in Fig. 2 (b), we include an Exponential Moving Average (EMA) self-supervision strategy to further refine the training outcomes and diminish the bias brought by imprecise pseudo labels. This approach is inspired by Huang et al. (Huang et al. 2021).

We initialize two networks: a teacher network and a student network. The consistency between their outputs, under different random data augmentations, provides a self-supervision mechanism. The teacher model serves as the regression target for the student model, with their parameters updated by the training loss and an EMA process, respectively.

The parameters of the teacher network (σ) and the student network (ϕ) are updated as follows:

$$\phi \rightarrow \tau\phi + (1 - \tau)\sigma, \quad (15)$$

where $\tau \in [0, 1]$ is the decay rate of the moving average. The EMA loss during training is computed as:

$$L_{EMA} = L_{total} + L_{o \rightarrow t} + L_{t \rightarrow o}, \quad (16)$$

where $L_{o \rightarrow t}$ denotes the student-to-teacher loss, and $L_{t \rightarrow o}$ is its inverse, computed as:

$$L_{a \rightarrow b} = 2 - 2 \cdot \frac{\langle z_a, z_b \rangle}{\|z_a\|_2 \cdot \|z_b\|_2}, \quad (17)$$

with z_a representing the parameters predicted by model a (i.e., the student model).

Experiment

Implementation Details

We implement our method using Pytorch. The model is optimized using the Adam optimizer. We train two models in our implementation. The first one is our baseline method: the ExPose Network, which uses three sub networks for body, head and hand parameters prediction respectively. The second one is set up by ourselves. In this model, only a sub network is used for all the body, head and hand reconstruction. Therefore, this model is more light weight than the ExPose Network, termed as a one stage method. The initial learning rate is 0.00001, then reduces by 10 times after training for 20 epochs. The batch size is set as 32. All the experiments are conducted on a single A6000 GPU.

Dataset

Training Data. The training datasets that we used include CuratedFits (Choutas et al. 2020), H3.6M(Ionescu et al. 2014), MPI(Mehta et al. 2017), Ochuman (Zhang et al. 2019), Posetrack (Joo, Neverova, and Vedaldi 2021), and EFT(Joo, Neverova, and Vedaldi 2021). Among them, the first three ones are indoor datasets while the remaining are outdoor images, which consist of diversified pose sources and scenarios. After data selection, H3.6M contains 234041 pictures, MPI contains 133522 pictures, Ochuman contains 3936 pictures, Posetrack contains 22895 pictures, and EFT contains 2074 pictures.

Evaluation Data. In order to compare our model with the state-of-the-art 3D human pose estimation models especially targeting at head, hand and body, together with several recently proposed full-body reconstruction models, we use several widely used public available test sets for experiments. For part-only evaluation, NoW (Feng et al. 2018) is used for the head evaluation, FreiHAND (Zimmermann et al. 2019) is used for hand evaluation, and 3DPW (Martinez et al. 2017) is used for body evaluation. As for full-body evaluation, we use EHF (Pavlakos et al. 2019) to test the comprehensive capture ability of the model and the effectiveness of the self-supervised learning strategy.

Full-body Evaluation

We first evaluate the performance of our method on full-body reconstruction. The experiments are conducted on EFH dataset. We follow (Pavlakos et al. 2019) to use Vertex-To-Vertex/Procrustes, PA-V2V Body, PA-V2V Left Hand,

Table 2: Performance Improvements through Incremental Data Addition.

Methods	V2V/Procrustes	PA-V2V Body	PA-V2V Left Hand	PA-V2V Right Hand	PA-V2V Face
curated fittings	53.3	52.4	13	12.4	5.3
+mpi	52.4	51.7	13	12.7	5.3
+3DPW	52.3	51.7	13	12.8	5.3
+Human3.6m	51.4	51.7	12.7	11.8	5.1
+ochuman	51.1	51.6	12.6	11.8	5.1
+posetrack	50.9	51.6	12.7	11.8	5.1
+EFT	50.6	51.6	12.4	11.7	5.1

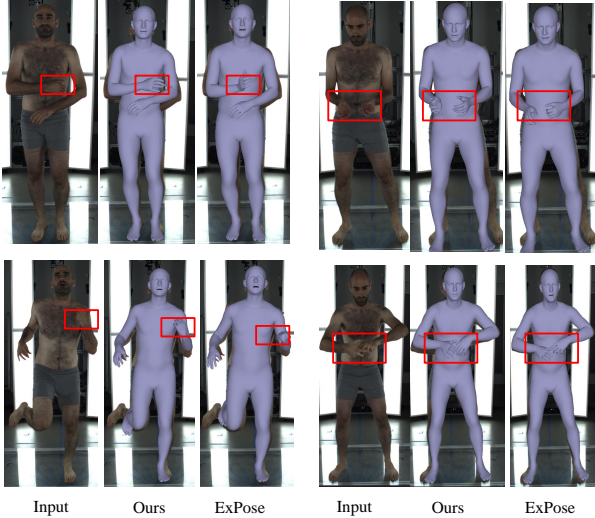


Figure 3: Visualization results on EFH dataset.

Table 3: Body reconstruction evaluation and comparison on 3DPW dataset.

Methods	Procrustes MPJPE (mm)	Pelvis MPJPE
ExPose	60.7	93.4
FrankMocap	60.0	94.3
PIXIE	61.3	-
SPIN	59.2	96.9
Ours+ExPose	59.0	90.2

PA-V2V Right Hand, and PA-V2V Face metrics to evaluate the performance of the models.

Comparison with the baseline: The comparison results between our proposed methods and the baseline model ExPose are presented in Table 1. Several conclusions can be drawn from these results: 1) Our FuPRE significantly improves the performance of ExPose in terms of body, head, and hands reconstruction. Specifically, the Vertex-To-Vertex/Procrustes metric improves from 55.1 to 50.6, demonstrating a substantial advancement in the full-body reconstruction task. 2) The substantial improvement mainly originates from the superior reconstruction of body parts. This suggests that the body part reconstruction benefits most from our FuPRE. This is plausible given that human body parts are the most visible in training images, and thus, the pseudo labels would be the most precise. 3) Both variants of our method, namely FuPRE+ExPose and FuPRE+ResNet, showcase outstanding performance. Particularly, it’s worth noting that our

Table 4: Face reconstruction evaluation and comparison on NoW dataset.

Methods	PA-P2S Median(mm)	mean	std
ExPose	1.38	1.74	1.47
PIXLE	1.18	1.49	1.25
Deep 3D Face	1.23	1.54	1.29
3DDFA	1.23	1.57	1.39
PRNet	1.5	1.98	1.88
DECA	1.10	1.38	1.18
MGCNet	1.31	1.87	2.63
RingNet	1.21	1.54	1.31
Ours+ExPose	1.18	1.50	1.27

FuPRE+ResNet, a one-stage method, achieves superior results compared to the two-stage ExPose, even without the need of specifically cropping out face, hand and body parts for separate processing. This validates the effectiveness of our FuPRE training framework and demonstrates its superiority over traditional two-stage methods.

Comparison with SOTA: We further compare our proposed methods with the state-of-the-art (SOTA) method, mainly with PIXLE, as illustrated in Table 1. 1) For body part reconstruction, both our two-stage method (FuPRE+ExPose) and one-stage method (FuPRE+ResNet) surpass PIXLE. Specifically, our two-stage method achieves a Vertex-To-Vertex/Procrustes metric of 50.6, and our one-stage method reaches 54.8, both outperforming the score of 55.0 attained by PIXLE. This demonstrates the superior performance of our methods in body part reconstruction. 2) For face and hand parts reconstruction, our methods are comparable to PIXLE. Our two-stage method obtains a score of 5.1 on face part and 12.4 on hand parts, while the one-stage method achieves 5.6 on face part and 12.4 on hand parts. Compared to PIXLE’s scores of 4.6 on face part and 11.0 on hand parts, our methods demonstrate competitive performance. These results confirm the robustness and effectiveness of our proposed methods, both in two-stage and one-stage scenarios. While we outperform PIXLE in body part reconstruction, we maintain competitive performance for face and hand parts. This demonstrates the balance our methods have achieved between performance and generalization across different body parts.

Qualitative comparison: Fig. 3 presents a visual comparison between our two-stage method and ExPose. A significant observation from this comparison is the pronounced accuracy of our method in predicting wrist poses. This can be attributed to the extensive use of pseudo labels in our network training process, enabling our models to learn and

Table 5: Hand reconstruction evaluation and comparison on Freihands dataset.

Methods	PA-MPJPE(mm)	PA-V2V(mm)	F@5mm	F@15mm
ExPose	12.2	11.8	0.48	0.92
MANO CNN	11.0	10.9	0.52	0.93
PIXIE	12.9	12.1	0.47	0.92
Ours+ExPose	12.6	12.2	0.46	0.92

Table 6: Ablation study on EFH dataset

Methods	V2V/Procrustes	PA-V2V Body	PA-V2V Left Hand	PA-V2V Right Hand	PA-V2V Face
ExPose	55.1	52.9	13.1	12.6	5.7
+ Psudeo ground-truth	53.8	52.5	12.8	12.7	5.3
+ Psudeo GT Selection	52.9	52.5	12.7	12.9	5.3
+ EMA(Ours+ExPose)	50.6	51.6	12.4	11.7	5.1

represent the pose distribution more effectively compared to ExPose. Moreover, our method exhibits superior performance in head reconstruction. For instance, in the bottom-right images of Fig. 3, the subject’s mouth is accurately depicted as closed by our method, whereas ExPose erroneously predicts it as open. This highlights the enhanced precision of our method in interpreting and reconstructing complex facial expressions. In essence, our method benefits substantially from the “learning from experts” pipeline, leading to superior performance in full-body reconstruction. This not only validates the effectiveness of our approach but also underscores the potential of utilizing pseudo labels to improve the learning process and the final reconstruction results.

Performance Improvements through Incremental Data

Addition: The underlying premise of FuRPE is the effective utilization of pseudo labels and expert-derived features, generated by part-experts, to facilitate model training. This presents an interesting proposition: as the volume of data and pseudo ground-truths increases, the model’s performance should correspondingly improve. To empirically validate this hypothesis, we incrementally augment the training set with additional data and observe the ensuing changes in performance. The results, as shown in Table 2, lend credence to our hypothesis, unambiguously demonstrating that an increase in the volume of training data corresponds to enhanced model performance. This can be attributed to the model’s ability to learn more generalized features and make increasingly accurate pose and shape predictions with a larger dataset.

Part-only Evaluation

Body Part Evaluation: We assess the body-part reconstruction performance on the 3DPW dataset using Procrustes MPJPE and Pelvis MPJPE metrics, as shown in Table 3. Our method not only outperforms state-of-the-art full-body models but also excels over the body-centric SPIN method. This substantiates our method’s superior body reconstruction capabilities, largely fostered by the use of pseudo labels and expert-derived features, thereby significantly improving upon the baseline.

Head Part Evaluation: We scrutinize the head-part reconstruction performance on the NoW dataset, utilizing the PA-P2S Median(mm), mean, and standard deviation metrics. The empirical results, presented in Table 4, reveal our

method’s superiority over all full-body reconstruction techniques, whilst achieving performance on par with specialized head-only reconstruction methodologies. Notably, our results closely align with the current state-of-the-art DECA method. This empirical evidence underscores the promising nature of our head-reconstruction results.

Hand Part Evaluation: Table 5 presents the hand-part reconstruction performance using PA-MPJPE(mm), PA-V2V(mm), F@5mm, and F@15mm metrics. All methods, including ours, exhibit similar performance, likely due to the low-resolution of training images limiting the model’s learning capacity for hand information. Future efforts should focus on incorporating high-resolution images into the training and testing sets to enhance hand reconstruction performance and benchmarking.

Ablation Study

In Table 6, we dissect the performance increments attributed to each key component added into our baseline network, ExPose. The implementation of the pseudo ground-truth generation module (including both pseudo labels and expert-derived features) reduces the overall V2V/Procrustes from 55.1 to 53.8, indicating the significant role of high-quality training data. Further inclusion of the pseudo ground-truth selection module refines the performance to 52.9, verifying its effectiveness in providing superior pseudo ground-truth.

The introduction of the MEA training scheme, denoted as “+ EMA (Ours+ExPose)” in the table, brings about the most dramatic performance enhancement, reducing the V2V/Procrustes to 50.6, the lowest among all configurations. This performance leap underscores the MEA scheme’s capability in imposing strong constraints through a self-supervised learning framework.

Comparing “Ours+ExPose” with other variants, it’s evident that the integration of all components results in the most pronounced performance improvement across all metrics, including PA-V2V for Body, Left Hand, Right Hand, and Face. This observation affirms the collective indispensability of all components, underscoring the superiority of our comprehensive approach.

Conclusion

In this paper, we present FuRPE, a framework that tackles the scarcity of annotated data in full-body reconstruction.

By utilizing part-experts and a pseudo ground-truth selection scheme, FuRPE generates high-quality pseudo labels, enhancing the learning process. Our novel training strategies further optimize the model’s robustness. FuRPE significantly surpasses existing methods on multiple benchmarks, demonstrating its potential to redefine the state-of-the-art.

References

- Alldieck, T.; Xu, H.; and Sminchisescu, C. 2021. imghum: Implicit generative models of 3d human shape and articulated pose. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5461–5470.
- Anwar, S.; Bascou, N. A.; Menekse, M.; and Kardgar, A. 2019. A systematic review of studies on educational robotics. *Journal of Pre-College Engineering Education Research (J-PEER)*, 9(2): 2.
- Choi, H.; Moon, G.; Chang, J. Y.; and Lee, K. M. 2021. Beyond static features for temporally consistent 3d human pose and shape from a video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1964–1973.
- Choutas, V.; Pavlakos, G.; Bolkart, T.; Tzionas, D.; and Black, M. J. 2020. Monocular expressive body regression through body-driven attention. In *European Conference on Computer Vision*, 20–40. Springer.
- Daněček, R.; Black, M. J.; and Bolkart, T. 2022. EMOCA: Emotion Driven Monocular Face Capture and Animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20311–20322.
- Dupont, P. E.; Nelson, B. J.; Goldfarb, M.; Hannaford, B.; Menciassi, A.; O’Malley, M. K.; Simaan, N.; Valdastrì, P.; and Yang, G.-Z. 2021. A decade retrospective of medical robotics research from 2010 to 2020. *Science Robotics*, 6(60): eabi8017.
- Fan, Z.; Zhu, Y.; He, Y.; Sun, Q.; Liu, H.; and He, J. 2021. Deep learning on monocular object pose detection and tracking: A comprehensive overview. *ACM Computing Surveys (CSUR)*.
- Feng, Y.; Choutas, V.; Bolkart, T.; Tzionas, D.; and Black, M. J. 2021a. Collaborative regression of expressive bodies using moderation. In *2021 International Conference on 3D Vision (3DV)*, 792–804. IEEE.
- Feng, Y.; Feng, H.; Black, M. J.; and Bolkart, T. 2021b. Learning an animatable detailed 3D face model from in-the-wild images. *ACM Transactions on Graphics (ToG)*, 40(4): 1–13.
- Feng, Z.-H.; Huber, P.; Kittler, J.; Hancock, P.; Wu, X.-J.; Zhao, Q.; Koppen, P.; and Raetsch, M. 2018. Evaluation of Dense 3D Reconstruction from 2D Face Images in the Wild. In *2018 13th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2018)*, 780–786.
- Ge, L.; Ren, Z.; Li, Y.; Xue, Z.; Wang, Y.; Cai, J.; and Yuan, J. 2019. 3d hand shape and pose estimation from a single rgb image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10833–10842.
- Huang, S.; Xie, Y.; Zhu, S.-C.; and Zhu, Y. 2021. Spatio-temporal Self-Supervised Representation Learning for 3D Point Clouds. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 6515–6525.
- Ionescu, C.; Papava, D.; Olaru, V.; and Sminchisescu, C. 2014. 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.
- Joo, H.; Neverova, N.; and Vedaldi, A. 2021. Exemplar Fine-Tuning for 3D Human Model Fitting Towards In-the-Wild 3D Human Pose Estimation. In *2021 International Conference on 3D Vision (3DV)*, 42–52.
- Kanazawa, A.; Black, M. J.; Jacobs, D. W.; and Malik, J. 2018. End-to-end recovery of human shape and pose. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7122–7131.
- Klinker, F. 2011. Exponential moving average versus moving exponential average. *Mathematische Semesterberichte*, 58(1): 97–107.
- Kocabas, M.; Athanasiou, N.; and Black, M. J. 2020. Vibe: Video inference for human body pose and shape estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5253–5263.
- Kocabas, M.; Huang, C.-H. P.; Hilliges, O.; and Black, M. J. 2021. PARE: Part attention regressor for 3D human body estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 11127–11137.
- Kolotouros, N.; Pavlakos, G.; Black, M. J.; and Daniilidis, K. 2019. Learning to reconstruct 3D human pose and shape via model-fitting in the loop. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2252–2261.
- Li, T.; Bolkart, T.; Black, M. J.; Li, H.; and Romero, J. 2017. Learning a model of facial shape and expression from 4D scans. *ACM Trans. Graph.*, 36(6): 194–1.
- Li, W.; Liu, H.; Ding, R.; Liu, M.; Wang, P.; and Yang, W. 2022a. Exploiting temporal contexts with strided transformer for 3d human pose estimation. *IEEE Transactions on Multimedia*.
- Li, W.; Liu, H.; Tang, H.; Wang, P.; and Van Gool, L. 2022b. Mhformer: Multi-hypothesis transformer for 3d human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13147–13156.
- Lin, J.; Zeng, A.; Wang, H.; Zhang, L.; and Li, Y. 2023. One-Stage 3D Whole-Body Mesh Recovery with Component Aware Transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 21159–21168.
- Lin, K.; Wang, L.; and Liu, Z. 2021. End-to-end human pose and mesh reconstruction with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1954–1963.
- Loper, M.; Mahmood, N.; Romero, J.; Pons-Moll, G.; and Black, M. J. 2015. SMPL: A skinned multi-person linear model. *ACM transactions on graphics (TOG)*, 34(6): 1–16.

- Martinez, J.; Hossain, R.; Romero, J.; and Little, J. J. 2017. A simple yet effective baseline for 3d human pose estimation. In *Proceedings of the IEEE international conference on computer vision*, 2640–2649.
- Mehta, D.; Rhodin, H.; Casas, D.; Fua, P.; Sotnychenko, O.; Xu, W.; and Theobalt, C. 2017. Monocular 3D Human Pose Estimation in the Wild Using Improved CNN Supervision. In *2017 International Conference on 3D Vision (3DV)*, 506–516.
- Mehta, D.; Sotnychenko, O.; Mueller, F.; Xu, W.; Sridhar, S.; Pons-Moll, G.; and Theobalt, C. 2018. Single-shot multi-person 3d pose estimation from monocular rgb. In *2018 International Conference on 3D Vision (3DV)*, 120–130. IEEE.
- Moon, G.; Chang, J. Y.; and Lee, K. M. 2019. Camera distance-aware top-down approach for 3d multi-person pose estimation from a single rgb image. In *Proceedings of the IEEE/CVF international conference on computer vision*, 10133–10142.
- Moon, G.; and Lee, K. M. 2020. Pose2pose: 3d positional pose-guided 3d rotational pose prediction for expressive 3d human pose and mesh estimation. *arXiv preprint arXiv:2011.11534*.
- Ortenzi, V.; Cosgun, A.; Pardi, T.; Chan, W. P.; Croft, E.; and Kulić, D. 2021. Object handovers: a review for robotics. *IEEE Transactions on Robotics*, 37(6): 1855–1873.
- Pavlakos, G.; Choutas, V.; Ghorbani, N.; Bolkart, T.; Osman, A. A.; Tzionas, D.; and Black, M. J. 2019. Expressive body capture: 3d hands, face, and body from a single image. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10975–10985.
- Pavlakos, G.; Zhou, X.; Derpanis, K. G.; and Daniilidis, K. 2017. Coarse-to-fine volumetric prediction for single-image 3D human pose. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7025–7034.
- Pavlo, D.; Feichtenhofer, C.; Grangier, D.; and Auli, M. 2019. 3d human pose estimation in video with temporal convolutions and semi-supervised training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7753–7762.
- Rempe, D.; Birdal, T.; Hertzmann, A.; Yang, J.; Sridhar, S.; and Guibas, L. J. 2021. Humor: 3d human motion model for robust pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 11488–11499.
- Rogez, G.; Weinzaepfel, P.; and Schmid, C. 2019. Lcrnet++: Multi-person 2d and 3d pose detection in natural images. *IEEE transactions on pattern analysis and machine intelligence*, 42(5): 1146–1161.
- Romero, J.; Tzionas, D.; and Black, M. J. 2022. Embodied hands: Modeling and capturing hands and bodies together. *arXiv preprint arXiv:2201.02610*.
- Rong, Y.; Liu, Z.; and Loy, C. C. 2022. Chasing the tail in monocular 3d human reconstruction with prototype memory. *IEEE Transactions on Image Processing*, 31: 2907–2919.
- Rong, Y.; Shiratori, T.; and Joo, H. 2021. Frankmocap: A monocular 3d whole-body pose estimation system via regression and integration. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1749–1759.
- Siriwardhana, Y.; Porambage, P.; Liyanage, M.; and Ylianttila, M. 2021. A survey on mobile augmented reality with 5G mobile edge computing: architectures, applications, and technical aspects. *IEEE Communications Surveys & Tutorials*, 23(2): 1160–1192.
- Weinzaepfel, P.; Brégier, R.; Combaluzier, H.; Leroy, V.; and Rogez, G. 2020. Dope: Distillation of part experts for whole-body 3d pose estimation in the wild. In *European Conference on Computer Vision*, 380–397. Springer.
- Xiong, J.; Hsiang, E.-L.; He, Z.; Zhan, T.; and Wu, S.-T. 2021. Augmented reality and virtual reality displays: emerging technologies and future perspectives. *Light: Science & Applications*, 10(1): 1–30.
- Xu, H.; Bazavan, E. G.; Zangir, A.; Freeman, W. T.; Sukthankar, R.; and Sminchisescu, C. 2020. Ghum & ghuml: Generative 3d human shape and articulated pose models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6184–6193.
- Zhan, Y.; Li, F.; Weng, R.; and Choi, W. 2022. Ray3D: ray-based 3D human pose estimation for monocular absolute 3D localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13116–13125.
- Zhang, H.; Tian, Y.; Zhou, X.; Ouyang, W.; Liu, Y.; Wang, L.; and Sun, Z. 2021. Pymaf: 3d human pose and shape regression with pyramidal mesh alignment feedback loop. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 11446–11456.
- Zhang, S.-H.; Li, R.; Dong, X.; Rosin, P.; Cai, Z.; Han, X.; Yang, D.; Huang, H.; and Hu, S.-M. 2019. Pose2Seg: Detection Free Human Instance Segmentation. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 889–898.
- Zhou, Y.; Habermann, M.; Habibie, I.; Tewari, A.; Theobalt, C.; and Xu, F. 2021. Monocular real-time full body capture with inter-part correlations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4811–4822.
- Zhou, Y.; Habermann, M.; Xu, W.; Habibie, I.; Theobalt, C.; and Xu, F. 2020. Monocular real-time hand shape and motion capture using multi-modal data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5346–5355.
- Zimmermann, C.; Ceylan, D.; Yang, J.; Russell, B.; Argus, M. J.; and Brox, T. 2019. FreiHAND: A Dataset for Markerless Capture of Hand Pose and Shape From Single RGB Images. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 813–822.